

Global structure of the time delay likelihood

Namu Kroupa^{1,2,3,*} and Will Handley^{2,4}¹*Astrophysics Group, Cavendish Laboratory, J.J. Thomson Avenue, Cambridge CB3 0HE, United Kingdom*²*Kavli Institute for Cosmology, Madingley Road, Cambridge CB3 0HA, United Kingdom*³*Engineering Laboratory, University of Cambridge, Cambridge CB2 1PZ, United Kingdom*⁴*Institute of Astronomy, University of Cambridge, Madingley Road, Cambridge CB3 0HA, United Kingdom*

(Received 13 September 2025; accepted 11 April 2026; published 18 May 2026)

We identify a fundamental pathology in the likelihood for time delay inference that challenges standard inference methods. By analyzing the likelihood for time delay inference with Gaussian process light curve models, we show that it generically develops a boundary-driven “W” shape with a global maximum at the true delay and gradual rises toward the edges of the observation window. This arises because time delay estimation is intrinsically extrapolative. In practice, global samplers such as nested sampling are steered toward spurious edge modes unless strict convergence criteria are adopted. We demonstrate this with simulations and show that the effect strengthens with higher data density over a fixed time span. To ensure convergence, we provide concrete guidance, notably increasing the number of live points. Further, we show that methods implicitly favoring small delays, for example, optimizers and local Markov chain Monte Carlo, induce a bias toward larger H_0 . Our results clarify failure modes and offer practical remedies for robust fully Bayesian time delay inference.

DOI: [10.1103/kj4s-nnyx](https://doi.org/10.1103/kj4s-nnyx)

I. INTRODUCTION

Time delay measurements constitute a primary observable for lensed-quasar time delay cosmography [1–5] and reverberation mapping of active galactic nuclei [6–8]. In both settings, the light curves are noisy, irregularly sampled, and stochastic, yielding likelihood surfaces that are difficult to explore and that will strain standard pipelines as samples grow. The Vera C. Rubin Observatory’s Legacy Survey of Space and Time (LSST) is expected to increase the number of known lensed quasars by roughly 2 orders of magnitude [9,10] and ongoing spectroscopic programmes, including the SDSS-V Black Hole Mapper, are expanding the number of reverberation mapping targets [11,12]. Automated, reliable inference of time delays with quantified uncertainties is therefore needed to control systematic errors in H_0 and to enable robust population studies.

A fully Bayesian framework is well suited to these problems as it represents uncertainty coherently and supports principled model comparison. Obtaining faithful posteriors, however, requires samplers that can navigate specific documented pathologies. Seasonal gaps and finite observing windows induce window-function aliases that generate multiple, well-separated modes in the delay posterior, often at separations comparable to the observing season [13–17]. The likelihood also exhibits strong, nonlinear degeneracies

between the delay and parameters controlling the intrinsic variability. In practice, this necessitates constraints on kernel flexibility in Gaussian process (GP) regression or spline smoothness [18,19]. In reverberation mapping, the convolutional response leads to pronounced parameter degeneracies, so that the recovered lag covaries with the width and shape of the transfer function and with hyperparameters of the continuum variability, and the broad-line region geometry and kinematics exhibit strong covariances and frequent multimodality [20–24]. Historical cases highlight this point. For Q0957+561, early analyses supported distinct delays that a modern global treatment would reveal as multimodality and were ultimately resolved by denser monitoring [25–29]. For B1608+656, degeneracies in the lens mass distribution were broken by combining time delays with high-resolution imaging and stellar kinematics in a unified Bayesian analysis that enabled evidence-based model comparison [30–32].

Methods based on optimization and locally exploring Markov chain Monte Carlo (MCMC) samplers frequently rely on favorable initialization to terminate in the correct mode. In practice, visual matching and restricted reinitializations impose implicit priors that limit exploration [33–36]. Global samplers, including nested sampling, offer an alternative by sampling the prior subject to progressively tighter likelihood constraints. They can, with sufficiently many live points and appropriate proposals, recover well-separated modes and complex degeneracies, and provide the Bayesian evidence for model comparison by construction [37–40]. Although MCMC can estimate evidence through specialized procedures [41–43], nested sampling delivers it natively. Nonetheless, the performance of any global method depends on configuration and diagnostics, and can degrade on highly structured posteriors of the kind encountered in time delay analyses.

*Contact author: nk544@cam.ac.uk

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International](https://creativecommons.org/licenses/by/4.0/) license. Further distribution of this work must maintain attribution to the author(s) and the published article’s title, journal citation, and DOI.

Here, we identify and analyze a further, intrinsic feature of the time delay likelihood that affects both local and global approaches. Under broad and commonly adopted assumptions, namely, stationary GP models for intrinsic variability, the likelihood as a function of delay exhibits a boundary-driven “W”-shaped structure. As the overlap of shifted light curves shrinks near the window edges, inference often reduces to extrapolation, which systematically elevates the likelihood toward the boundaries. We show that while the boundary rise is generic, the central peak remains dominant under typical conditions, so the inference problem is well posed in principle. The practical difficulty is algorithmic. Samplers must be configured to discover and correctly weight all modes in the presence of gradual boundary increases. Through numerical simulations, we demonstrate how this structure can mislead Nested sampling (NS) unless care is taken with live-point budgets and proposal strategies. We quantify the consequence of the common tendency of optimization and local samplers to bias delays to low values, which translates into a positive bias in H_0 .

Beyond astrophysics, time delay estimation is a recurring statistical problem in underwater acoustics [44], wireless communications [45], signal processing [46,47], radar [48], biology [49], geophysics [50], finance [51], and control applications [52]. From the perspective of sampling, NS applied to the calculation of phase diagrams in materials science [53] faces a similar computational challenge. There, NS is used to generate samples from the Boltzmann distribution. One of the primary aims lies in the calculation of the critical temperature of a phase transition. Unbiased estimation thereof relies on proper sampling of an ordered low-entropy phase, corresponding to a narrow deep potential energy basin of, for example, a solid, and a disordered high-entropy phase, corresponding to a wide shallow basin of, for example, a liquid. Since NS is initialized from the high-temperature disordered phase, discovering the solid phase can be computationally challenging [54]. Analogously, in the time delay problem, it is difficult to discover the true central mode while it is easy to sample the wider edge modes.

The paper is organized as follows. Section II presents time delayed GPs, Sec. III describes the toy model, Sec. IV derives analytical results, and Sec. V contains a numerical investigation. Finally, Sec. VIII outlines the conclusions.

II. TIME DELAYED GAUSSIAN PROCESSES

Typically [55], GPs are defined as a collection of random variables, any finite subset of which has a joint Gaussian distribution. We take the alternative definition in weight-space here, in order to make the underlying parametric model explicit. That is, we view a GP with mean function m as linear regression with a possibly infinite number of basis functions ϕ_i ,

$$f(t) = m(t) + \sum_{i=1}^{\infty} w_i \phi_i(t), \quad (1)$$

where w_i are the model parameters. The mean function m and the basis ϕ_i may depend on further (hyper)parameters that we take at this stage as given. Placing a zero-mean Gaussian prior

with diagonal covariance matrix on w_i thus induces a prior in function space.

In the following, we will also be interested in the time delayed function,

$$h(t) = f(t + \Delta t) = m(t + \Delta t) + \sum_{i=1}^{\infty} w_i \phi_i(t + \Delta t), \quad (2)$$

for a given time delay Δt . Since the parameters w_i are shared between the random variables $f(t)$ and $h(t')$, they are jointly distributed as a Gaussian. The mean directly follows from computing $\mathbb{E}[f(t)]$ and $\mathbb{E}[h(t')]$, while the covariance may be obtained similarly. For example, the off-diagonal entries are computed from $\mathbb{E}[(f(t) - \mathbb{E}f(t))(h(t') - \mathbb{E}h(t'))]$. Thus, integrating out the parameters leaves us with the marginal likelihood

$$\begin{aligned} \begin{bmatrix} f(t) \\ h(t') \end{bmatrix} &\sim \mathcal{N} \left(\begin{bmatrix} m(t) \\ m(t' + \Delta t) \end{bmatrix}, \right. \\ &\quad \left. \times \begin{bmatrix} k(t, t) & k(t, t' + \Delta t) \\ k(t + \Delta t, t') & k(t' + \Delta t, t' + \Delta t) \end{bmatrix} \right) \end{aligned} \quad (3)$$

for given t , t' , and Δt . Here, we have defined the kernel

$$k(t, t') = \sum_{i=1}^{\infty} \sum_{j=1}^{\infty} \mathbb{E}[w_i w_j] \phi_i(t) \phi_j(t') \quad (4)$$

$$= \sum_{i=1}^{\infty} \lambda_i^2 \phi_i(t) \phi_i(t'), \quad (5)$$

where λ_i^2 is the prior variance of each parameter. Specific choices of ϕ_i and corresponding priors on w_i return common kernels such as Matérn or squared exponential kernels. The inductive bias of the model arises both from the choice of the basis and the rate of decay of the prior variance λ_i^2 as $i \rightarrow \infty$. Different choices for these result in the function f having different smoothness properties. In practice, certain functional forms of the kernel are invoked that guarantee the existence of such an expansion in basis functions and the kernel hyperparameters then control ϕ_i and λ_i , so that one only has to work with Eq. (3).

As expected, infinitely many parameters render GPs flexible models. In fact, if ϕ_i is a complete basis, it is in principle possible to re-expand the mean function m in this basis and absorb it in the sum in Eq. (1). It is therefore reasonable to ask whether the inclusion of a separate mean function is necessary in the first place.

The necessity arises from noting that, for a finite amount of data and infinitely many basis functions, the regularization imposed by the prior can indeed be very strong, in the sense that the posterior f is completely prior dominated as we move away from a data point. This problem is exacerbated by the fact that the task of inferring the time delayed function h is intrinsically extrapolative. That is, if the time delay Δt is large, h is far away from f along the t axis so that h and f are effectively uncorrelated.

To be more precise, suppose we are given noise-free data $y_1(t_1), \dots, y_1(t_{n_{\text{data}}})$ and $y_2(t_1), \dots, y_2(t_{n_{\text{data}}})$ for f and its time-shifted copy h , respectively, observed at the times $t_1, \dots, t_{n_{\text{data}}}$. To aid interpretation, the posterior $p(\Delta t \mid y_1, y_2)$ of the time

delay may be written as (Appendix A)

$$p(\Delta t | y_1, y_2) \propto \int d\mathbf{w} p(y_2 | \Delta t, \mathbf{w}) p(\mathbf{w} | y_1), \quad (6)$$

where we assumed a uniform prior on Δt and the integration is performed over all parameters $\mathbf{w}$. This states that we first compute the posterior $p(\mathbf{w} | y_1)$ of the parameters given only light curve y_1 . We then draw a set of parameters $\mathbf{w}$ from this posterior, which completely determine the function f . Given the time delay Δt , we now have full information to construct the time-shifted copy h (by simply shifting f along the t axis), which allows us to evaluate the likelihood $p(y_2 | \Delta t, \mathbf{w})$ at the second light curve y_2 . Since the time shift implies that we do not evaluate f at the times $t_1, \dots, t_{n_{\text{data}}}$ but instead at $t_1 + \Delta t, \dots, t_{n_{\text{data}}} + \Delta t$, extrapolation and interpolation are required. The amount of extrapolation depends on Δt .

In general, the extrapolated function averaged over $\mathbf{w}$, which is the posterior mean function, is a linear combination of kernel functions centered on each data point [55]. For a stationary GP, the kernels are localized; i.e., the extrapolation effectively does not depend on the data beyond a certain decorrelation length scale. Instead, the GP returns to the prior, as previously claimed. We will see in Sec. IV that the lack of extrapolative capability of a GP also renders the structure of the time delay posterior [Eq. (6)] pathological, causing additional practical problems.

Because time delay inference necessarily compares shifted light curves beyond their overlap window, extrapolation is unavoidable. This necessitates adding nonstationary structure, for example, via nonstationary kernels or explicit mean functions, so that trends can be represented. Here, we intentionally restrict the investigation to the simplest class, namely, a stationary exponential kernel and a constant mean, so that the likelihood structure and its boundary effects can be isolated and analyzed cleanly. The mechanism we uncover is model agnostic in the sense that it is driven by finite-window overlap rather than the particular functional prior. Extensions to more flexible models are left to future work.

III. TOY MODEL

In this section, we analyze the posterior of the time delay by considering the following synthetic data setting. We consider two light curves, $\mathbf{y}_1$ and $\mathbf{y}_2$, both sampled from a GP with zero mean and an exponential kernel,

$$k(t, t') = A^2 \exp\left(-\frac{|t - t'|}{\ell}\right), \quad (7)$$

as this corresponds to the damped random-walk model of quasar variability [56,57]. The known hyperparameters are $\theta = (A, \ell, \sigma, \Delta t)$, with the amplitude A , length scale ℓ , noise σ , and time delay Δt , respectively. The observation times $\mathbf{t}$ are chosen to be n_{data} equally spaced points between some minimum and maximum observation time, which we set to $T_{\text{min}} = 0$ days and $T_{\text{max}} = 10^3$ days, respectively, as this reflects realistic total observation times. We set $A = 1$ in units of magnitude without loss of generality, as any other value may be reduced to this choice by appropriate rescaling. Furthermore, we set the noise to be negligible in comparison to the

signal, $\sigma = 10^{-2} \ll A$. Thus, the data are a sample from

$$\begin{bmatrix} \mathbf{y}_1 \\ \mathbf{y}_2 \end{bmatrix} \sim \mathcal{N}\left(\mathbf{0}, \begin{bmatrix} \mathbf{K}(\mathbf{t}, \mathbf{t}) + \sigma^2 \mathbf{I} & \mathbf{K}(\mathbf{t}, \mathbf{t} - \Delta t \mathbf{1}) \\ \mathbf{K}(\mathbf{t} - \Delta t \mathbf{1}, \mathbf{t}) & \mathbf{K}(\mathbf{t}, \mathbf{t}) + \sigma^2 \mathbf{I} \end{bmatrix}\right), \quad (8)$$

where the covariance matrix in each block is given by $[\mathbf{K}(\mathbf{t}, \mathbf{t})]_{ij} = k(t_i, t_j)$ and $\mathbf{I}$ is the identity matrix. Given such data, we would in practice sample parameters using the likelihood

$$L(\theta) = \frac{1}{\sqrt{|2\pi\mathbf{K}_\theta|}} \exp\left(-\frac{1}{2}\mathbf{y}^\top \mathbf{K}_\theta^{-1} \mathbf{y}\right), \quad (9)$$

where $\mathbf{y}_1$ and $\mathbf{y}_2$ are concatenated into $\mathbf{y}$ and $\mathbf{K}_\theta$ is the full covariance matrix in Eq. (8), parametrized by the hyperparameters θ . Furthermore, uniform priors are set on all GP hyperparameters. In general, we observe [58] that for real datasets, the bulk of the posterior mass is well contained within the chosen prior ranges, except for the length scale of the exponential kernel, which we attribute to the data (Appendix B).

IV. ANALYTICAL RESULTS

In this section, we use the toy model to investigate the log-likelihood structure of the time delay. The advantage in this setting is that we know the exact data-generating distribution and the values of the true (hyper)parameters. By applying the analytical understanding, we then show that the pathological likelihood structure may in principle be removed by a suitable prior.

A. The log-likelihood structure

In order to make statements about L for any data sample, we consider the data-averaged log-likelihood and its variance, respectively,

$$\mathbb{E}_{\mathbf{y}} \ln L(\theta) = -\frac{1}{2} \text{Tr}(\mathbf{K}_\theta^{-1} \mathbf{K}) - \frac{1}{2} \ln |2\pi \mathbf{K}_\theta|, \quad (10)$$

$$\mathbb{V}_{\mathbf{y}} \ln L(\theta) = \frac{1}{4} \sum_{ij} (\mathbf{K}_\theta^{-1})_{ij} [\mathbf{K}_{ii} \mathbf{K}_{jj} + (\mathbf{K}_{ji})^2], \quad (11)$$

where $\mathbf{K} = \mathbb{E}_{\mathbf{y}} \mathbf{y} \mathbf{y}^\top$ is the true covariance of the data, i.e., the covariance matrix in Eq. (8). The first term is typically interpreted as representing the goodness of fit, whereas the second term penalizes model complexity [55]. Note that this decomposition into goodness of fit and model complexity penalty is distinct from Occam's razor equation, $\ln Z = \langle \ln L \rangle_P - D_{\text{KL}}$ [59], where Z is the Bayesian evidence, $\langle \ln L \rangle_P$ the posterior-averaged log-likelihood, and D_{KL} the Kullback-Leibler divergence.

The top subplot of Fig. 1 shows $\mathbb{E}_{\mathbf{y}} \ln L$. While $\mathbb{E}_{\mathbf{y}} \ln L$ is highly oscillatory, it has a global maximum at the true time delay. Away from this peak, the log-likelihood gradually rises to local maxima at the edges of the time delay prior. Increasing the number of data points both sharpens the peak at the true time delay and the local maxima at the edges. As stated previously, $\mathbb{E}_{\mathbf{y}} \ln L$ lower bounds the log of the true likelihood, so that we cannot necessarily deduce whether the peak at the true time delay is dominant. However, while the variance spikes at particular values of Δt , as shown in the middle subplot, it is




FIG. 1. Top: Data-averaged log-likelihood $\mathbb{E}_y \ln L$ against the time delay Δt . The true time delay for this figure is $\Delta t_{\text{true}} = 10$ days. The log-likelihood is locally highly oscillatory and globally follows a “W” shape with a global maximum at the true time delay. Middle: The standard deviation $\sigma_{\ln L} = \sqrt{\nabla_y \ln L}$ is negligible compared to the expectation. Bottom: Data-averaged regularized log-likelihood $\mathbb{E}_y \ln L_{\text{reg}}$ formed by subtracting the trend favoring large time delays. Only the maxima in the vicinity of the true time delay remain.

overall negligible and $\mathbb{E}_y \ln L$ should be representative of an individual dataset sample. As we will show in Sec. VB, the gradual increase toward large time delays causes NS to falsely sample large time delays.

What causes this gradual increase in log-likelihood? As we show in Appendix C, the log-likelihood at the true time delay is much larger than the log-likelihood at the edges. Further, we show that the dominant contribution to the log-likelihood is given by the complexity penalty term. This renders large time delays unfavorable compared to the true time delay. In this sense, the likelihood is not pathological and the inference problem is well specified. Namely, if we could sample perfectly from the posterior, we would always recover the true time delay. The difficulty of recovering the true time delay is therefore a problem with the method of sampling, more specifically a problem with the internal dynamics of the global sampling method induced by the gradual likelihood increase.

As shown in Appendix C, “ χ^2 ”, representing the goodness of fit, enters as a subdominant contribution and, in expectation, is of similar size for both the true and large time delays. Any account of the gradual drift would therefore enter as an even lower-order contribution.

From an intuitive standpoint, however, the drift can be understood as follows (Fig. 2). Following the approach in Appendix C, we first decompose the likelihood as

$$p(\mathbf{y}_1, \mathbf{y}_2 | \Delta t) = p(\mathbf{y}_2 | \mathbf{y}_1, \Delta t) p(\mathbf{y}_1). \quad (12)$$

The term $p(\mathbf{y}_1)$ is independent of Δt and can be ignored. Only the predictive distribution $p(\mathbf{y}_2 | \mathbf{y}_1, \Delta t)$ depends on the time delay. Thus, evaluating the likelihood is equivalent to asking how well $\mathbf{y}_1$ predicts $\mathbf{y}_2$.

We now imagine first fitting the light curve $\mathbf{y}_1$ with a GP. This fit will interpolate the data well since we are using

the correct GP model that also generated the data. However, outside of the domain of the data, the GP reverts to the prior and we effectively see a zero-mean function with error bars $\pm A$. As we start increasing Δt and monitor the overlap of the prediction $p(\mathbf{y}_2 | \mathbf{y}_1, \Delta t)$ with $\mathbf{y}_2$, we will see very good agreement at the true time delay, corresponding to the true mode in the posterior. As Δt increases further, the match with the interpolation clearly becomes worse. However, the extrapolated part with the wide error bars starts to overlap with $\mathbf{y}_2$ on one end. In this region, the fit does agree with the data, as all data points lie within the error bars of the predictive distribution. As Δt is increased, this region grows while the region with the bad fit shrinks. Formulated differently, the number of data points ignored grows proportionally to Δt . Therefore, the likelihood increases approximately linearly with Δt , as seen in Figs. 1 and 2.

As we increase the number of data points for a fixed total observation time, the data density increases. Therefore, when Δt is increased, the rate at which the fit ignores data points increases as well. Consequently, the slope of the linear increase also increases with the data density, as observed previously in Fig. 1.

B. Constructing a counterterm prior

Having understood the structure of the log-likelihood, we ask at this point whether we can design a prior that cancels out the gradual increase toward the edges.

Above, we have argued that the likelihood can effectively be thought of the superposition of two competing models, namely, a narrow tall peak at $\Delta t = 0$ and a broad shallow peak at $\Delta t = \pm t_{\text{range}}$. Hence, we can try to isolate just the $\Delta t = 0$ peak by subtracting an appropriate term.




FIG. 2. Illustration of the gradual increase in $\ln L$ as the time delay Δt increases. Two synthetic datasets shown on the right are sampled from the distribution defined in Eq. (8) with true time delay $\Delta t = 10$ days. Subsequently, the log-likelihood $\ln L$ of the GP model is evaluated as a function of the time delay, shown on the left. The evaluation of $\ln L$ is reformulated as visually comparing the predictive distribution $p(\mathbf{y}_2|\mathbf{y}_1)$ conditioned on $\mathbf{y}_1$ with the dataset $\mathbf{y}_2$. This is shown in the sequence of plots on the right for different time delays. At the true time delay, the predictive distribution evidently matches $\mathbf{y}_2$. As the time delay is increased, the mean of $p(\mathbf{y}_2|\mathbf{y}_1)$ does not pass through the data $\mathbf{y}_2$ and $\ln L$ decreases significantly. However, the GP increasingly treats the nonoverlapping data as noise, which provides a good fit. As Δt is increased further, more data are fit as noise and $\ln L$ rises gradually, leading to the “W” shape of $\ln L$ (Fig. 1).

To construct this term, we change our perspective and imagine two fictitious light curves $\mathbf{y}'_1$ and $\mathbf{y}'_2$ that are decorrelated, i.e., have a true time delay of $\Delta t'_{\text{true}} = t_{\text{range}}$. In this case, the largest likelihood occurs when the predictive distribution conditioned on $\mathbf{y}'_1$ extrapolates to noise and $\mathbf{y}'_2$ resides in this noise band. However, the likelihood for our original light curves $\mathbf{y}_1$ and $\mathbf{y}_2$ must be the same at $\Delta t = t_{\text{range}}$ because the noise band has lost its memory of the first light curve. Hence, we should expect perfect cancellation at $\Delta t = t_{\text{range}}$ when we subtract the two likelihoods: $\ln L(t_{\text{range}}) - \ln L'(t_{\text{range}}) = 0$, where L' is the likelihood of the fictitious data.

As we decrease the time delay, the log-likelihood in our fictitious system should decrease because the region of the “good” fit, in which $\mathbf{y}'_2$ resides in the noise band, shrinks. This decrease in log-likelihood should be proportional to Δt as we are losing data points that fit the prediction well. At least in the vicinity of $\Delta t = t_{\text{range}}$, we can therefore expect that the decrease in $\ln L$ and $\ln L'$ is the same.

This discussion motivates the definition of the regularized log-likelihood,

$$\ln L_{\text{reg}}(\Delta t) = \ln L(\Delta t) - \mathbb{E}_{\mathbf{y}_{t_{\text{range}}}} \ln L(\Delta t). \quad (13)$$

Here, we have essentially subtracted the likelihood of the fictitious decorrelated data from the likelihood of the real data. That is, the subtracted term is given by Eq. (10) but where $\mathbf{K}$ uses $\Delta t = t_{\text{range}}$ as the true time delay. Note that we also average the subtracted term over the fictitious data, as otherwise we would have to select a particular fictitious dataset.

To investigate this likelihood, we average over the data in order to be able to make generic statements for different realizations of the data, as done previously. Thus, we obtain the data-averaged regularized log-likelihood,

$$\mathbb{E}_{\mathbf{y}} \ln L_{\text{reg}}(\Delta t) = \mathbb{E}_{\mathbf{y}} \ln L(\Delta t) - \mathbb{E}_{\mathbf{y}_{t_{\text{range}}}} \ln L(\Delta t). \quad (14)$$

This is shown in the bottom subplot in Fig. 1. As expected, the trend in the log-likelihood toward $\pm t_{\text{range}}$ is removed, leaving just the peak near the true time delay. Notably, the entire gradual drift is subtracted, suggesting that the above motivating argument for the construction of the regularized likelihood holds for a larger range of Δt than the argument initially suggests.

It is now tempting to interpret the subtracted term in Eq. (13) as a prior on Δt that removes the full separation of the light curves and therefore acts as a counterterm. However, such a prior can overregularize the likelihood and result in time delays smaller in absolute value than those obtained with flat priors. This happens because this prior does not distinguish between the region around the true peak, where an unbiased uniform prior is desirable, and large time delays, where prior constraints may be applicable. Equivalently, the shape of the prior is essentially a triangle with the slope of the sides matched to the gradual increase of the likelihood. Because the triangle is sloped everywhere, it also regularizes in the vicinity of the true time delay. Ideally, we would like a triangle prior with a flat top, which would cancel the edge modes and act as a uniform prior near the central mode. This begs the question in which region to flatten the triangle, which requires knowledge of where the true time delay is. However,

this is what we intend to infer in the first place and hence do not know *a priori*.

In any case, we generally do not advocate for biased priors to change pathological likelihood functions and instead seek changes in the method to sample unbiased yet pathological likelihoods.

As a final remark, we note that the regularized likelihood was constructed using knowledge about the functional form of the likelihood itself. From a Bayesian point of view, this may be an undesirable feature of a prior. On the other hand, this approach bears similarities to a Jeffreys prior [60], which is also constructed from the likelihood.

C. Further considerations on the likelihood

In this section, we discuss caveats to the above arguments regarding the gradual drift of the log-likelihood.

1. Likelihood or log-likelihood?

First, we should actually consider the likelihood instead of the log-likelihood since by Jensen’s inequality, $\mathbb{E}_{\mathbf{y}} L \geq \exp(\mathbb{E}_{\mathbf{y}} \ln L)$, the data-averaged log-likelihood only provides a lower bound on the log of the data-averaged likelihood. Therefore, the data-averaged likelihood may be larger at any time delay, in particular at the boundary. This means that the likelihood at the boundary could in principle still lie above the likelihood at the true time delay, thus invalidating the above conclusions. As we show in Appendix D, the data-averaged likelihood favors the true time delay, consistent with our conclusions. However, the variance diverges, so that it is not possible to draw conclusions that hold for any dataset, in contrast to the data-averaged log-likelihood, for which the variance remains small. A similar divergence of the variance is displayed by the data-averaged Bayes factor (Appendix E). Furthermore, the data-averaged likelihood, $\int d\mathbf{y} p(\mathbf{y} | \boldsymbol{\theta})$, misses the normalization of the posterior computed by nested sampling, $\int d\mathbf{y} \frac{p(\mathbf{y} | \boldsymbol{\theta}) p(\boldsymbol{\theta})}{Z(\mathbf{y})}$, where $Z(\mathbf{y})$ denotes the Bayesian evidence. As we use uniform priors, the factor $p(\boldsymbol{\theta})$ can be neglected. At best, we can expect the factor $\frac{1}{Z(\mathbf{y})}$ to suppress the contribution to the posterior from datasets for which the peak at the true time delay is subdominant. However, we show numerically in Sec. VB that this is not the case.

2. Nonperturbative effects

Second, the above derivation assumed that the noise σ is small, allowing a perturbative treatment. If σ is comparable to A or larger, $\mathbf{y}_1$ and $\mathbf{y}_2$ lose their mutual correlation, which does not align with the real datasets considered here. A more pressing concern, however, is that we fixed all (hyper)parameters to their true values except for the time delay. We address this question in Sec. VD below.

3. Numerical conditioning

Third, we note that the covariance matrix used in the calculation of the likelihood can become rather ill-conditioned. This error propagates to the numerical value of the likelihood itself. We rule out a systematic influence of this error in Appendix F, showing that the full covariance matrix $\mathbf{K}$

exhibits large condition numbers at regular values of Δt but not at the edges. If the numerical conditioning did have an influence, it would be important to ensure that our analysis reflects the same numerical error. Indeed, the calculation of the data-averaged log-likelihood [Eq. (10)] involves the inverse of $\mathbf{K}_\theta$, which is implemented by a Cholesky decomposition, following the standard GP implementation [55]. In contrast, the data-averaged likelihood [Eq. (D1)] involves the determinant of the matrix and is therefore not numerically representative of the log-likelihood used in NS runs on real data.

4. Beyond stationarity

Fourth, if we allow for a nonstationary kernel, the error bars of the interpolation may not necessarily increase to their maximum value over the decorrelation length scale. More generally, nonstationary GPs with nonconstant mean functions can extrapolate trends and we could expect that the problem with the drift disappears. However, incorporating nonstationary GPs in fits on real data [58] showed that the problem of the large time delays persists, showing that it is more severe than anticipated and that the assumption of stationarity in this section can be relaxed.

Finally, on real data, we do not know the data-generating distribution. Therefore, the GP model may be misspecified. This certainly adds a further layer of complexity. Indeed, this may contribute to the fact that the large time delays occur empirically for real data.

V. NUMERICAL INVESTIGATION

A. Method

We sample the posterior with NS, primarily for three reasons. First, NS naturally incorporates uniform priors and priors with hard constraints, such as those from the flexknots, as NS can operate without gradients of the likelihood. Second, the GP likelihood with a joint uniform prior on all hyperparameters may include singular covariance matrices that lead to divergent negative log-likelihoods, i.e., zero likelihood. NS, in particular POLYCHORD [61,62], adaptively ignores such regions without changing the evidence. Third, kernel, flex-knot, and time delay posteriors are observed to be typically multimodal and exhibit extended degeneracies, which global methods such as NS handle automatically.

In practice, we use the POLYCHORD implementation of NS and the ANESTHETIC package [63] for sample postprocessing, unless stated otherwise. All calculations are performed on the original GP likelihood [Eq. (9)] with uniform priors; i.e., we do not regularize with the counterterm prior (Sec. IV B).

B. Sampling Δt

In this section, we create synthetic datasets of random GP realizations with known time delay and infer the posterior time delay from the synthetic data. We sample Δt using NS with a uniform prior on $[-t_{\text{range}}, t_{\text{range}}]$, where $t_{\text{range}} = T_{\text{max}} - T_{\text{min}}$. Default POLYCHORD settings are used. The sampling is performed for a dataset drawn from Eq. (8) with $\Delta t_{\text{true}} = 10$ days. This is repeated for 25 independent and identically distributed datasets. We thus obtain a data-averaged posterior, formally defined as $\int dy p(\theta | \mathbf{y}) p(\mathbf{y} | \theta_{\text{true}})$, where θ_{true} are

the parameters generating $\mathbf{y}$. In practice, this means that we merge the sets of equally weighted posterior samples obtained from each dataset and perform any further inference with this enlarged set of samples.

The results are shown in Fig. 3. The top subplot shows the data-averaged posterior for $\ell = 10$ days. As the number of observations n_{data} is increased, the uncertainty in the posterior is suppressed and the peak at the true time delay becomes increasingly sharp. This is consistent with the expectation that more data should improve the inference. More precisely, the data density on the time axis increases as n_{data} is increased. At the critical value of $n_{\text{data}} = 100$, there is a data point every $\frac{t_{\text{range}}}{n_{\text{data}}} = 10 \text{ days} = \ell$; i.e., samples start to become correlated. This correlation structure in the data is expected to improve time delay inference. Stated differently, the data points are effectively uncorrelated for $n < 100$ so that both light curves appear as two uncorrelated random draws so that there is nonzero probability density away from the true time delay.

Counterintuitively, in addition to the peak around the true value, the posterior becomes sharper around the edges of the distribution, i.e., at $\pm t_{\text{range}}$, as n_{data} is increased. To obtain a more refined view of the posterior, the middle subplot shows the decomposition of the data-averaged posterior into the contribution from individual datasets for $n_{\text{data}} = 100$. The posterior is peaked at either the true time delay, in which case the error bar on Δt is small, or at both edges, $\pm t_{\text{range}}$, in which case the error bar on Δt is large.

It remains to ask whether these large time delays consistently appear across different values of ℓ . To quantify the width of the data-averaged posterior, we compute the standard deviation $\sigma_{\Delta t}$ of Δt . In agreement with the above description of the data-averaged posterior, the bottom subplot shows that $\sigma_{\Delta t}$ remains roughly unchanged for $\ell = 10$ days as n_{data} is increased. For $\ell < 10$ days, $\sigma_{\Delta t}$ remains large but tends to become smaller for $\ell > 10$ days. This is consistent with the above argument that correlation between data points improves time delay inference. However, $\sigma_{\Delta t}$ tends to increase with n_{data} for $\ell = 100$ days and $\ell = 125$ days as more data points increase the concentration of the posterior at the edges. Note that the fluctuations between the datasets also remain large since $\sigma_{\Delta t}$ is an unstable estimator and hence is expected to converge very slowly with the number of datasets. Overall, the above results suggest that, depending on the randomly drawn dataset, the inferred time delay either agrees with the true time delay or is significantly far away. This observation is systematic across different (hyper)parameters of our setup.

Increasing the number of live points by a factor of 10 fixes the problem by replacing the inferred large time delays with the true time delay. Figure 1 provides the underlying reason. An NS run is initialized by populating the Δt axis uniformly with live points. The probability of a live-point landing in the sharp mode at the true time delay depends on the prior volume of that mode and becomes smaller as the number of data points is increased. Suppose that no live point lands in the mode. As the NS run then proceeds, the live points are pushed up in likelihood while retaining a uniform distribution. The probability of a live point hopping into the mode decreases as the NS iteration number increases, partly because the prior volume at the current likelihood contour decreases. But more importantly, the gradual drift in the log-likelihood (Sec. IV A)




FIG. 3. Top: Data-averaged posterior for $\ell = 10$ days. As n_{data} is increased, the central peak of the posterior becomes sharper around the true time delay, as expected. However, the posterior density also increases around the edges, $\pm t_{\text{range}}$. Middle: Posterior time delays of each dataset for $\ell = 10$ days and $n_{\text{data}} = 100$. The posterior is peaked at either the true time delay or $\pm t_{\text{range}}$. Bottom: Standard deviation $\sigma_{\Delta t}$ of Δt from the data-averaged posterior, for different values of ℓ and n_{data} . The standard deviation is systematically around the same order of magnitude as t_{range} , indicating that the inference of large time delays (as in the top subplot) is systematic across ℓ and n_{data} . The error bars on each value indicate the variation of $\sigma_{\Delta t}$ between different dataset realizations. Evidently, $\sigma_{\Delta t}$ can vary significantly between datasets.

pushes the live points away from the mode, toward the edges $\Delta t = \pm t_{\text{range}}$, so that they aggregate in the local modes at the edges. This explains why NS either infers the true time delay (if a live point happens to land in the true mode at initialization) or a large time delay (if no live points lands in that mode and they are pushed away from it). Moreover, this also explains why the convergence failure becomes more extreme when the number of data points increases. Consequently, increasing the number of live points increases the probability of finding the true mode at the initialization of an NS run as the resolution of the parameter space is increased.

For real datasets and GP models with more parameters, this problem becomes more severe since the fractional parameter space volume of the true mode is even smaller.

C. Convergence by number of live points

To make the previous point about increasing the number of live points to ensure convergence more precise, we perform 50 independent and identically distributed NS runs in the setting of Sec. VB for each value of the number of live points. For each ensemble of runs, we compute the fraction of unconverged runs.

An unconverged run is defined as follows. Given a likelihood (induced by a randomly drawn dataset), we perform a Laplace approximation at the local model corresponding to the true time delay. The location of the mode is obtained by local optimization, initialized at the true time delay that generated the data, and the standard deviation around this mode is obtained via the second derivative. An NS run is said to be unconverged if the posterior mean of the time delay is more than a certain number f of standard deviations away from the Laplace approximation. Note that the standard deviation used here is the one computed from the Laplace approximation. We do not use the NS posterior standard deviation, as this can be very large and falsely indicate convergence. As the Laplace approximation is only a good approximation when the number of observations is large, we expect that even a converged NS posterior may not be in good agreement (i.e., within, say, 3 standard deviations) with the Laplace approximation. We must therefore calibrate f to reflect the level of agreement defining convergence. For this, we increase the number of live points until we determine visual agreement between the two time delay posteriors (as in the previous sections) for multiple NS runs. We then read off the maximum disagreement measured in standard deviations and set f to this value. Hence, for this large value of the live points, there are by definition no unconverged runs.

To reduce the computational cost of the NS runs, we use the BLACKJAX implementation of NS [64,65] on a GPU, which internally uses slice sampling like POLYCHORD. We choose the same termination criterion as in POLYCHORD and, like POLYCHORD, the number of slice sampling steps to be 5 times the number of dimensions, which is 1 since we are only sampling the time delay. These settings make the calculations in this section comparable to the previous ones. The Laplace approximation is implemented via automatic differentiation in JAX [66]. The local optimization is performed via Newton's method [67].




FIG. 4. Fraction of unconverged NS and SMC runs against number of live points and particles, respectively. As the number is increased, nonconvergence becomes less probable since finding the central mode becomes more probable. Thus, the fraction decreases. Increasing the observation window t_{range} from 10^3 to 10^4 days increases the probability of nonconvergence as the central mode occupies a smaller fraction of parameter space and the gradual rise of the likelihood toward the edge modes drives the population of live points or particles further away from the central mode. Similar trends are seen for both NS and SMC. The plotted lines are only visual guides.

Figure 4 shows the fraction of unconverged runs against the number of live points, for light curves with 100 observations and observation windows t_{range} of 10^3 and 10^4 days. Here, we discuss the results for NS. The results for sequential Monte Carlo are discussed in Sec. VII B. As expected, for 10^3 days, the fraction decreases as the number of live points is increased. This is because the probability of initializing a live point in the central mode increases with the number of live points. We see that even at 250 live points, there is still a nonzero probability of missing the central mode. Note that 250 live points is an order of magnitude larger than the standard setting adopted by POLYCHORD. To show that the pathology is driven by the relative volumes of the central mode to the edge modes, as well as the gradual rise that guides live points slowly away from the central mode during an NS run, we increase the observation window of the light curves to 10^4 days. The fraction of unconverged runs increases for all considered values of live points. While a slow decrease in the fraction is visible, consistent with our expectation, this implies that the number of live points must be substantially increased to ensure proper convergence.

Finally, we note that the calculations in Fig. 4 are visibly noisy, since the fraction does not decrease monotonically. The noise can be decreased further by increasing the number of independent NS runs or improving the quality of the Laplace approximation by considering datasets with substantially more observations. Both of these options become quickly infeasible due to the computational costs.

D. Sampling Δt , ℓ , and σ

Here, we address the question of whether the structure of the log-likelihood is preserved if the length scale ℓ and noise σ are sampled as well. For each of these hyperparameter values, the full joint posterior over $\theta = (\Delta t, \ell, \sigma)$ is sampled with NS with standard settings. The priors are set to be uniform in the ranges $\ell \in [0, t_{\text{range}}]$, $\Delta t \in [-t_{\text{range}}, t_{\text{range}}]$, and $\sigma \in [0, 1]$, where t_{range} is the difference between the maximum and minimum observation times. The parameters are set to $\ell = 10$ days, $\Delta t = 10$ days, and $n_{\text{data}} \in \{10, 100, 200, 400, 600, 800\}$. Due to the increased computational cost, we perform the inference only on 12 independent and identically distributed datasets.

We observe that the inferred σ is larger than the true value and the posterior modes at the edges disappear, leaving just the mode around the true value of Δt . This is consistent with the caveat discussed in Sec. IV C. However, due to the small sample size of 12 datasets, we cannot exclude the possibility that there exist datasets for which the edge modes are not suppressed.

Instead, we show that the likelihood still exhibits a “W” shape by reducing the convergence parameters on our NS runs. Specifically, the POLYCHORD settings `num_repeats` and `nlive` are now set to 3 and 15 instead of the default values 15 and 75, respectively. Lowering `nlive` has the effect of coarse graining the parameter space so that the NS run is more susceptible to missing modes with small parameter space volume. Lowering `num_repeats` induces correlations between NS dead points and makes the set of live points susceptible to following funnel-shaped likelihoods during the NS run. This allows us to investigate the structure of the likelihood surface in statistically less significant regions (which a proper NS run would ignore) while circumventing a direct plot of the now higher-dimensional log-likelihood. We refrain from plotting the log-likelihood directly as such plots are visually dominated by numerically highly unstable, albeit statistically negligible, low log-likelihood regions.

The top subplot of Fig. 5 shows the data-averaged time delay posterior for $\Delta t = 10$ days and $\ell = 10$ days and 1 day. For $n_{\text{data}} = 10$ data points, the posterior is noisy and the uncertainty around Δt_{true} is correspondingly large. As n_{data} is further increased, a sharp peak at Δt_{true} develops. However, two additional modes at the edges of the observation time, $\pm 10^3$ days, emerge. Moreover, if we inspect the posterior for each dataset individually, we observe that each posterior is peaked at exactly one of the three modes of the data-averaged posterior.

This means that, for some datasets, the inference of Δt_{true} is possible, whereas it is not possible for other datasets. Since all datasets are random draws from the same GP distribution, this implies that the random fluctuations in a particular GP realization control whether the true time delay can be inferred. In any case, the underlying structure of the log-likelihood still exhibits the drift to the edges.

We note that the synthetic light curves have fewer effective samples than the number of data points suggests as the observations are correlated with a decorrelation length scale ℓ_{true} . We can thus define the effective sample size as $\frac{T}{\ell_{\text{true}}}$, where T is the total observation time. Hence, increasing the number of




FIG. 5. Data-averaged posteriors with reduced nested sampling convergence settings. Top: The data-averaged posterior of the time delay converges to a trimodal distribution as the density of data points increases: While the central peak at the true time delay is visible, two unphysical additional modes at the edges of the observation time interval emerge. Bottom: The unphysical modes are suppressed by increasing the effective sample size of the GP. This is achieved by decreasing the length scale by a factor of 10. In both figures, the inference was performed on synthetic light curves generated with known time delay Δt_{true} and GP length scale ℓ .

effective samples further can be achieved by decreasing ℓ . The bottom subplot of Fig. 5 shows the same calculation but with ℓ_{true} decreased by a factor of 10. This partially suppresses the modes at the edges of the data-averaged posterior, consistent with the expectation that more independent data points should improve the time delay inference; however, further spikes appear away from the true time delay.

VI. REAL DATA AND COSMOLOGICAL PRIORS

In this section, we move beyond toy problems and investigate the likelihood induced by real data and whether priors from cosmology and lens models are sufficient to alleviate the problems presented previously.

A. Real light curve data

The numerical investigation so far has been limited to synthetic data from a toy model. Time delay inference calculations [58] showed that this problem persists on real




FIG. 6. Cuts through the log-likelihood surface for light curve data from the lensed quasar HE0435–1223. The log-likelihood exhibits a “W” shape, with the value of $\ln L$ at the edges comparable to $\ln L$ near the central peak. This drives nested sampling toward the edges of the time delay prior, explaining systematic convergence problems [58].

light curve data under a broad choice of mean functions and kernels.

To show that the log-likelihood still exhibits the global structure described previously, we plot an exemplary cut through the high-dimensional log-likelihood surface for a pair of light curves of the lensed quasar HE0435–1223. While the data-generating GP is unknown in this case, we choose the exponential kernel and remove the choice of the mean function and amplitude hyperparameter by whitening the pair of light curves; i.e., for each light curve, we subtract the mean and divide the magnitudes by the standard deviation. We then plot the log-likelihood as a function of the time delay for different kernel length scales ℓ , as shown in Fig. 6. Manifestly, the log-likelihood qualitatively shows the same global structure. Moreover, the values of the log-likelihood at the edges of the time delay prior are comparable to those at the central peak, which we associate with the true time delay, i.e., the globally dominant mode of the posterior. Note that the likelihood being comparable or higher at the edges than the one at the center is not an issue in a fully Bayesian framework as the parameter space volume of each mode must also be taken into account.

B. Time delay priors from lens models and cosmology

The priors considered so far are based solely on the light curves. Cosmological parameter inference additionally requires a prior on the cosmology and the lens model. We may therefore ask whether the cosmological and lens model priors induce a sufficiently constrained time delay prior that excludes the preference for large time delays.

For this, we take a simplified version of a lens model for HE 0435–1223 described in Wong *et al.* [68] as an example. We adopt a singular power-law elliptical mass distribution

[69] for the lens galaxy, combined with an external shear component. This model includes the following free parameters: the Einstein radius $\theta_E \in [0, 100]''$, power-law slope $\gamma \in [0, 100]$, ellipticity components $e_1 \in [-\frac{1}{2}, \frac{1}{2}]$ and $e_2 \in [-\frac{1}{2}, \frac{1}{2}]$, smoothing scale $s_{\text{scale}} \in [0, 100]''$, and lens centroid position $\theta_{x,y} \in [-2, 2]''$. The external shear is parametrized by its two components $\gamma_{1,2} \in [-\frac{1}{2}, \frac{1}{2}]$. We fix the lens and source redshifts to $z_d = 0.4546$ [70] and $z_s = 1.693$ [71], respectively. For the cosmological priors, we adopt a flat Lambda cold dark matter (Λ CDM) model with uniform distributions for the Hubble parameter $H_0 \in [0, 100] \text{ km s}^{-1} \text{ Mpc}^{-1}$ and matter density $\Omega_m \in [0, 1]$.

We generate samples of lens configurations by drawing parameter values from uniform prior distributions over their respective ranges. A configuration is accepted if exactly four images are formed from a central source position ($0'', 0''$), consistent with the observed quad configuration in HE0435–1223. The time delays Δt are computed by taking pairwise differences between the arrival times of the images. The arrival times are calculated using the `lenstronomy` package [72,73]. We also show the results for the same calculations performed for the narrower prior $H_0 \in [60, 80] \text{ km s}^{-1} \text{ Mpc}^{-1}$ and discuss this further below.

The resulting time delays are collected in a histogram, shown in the top plot of Fig. 7. We observe that a significant fraction of the probability mass lies in the range $[0, 2000]$ days. Hence, we do not expect the lens model and cosmological priors to constrain the time delays sufficiently for time delay inference.

However, we indeed observe that the prior density decreases as $|\Delta t|$ increases. This shows that large time delays, i.e., on the order of 10^3 days for this system, are unphysical, as previously anticipated. The shape of the distribution is expected since the lens model and geometry of the system only constrain the product $H_0 \Delta t$ [74], so that a uniform prior on H_0 induces a distribution on Δt that diverges at $\Delta t = 0$. Sampling lens models and the cosmologies amounts to forming a weighted superposition of these divergent distributions, leading to the observed shape.

As discussed in the Introduction, we assess the impact of setting an implicit prior on the time delay by repeating the above calculation of the histogram but retaining only those configurations that produce time delays within ± 100 days. The resulting histogram is shown in the middle plot of Fig. 7. This is the restriction of the histogram in the top plot to the desired range. We additionally show the distribution of the lens model and cosmology parameters induced by this filtering in the bottom plot. These distributions show which priors we would have had to set in order to produce the restricted time delay distribution.

As follows from the above discussion, we expect that small values of H_0 are removed by removing large time delays from the sample set. This is confirmed by the distribution shown in the corner plot, in which H_0 acquires an approximately linear bias to larger values. This bias propagates to posteriors when performing inference and we leave an analysis of implications thereof to further work. In contrast, the prior on Ω_m remains roughly uniform.

Following this result, we could attempt to narrow the range of the H_0 prior in order to constrain the Δt prior.

The results are also shown in Fig. 7. Narrowing the range to $[60, 80] \text{ km s}^{-1} \text{ Mpc}^{-1}$ still produces time delays of order 2000 days and thus does not sufficiently constrain the time delay. The tails of the time delay distributions become lighter, although this difference is negligible. We again show the filtered prior samples in Fig. 7 and the corresponding corner plot, which reinforce that the narrow prior has negligible influence on the previous conclusions.

We remark that the distributions shown in the corner plot are manifestly nonuniform. For all lens model parameters, most of the nonuniformity comes from conditioning on the quad configuration of HE0435–1223. For the cosmological parameters, this conditioning has negligible influence and most of the nonuniformity is a result of restricting the time delay prior range. This strongly suggests that the lens model parameters are primarily constrained by imaging data, not the light curves.

VII. OTHER GLOBAL SAMPLERS

In this section, we discuss other global sampling algorithms, namely, Parallel Tempering (PT) [75,76] and sequential Monte Carlo (SMC) [77,78] and diffusive NS [79].

A. Discussion

PT and SMC are tempering methods, which sample a sequence of K intermediate distributions $p_{\beta_i} \propto \pi L^{\beta_i}$ with $0 = \beta_1 < \beta_2 < \dots < \beta_K = 1$ between prior and posterior corresponding to $\beta_1 = 0$ and $\beta_K = 1$, respectively.

PT maintains such a ladder of temperatures and runs Markov chain Monte Carlo at each temperature, proposing and accepting swaps between temperatures. Hot chains ($\beta \ll 1$) cross barriers and swaps exchange information between hot and cold chains ($\beta \sim 1$). In the case of time delay inference, the true mode is extremely narrow in prior volume and there is a broad, gently rising region toward the edges that provides moderately large likelihood over a large prior volume. PT may still struggle because the probability a chain ever visits a region with tiny volume can still be vanishingly small, even if barriers are flattened.

Moreover, PT relies on sufficient overlap between the distributions p_{β_i} of adjacent temperatures as this controls the swap acceptance rate. If a cold distribution has a very narrow spike at the true delay but hotter distributions are dominated by the edges, then adjacent MCMC chains can have little overlap in log-likelihood. This requires many temperatures, $K \gg 1$, in order to increase the overlap and long mixing at each temperature, which increases the computational cost in a way analogous to increasing the number of live points in the NS setting.

SMC maintains a distribution of particles starting at the prior and evolves them simultaneously from hot to cold distributions. It does this by reweighting and resampling particles, and subsequently mutation via MCMC. If the central mode has a tiny measure under intermediate β (i.e., the probability of finding a particle in the central mode is tiny), while the boundary region has a huge measure, then before rejuvenation is able to discover the central mode, resampling preferentially resamples boundary particles. After a few iterations through




FIG. 7. Top: Histogram of time delays induced by uniform lens model and cosmology priors, filtering for four images. Larger time delays are visibly penalized but there is significant probability mass at ± 2000 days. Hence, these priors do not sufficiently constrain the time delay. Narrowing the prior of H_0 to the range $[60, 80]$ induces a narrower distribution on the time delay, as evidenced by the lighter tails of the time delay distribution. This effect is negligible, however. Middle: Same histogram, but with time delays filtered to ± 100 days. Bottom: Corresponding distributions on the lens model and cosmology parameters induced by the filtering to ± 100 days. These distributions show the priors that would have given rise to the restricted time delays. Note that H_0 acquires a bias to larger values. Restricting the H_0 prior range to $[60, 80]$ again barely changes the distributions.

different values of β , the population becomes dominated by boundary ancestors. This problem is resolved by sufficiently increasing the number of particles. Overall, this is analogous to the problem of NS, in which the central mode is missed at initialization, and where the number of particles serves an analogous role as the number of live points.

Diffusive NS is a variation of NS that aims to improve exploration of the parameter space by sampling a mixture over likelihood levels, in a fashion that is similar to simulated annealing over the constrained NS level sets. However, the failure mode in the case of time delay inference is that the algorithm never sees the true basin in the first place. Once the level structure and the MCMC moves are built from samples that mostly live near the edges, we still end up with levels that are constructed from the edge modes and MCMC moves within a level that rarely hit the tiny central mode. That is, diffusive NS does not remove the need for at least one particle or trajectory to enter the tiny central mode at some point. This problem is again resolved by increasing the number of levels or equivalently increasing the number of live points.

In summary, these global sampling algorithms, including NS, are built on the same principle, namely, by bridging a unimodal and easy-to-sample prior and a multimodal ergodically broken posterior. If the spurious boundary modes posed local traps, then these methods make it possible to escape them. In the setting of time delay inference, the edge modes are not traps but instead supported by a global trend and large phase space volume. The bridging makes the distribution more volume dominated, so the fundamental problem is finding and retaining the true mode.

B. Numerical example

To illustrate the preceding discussion, we provide a numerical example with SMC with an identical setup as in Sec. V C. Importantly, we observed that the fraction of unconverged runs was positive for all tested hyperparameter settings of SMC. As we are interested in the behavior of the fraction as the number of particles is increased, in particular its decay as the number of particles is increased, the specific value of the fraction is less important. Therefore, the hyperparameters of SMC are chosen to roughly match the fraction of unconverged runs of NS at 25 particles and live points, respectively.

Specifically, we use an adaptive tempered SMC sampler implemented in BLACKJAX [80]. For each run, a fixed number of particles are drawn from the prior. At each iteration i , the next temperature β_{i+1} is selected automatically to achieve a target relative effective sample size of 0.5. A root-finding procedure via bisection is used to determine the increment $\Delta\beta = \beta_{i+1} - \beta_i$. Particles are reweighted by the incremental likelihood power with weights $w_{i+1} \propto w_i L^{\Delta\beta}$, followed by resampling to control weight degeneracy. After resampling, particles are mutated using a random-walk Metropolis MCMC chain targeting the current tempered distribution π_{β_i} . We use a Gaussian proposal with mean zero and covariance adapted at each SMC stage from the empirical weighted particle covariance, as done in the BLACKJAX implementation of NS [64] that we are comparing with in Sec. V C. The covariance is scaled by a heuristic value of 0.125. The algorithm iterates

until $\beta = 1$ is reached. A final particle approximation of the posterior is given by standard SMC estimators.

Figure 4 shows the convergence results for SMC. The decreasing trend of the fraction of unconverged runs with increasing particle number is shared with NS, as expected. Similarly, increasing the length of the observation window t_{range} increases the fraction as particles are more likely to be driven to the edges. This increase is roughly of the same size as for NS, pointing to the fact that the pathology is intrinsic to the structure of the time delay likelihood and independent of the sampling algorithm.

VIII. CONCLUSIONS

The time delay likelihood possesses a fundamental pathological structure, appearing as a “W”-shaped landscape with a dominant central peak at the true value and spurious modes rising at the boundaries of the observation window. Denser sampling over a fixed time span sharpens both the true peak and these edge modes. This structure is not a shortcoming of a particular machine learning model but instead a fundamental consequence of time delay inference itself. In particular, because a light curve is shifted out of the region of overlapping data, the problem is intrinsically extrapolative. GPs make this issue explicit, as they revert to their prior away from conditioning data and must effectively extrapolate a random-walk process. However, the effect is model-independent, as any finite observation window necessitates extrapolation.

In principle, the inference problem is well posed, as the peak at the true delay dominates when sampled correctly. The practical difficulty, however, lies in the sampling methodology. Global methods like nested sampling require stringent settings, including substantially more live points, tighter termination criteria, multiple seeds, and stability checks, to overcome this structure of the likelihood. Otherwise, the live points tend to drift along the boundary-induced likelihood gradient and fail to resolve the central peak. Conversely, local methods and optimizers that implicitly favor small values of $|\Delta t|$ can bias H_0 to larger values. While external priors from lens models and cosmology rarely suppress the large delay modes sufficiently, informative, system-specific priors can offer partial mitigation, though they do not represent a universal solution.

Beyond astrophysics, the problem analyzed in this work is at the center of scalability issues of NS in materials science [54], as explained in the Introduction. In particular, this has seeded efforts for the development of extensions to NS [81–83]. In contrast to the materials case, however, where the dimension of the sampling space is often of order 10^2 – 10^3 , with the goal of letting the dimensionality become as large as possible to approach the thermodynamic limit, the time delay problem is low dimensional as it scales with the number of gravitationally lensed images, and therefore vastly more interpretable. Sampling strategies designed to alleviate this difficulty may therefore be carefully tested and analyzed for the time delay problem, which is significantly more difficult in the case of materials science. Further, a solution to one problem may imply a solution for the other. Together with recent advances in understanding the challenge of incorporat-

ing gradient information into NS [84], this work contributes to ongoing NS method development.

From an algorithmic perspective, NS may be viewed as a particular variation of SMC [85], which constitutes a much broader design space of global sampling algorithms. Again, a solution to the time delay problem for NS may imply a transferable and wider-reaching solution in the larger class of SMC algorithms, thus rendering the time delay problem a test bed for global sampling algorithms.

Overall, we envision that our work provides a foundation for developing further solutions to this sampling instability. Coupled with continued algorithmic development, these efforts will pave the way for fully Bayesian time delay inference, light curve regression, and beyond.

ACKNOWLEDGMENTS

N.K. thanks Gábor Csányi for helpful discussions. N.K. was supported by the Harding Distinguished Postgraduate Scholarship. This work was performed using the Cambridge Service for Data Driven Discovery (CSD3), part of which is operated by the University of Cambridge Research Computing on behalf of the STFC DiRAC HPC Facility [86]. The DiRAC component of CSD3 was funded by BEIS capital funding via STFC capital Grants No. ST/P002307/1 and No. ST/R002452/1 and STFC operations Grant No. ST/R00689X/1. DiRAC is part of the National e-Infrastructure.

DATA AVAILABILITY

The data that support the findings of this article are not publicly available upon publication because it is not technically feasible and/or the cost of preparing, depositing, and hosting the data would be prohibitive within the terms of this research project. The data are available from the authors upon reasonable request.

APPENDIX A: DERIVATION FOR THE TIME DELAY POSTERIOR

Here, we give a derivation of Eq. (6). By Bayes' theorem, we have

$$p(\Delta t | y_1, y_2) = \frac{p(y_1, y_2 | \Delta t)p(\Delta t)}{p(y_1, y_2)}, \quad (\text{A1})$$

with prior $p(\Delta t)$, likelihood $p(y_1, y_2 | \Delta t)$, and evidence $p(y_1, y_2) = p(y_2 | y_1)p(y_1)$. The likelihood may be written in the following marginalized form:

$$\begin{aligned} p(y_1, y_2 | \Delta t) &= \int d\mathbf{w} p(y_1, y_2, \mathbf{w} | \Delta t) \\ &= \int d\mathbf{w} p(y_2 | \Delta t, y_1, \mathbf{w})p(y_1, \mathbf{w} | \Delta t), \end{aligned} \quad (\text{A2})$$

$$(\text{A3})$$

where the integration is performed over all parameters $\mathbf{w} = (w_1, w_2, \dots)^T$. The last term decomposes further as

$$p(y_1, \mathbf{w} | \Delta t) = p(\mathbf{w} | y_1, \Delta t)p(y_1 | \Delta t). \quad (\text{A4})$$

Here, $p(y_1 | \Delta t)$ is the prior probability of y_1 given Δt . However, y_1 only depends on Δt if we are given y_2 . Therefore, we may simplify this as $p(y_1 | \Delta t) = p(y_1)$. Further, $p(\mathbf{w} | y_1, \Delta t)$ is the posterior of $\mathbf{w}$ given only the first light curve y_1 . This is independent of Δt and hence simplifies to $p(\mathbf{w} | y_1)$, as well.

Considering $p(y_2 | \Delta t, y_1, \mathbf{w})$, we note that y_2 is completely determined by Δt and $\mathbf{w}$ and independent of y_1 as Δt and $\mathbf{w}$ determine the underlying function f . In particular, we may write this as

$$p(y_2 | \Delta t, y_1, \mathbf{w}) = \prod_{i=1}^{n_{\text{data}}} \delta(y_2(t_i) - f(t_i + \Delta t, \mathbf{w})), \quad (\text{A5})$$

where δ is the Dirac delta function, but we choose to keep this term in the more general form $p(y_2 | \Delta t, \mathbf{w})$.

Collecting the above equations, we are left with

$$p(\Delta t | y_1, y_2) = \frac{p(\Delta t)}{p(y_2 | y_1)} \int d\mathbf{w} p(y_2 | \Delta t, \mathbf{w})p(\mathbf{w} | y_1). \quad (\text{A6})$$

We usually choose the prior $p(\Delta t)$ to be uniform so that the prefactor of the integral is a constant. This leads to Eq. (6), as required.

APPENDIX B: WEAKLY DETERMINED LENGTH SCALES

We consistently observe that the length scale posterior of the exponential kernel increases for small values of the length scale and subsequently exhibits a degenerate flat posterior on the scale of the time range of the dataset, up to the end of the prior range. In contrast, the other kernels have a much more constrained posterior. It is therefore reasonable to ask whether the problem lies with the data or with the exponential kernel itself. In the following, we conclude that it must be the data and that the posterior is in principle not pathological under optimal assumptions.

For a wide uniform prior, the posterior is proportional to the likelihood. Hence, we only consider the likelihood from here on. We assume that the given data $\mathbf{y}$ is sampled in regular intervals δt from a zero-mean GP with the true exponential kernel $\Sigma_{ij} = b^{|i-j|}$ with unit amplitude and true length scale ℓ_{true} , where $b = e^{-\delta t/\ell_{\text{true}}}$. The log-likelihood is then

$$\ln L(\ell) = -\frac{1}{2} \mathbf{y}^T \mathbf{K}^{-1} \mathbf{y} - \frac{1}{2} \ln |\mathbf{K}| - \frac{n}{2} \ln(2\pi), \quad (\text{B1})$$

where n is the total number of data points (i.e., number of light curves times the number of data points n_{data} per light curve), the covariance matrix $\mathbf{K}_{ij} = a^{|i-j|}$ with $a = e^{-\delta t/\ell}$ depends on the parameter ℓ , and we have set the amplitude in the covariance matrix to its true value for simplicity. We have $\mathbf{y} = \mathbf{K}^{1/2} \hat{\mathbf{y}}$, where $\hat{\mathbf{y}}$ is a sample from a Gaussian with an identity covariance matrix. We now consider the large data limit, in which $n \rightarrow \infty$. In this limit, the individual terms in Eq. (B1) become self-averaging and we may use random matrix theory [87–89] to proceed further. To this end, we introduce the normalized trace operator

$$\tau(\mathbf{A}) = \lim_{n \rightarrow \infty} \frac{1}{n} \mathbb{E}[\text{Tr} \mathbf{A}], \quad (\text{B2})$$

for an $n \times n$ matrix $\mathbf{A}$.

Using the cyclic property of the trace, the first term in Eq. (B1) can now be written as

$$-\frac{1}{2}\mathbf{y}^\top \mathbf{K}^{-1} \mathbf{y} = -\frac{n}{2} \tau(\boldsymbol{\Sigma}^{1/2} \mathbf{K}^{-1} \boldsymbol{\Sigma}^{1/2} \hat{\mathbf{y}} \hat{\mathbf{y}}^\top) \quad (\text{B3})$$

$$= -\frac{n}{2} \tau(\boldsymbol{\Sigma}^{1/2} \mathbf{K}^{-1} \boldsymbol{\Sigma}^{1/2}) \tau(\hat{\mathbf{y}} \hat{\mathbf{y}}^\top). \quad (\text{B4})$$

In the last line, we have used the fact that τ factorizes to leading order for the free product of a deterministic and random matrix. Notably, $\tau(\hat{\mathbf{y}} \hat{\mathbf{y}}^\top) = 1$ since $\hat{\mathbf{y}} \hat{\mathbf{y}}^\top$ is a Wishart matrix. Using the cyclic property of the trace again, we have $\tau(\boldsymbol{\Sigma}^{1/2} \mathbf{K}^{-1} \boldsymbol{\Sigma}^{1/2}) = \tau(\mathbf{K}^{-1} \boldsymbol{\Sigma})$. The inverse $\mathbf{K}^{-1}$ is known analytically,

$$(\mathbf{K}^{-1})_{ij} = \frac{1}{1-a^2} \begin{cases} 1, & \text{if } i = j = 1 \text{ or } i = j = n, \\ 1+a^2, & \text{if } 1 < i = j < n, \\ -a, & \text{if } |i-j| = 1, \\ 0 & \text{else,} \end{cases} \quad (\text{B5})$$

from which we obtain $\tau(\mathbf{K}^{-1} \boldsymbol{\Sigma}) = \frac{1+a^2-2ab}{1-a^2}$.

To handle the second term in Eq. (B1), we can approximate $\mathbf{K}$ by a circulant matrix since boundary effects become negligible in the large data limit,

$$\mathbf{K} = \begin{bmatrix} 1 & a & \cdots & a^{n-1} \\ a & 1 & \cdots & a^{n-2} \\ \vdots & \vdots & \ddots & \vdots \\ a^{n-1} & a^{n-2} & \cdots & 1 \end{bmatrix} \quad (\text{B6})$$

$$\approx \begin{bmatrix} 1 & a & a^2 & \cdots & a^2 & a \\ a & 1 & a & \cdots & a^3 & a^2 \\ a^2 & a & 1 & \cdots & a^4 & a^3 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ a & a^2 & a^3 & \cdots & a & 1 \end{bmatrix}. \quad (\text{B7})$$

This is equivalent to considering a GP on a circle with periodic boundary conditions and letting the radius become large.

The spectrum is then given by

$$\lambda(x_i) = \frac{1-a^2}{1+a^2-2a\cos(\pi x_i)}, \quad (\text{B8})$$

where $x_i = 2i/n$ with $i \in \{0, \dots, n/2\}$. Using this, we have

$$-\frac{1}{2} \ln |\mathbf{K}| = -\frac{1}{2} \text{Tr} \ln \mathbf{K} \quad (\text{B9})$$

$$= -\frac{n}{2} \tau(\ln \mathbf{K}) \quad (\text{B10})$$

$$= -\frac{n}{2} \int_0^1 \ln \lambda(x) dx \quad (\text{B11})$$

$$= -\frac{n}{2} \ln(1-a^2), \quad (\text{B12})$$

where in the third line the sum over eigenvalues was replaced by an integral and in the final line $\int_0^1 \ln(1+a^2-2a\cos(\pi x)) dx = 0$ for all $a \in [0, 1]$ was used.

Therefore, the log-likelihood per data point now becomes

$$\frac{\ln L(\ell)}{n} = -\frac{1}{2} \left(\frac{1+a^2-2ab}{1-a^2} + \ln(1-a^2) + \ln(2\pi) \right). \quad (\text{B13})$$

Maximizing this log-likelihood by setting the derivative to zero yields $\ell = \ell_{\text{true}}$, as desired.

We can now investigate the large ℓ limit. To leading order in ℓ , we have $\ln L(\ell)/n \approx -\frac{1}{2} \frac{1-b}{\delta t} \ell$. Thus, $\ln L(\ell) \rightarrow -\infty$ linearly in ℓ . Hence, the likelihood is well constrained and we do not expect any degeneracies under optimal data assumptions.

APPENDIX C: A REFINED ANALYSIS OF THE LOG-LIKELIHOOD

Here, we offer more clarity on the intuitive description of the log-likelihood given in Sec. IV A.

First, we rewrite Eq. (8) in the following sequential form:

$$\mathbf{y}_1 \sim \mathcal{N}(\mathbf{0}, \mathbf{K}(\mathbf{t}, \mathbf{t}) + \sigma^2 \mathbf{I}), \quad (\text{C1})$$

$$\mathbf{y}_2 | \mathbf{y}_1 \sim \mathcal{N}(\mathbf{m}_{\mathbf{y}_2|\mathbf{y}_1}, \mathbf{K}_{\mathbf{y}_2|\mathbf{y}_1}), \quad (\text{C2})$$

where

$$\mathbf{m}_{\mathbf{y}_2|\mathbf{y}_1} = \mathbf{K}(\mathbf{t} - \Delta t \mathbf{1}, \mathbf{t}) [\mathbf{K}(\mathbf{t}, \mathbf{t}) + \sigma^2 \mathbf{I}]^{-1} \mathbf{y}_1, \quad (\text{C3})$$

$$\mathbf{K}_{\mathbf{y}_2|\mathbf{y}_1} = \mathbf{K}(\mathbf{t}, \mathbf{t}) + \sigma^2 \mathbf{I}$$

$$- \mathbf{K}(\mathbf{t} - \Delta t \mathbf{1}, \mathbf{t}) [\mathbf{K}(\mathbf{t}, \mathbf{t}) + \sigma^2 \mathbf{I}]^{-1} \mathbf{K}(\mathbf{t} - \Delta t \mathbf{1}, \mathbf{t}). \quad (\text{C4})$$

The conditional form of a multivariate normal distribution was used here [see Ref. [55], Eq. (A.6)]. The interpretation of Eq. (C2) is that $\mathbf{y}_1$ was interpolated, giving the smooth interpolant $\mathbf{m}_{\mathbf{y}_2|\mathbf{y}_1}$ with error bars $\mathbf{K}_{\mathbf{y}_2|\mathbf{y}_1}$, and we look at the interpolant at the shifted times $\mathbf{t} - \Delta t \mathbf{1}$. Thus, we pretend that the data $\mathbf{y}_2$ has observation times $\mathbf{t} - \Delta t \mathbf{1}$ and compare this to the interpolant. Also, note that $\mathbf{m}_{\mathbf{y}_2|\mathbf{y}_1}$ is a linear combination of kernels centered on each data point $\mathbf{y}_1$.

The likelihood correspondingly factorizes,

$$p(\mathbf{y}_1, \mathbf{y}_2 | \Delta t) = p(\mathbf{y}_2 | \mathbf{y}_1, \Delta t) p(\mathbf{y}_1 | \Delta t), \quad (\text{C5})$$

which states that we first compare $\mathbf{y}_1$ with the GP prior; i.e., we evaluate $p(\mathbf{y}_1 | \Delta t)$ at the dataset $\mathbf{y}_1$. Then, we interpolate $\mathbf{y}_1$, shift the dataset $\mathbf{y}_2$ to the times $\mathbf{t} - \Delta t \mathbf{1}$, and compare with the interpolant by evaluating $p(\mathbf{y}_2 | \mathbf{y}_1, \Delta t)$. Since the term $p(\mathbf{y}_1 | \Delta t)$ actually does not depend on Δt and therefore only contributes a constant factor to the likelihood, we focus on the behavior of $p(\mathbf{y}_2 | \mathbf{y}_1, \Delta t)$. The log of this term can be decomposed into a χ^2 term measuring the goodness of fit and a penalty term:

$$\begin{aligned} \ln p(\mathbf{y}_2 | \mathbf{y}_1, \Delta t) &= -\frac{1}{2} \underbrace{(\mathbf{y}_2 - \mathbf{m}_{\mathbf{y}_2|\mathbf{y}_1})^\top \mathbf{K}_{\mathbf{y}_2|\mathbf{y}_1}^{-1} (\mathbf{y}_2 - \mathbf{m}_{\mathbf{y}_2|\mathbf{y}_1})}_{\chi^2} \\ &\quad - \frac{1}{2} \underbrace{\ln |\mathbf{K}_{\mathbf{y}_2|\mathbf{y}_1}|}_{\text{penalty}} - \frac{n}{2} \ln(2\pi), \end{aligned} \quad (\text{C6})$$

where we recall that n is the total number of data points, i.e., the number of light curves times the number of data points n_{data} per light curve.

Suppose now that the true time delay is $\Delta t_{\text{true}} = 0$, as this simplifies the analysis, and that the noise is negligible compared to the kernel amplitude, $\sigma \ll A$. We can then consider the interpolated mean and covariance at the true time delay,

$\Delta t = 0$, and for large time delays, meaning that Δt is much larger than the length scale ℓ of the kernel.

For $\Delta t = 0$, we have to first order in σ ,

$$\mathbf{m}_{y_2|y_1} \approx \mathbf{y}_1 - \sigma^2 \mathbf{K}(\mathbf{t}, \mathbf{t})^{-1} \mathbf{y}_1, \quad (\text{C7})$$

$$\mathbf{K}_{y_2|y_1} \approx 2\sigma^2 \mathbf{I}, \quad (\text{C8})$$

meaning that $\mathbf{y}_2$ is modeled as a copy of $\mathbf{y}_1$ with a small amount of distortion of order $\mathcal{O}(\sigma^2)$ and random noise of order $\mathcal{O}(\sigma)$. Using $\mathbf{K}_{y_2|y_1} \approx \mathcal{O}(\sigma^2)$, it can be shown that $\mathbf{y}_2 - \mathbf{y}_1 = \mathcal{O}(\sigma)$. Thus, we have

$$\mathbf{y}_2 - \mathbf{m}_{y_2|y_1} = \mathcal{O}(\sigma) + \sigma^2 \mathbf{K}(\mathbf{t}, \mathbf{t})^{-1} \mathbf{y}_1 = \mathcal{O}(\sigma) \quad (\text{C9})$$

and hence $\chi^2 = \mathcal{O}(n)$ to leading order in σ . The penalty term is $\mathcal{O}(n \ln(2\sigma^2))$, which is large and negative. Therefore, the penalty dominates and we observe a large and positive value of $\ln p(\mathbf{y}_2 | \mathbf{y}_1, \Delta t)$, leading to the peak at the true time delay.

For Δt large, we first note that $\mathbf{K}(\mathbf{t} - \Delta t \mathbf{1}, \mathbf{t}) \approx \mathbf{0}$ since the kernel is stationary. Hence, we immediately obtain

$$\mathbf{m}_{y_2|y_1} \approx \mathbf{0}, \quad (\text{C10})$$

$$\mathbf{K}_{y_2|y_1} \approx \mathbf{K}(\mathbf{t}, \mathbf{t}) + \sigma^2 \mathbf{I}, \quad (\text{C11})$$

meaning that the light curves are modeled as decorrelated from each other. Since the (unconditional) distribution of $\mathbf{y}_2$ is $\mathcal{N}(\mathbf{0}, \mathbf{K}(\mathbf{t}, \mathbf{t}) + \sigma^2 \mathbf{I})$, we have that $\chi^2 = \mathcal{O}(n)$.

Further, the penalty term is $\ln |\mathbf{K}(\mathbf{t}, \mathbf{t}) + \sigma^2 \mathbf{I}| \approx \ln |\mathbf{K}(\mathbf{t}, \mathbf{t})|$, which in general is a function of A , ℓ , and n . For example, for the exponential kernel, comparison with the result from Eq. (B12) gives $\ln |\mathbf{K}(\mathbf{t}, \mathbf{t})| = n(2 \ln A + \ln(1 - e^{-2\delta t/\ell}))$ in the large data limit, where we had to pull the factor of A^2 out of the matrix to use the result, which holds for $A = 1$. Since $\ln(1 - e^{-2\delta t/\ell}) \geq 0$, we have that $\ln |\mathbf{K}(\mathbf{t}, \mathbf{t})| \geq 2n \ln A = \mathcal{O}(n \ln A)$. We assume that a similar scaling holds for other stationary kernels. The important point is that the penalty can be taken to be of order $\mathcal{O}(n \ln A)$.

We can now collect the above results and compare the size of terms. For $\Delta t = 0$ and large Δt , the penalty term dominates with values $\mathcal{O}(n \ln(2\sigma^2))$ and $\mathcal{O}(n \ln A)$, respectively. Since $\sigma \ll A$, the former value is much more negative than the latter and this is what ultimately favors the time delay $\Delta t = 0$. For both values of the time delay, we have comparable values of $\chi^2 = \mathcal{O}(n)$. This is subdominant compared to the penalty term.

APPENDIX D: DATA-AVERAGED LIKELIHOOD

The integral of the likelihood [Eq. (9)] with $\mathbf{y}$ distributed according to Eq. (8) gives directly

$$\mathbb{E}_{\mathbf{y}} L(\boldsymbol{\theta}) = |2\pi(\mathbf{K}_{\boldsymbol{\theta}} + \mathbf{K})|^{-1/2}, \quad (\text{D1})$$

$$\mathbb{E}_{\mathbf{y}} L^2(\boldsymbol{\theta}) = |4\pi^2 \mathbf{K}_{\boldsymbol{\theta}}(\mathbf{K}_{\boldsymbol{\theta}} + 2\mathbf{K})|^{-1/2}, \quad (\text{D2})$$

where $\mathbf{K}_{\boldsymbol{\theta}}$ and $\mathbf{K}$ are $n \times n$ matrices. Here, n is the total number of data points, i.e., number of light curves times the number of data points n_{data} per light curve.

Taking the true time delay to be zero, we have at $\Delta t = 0$, $\mathbf{K}_{\boldsymbol{\theta}} = \mathbf{K}$. Hence,

$$\mathbb{E}_{\mathbf{y}} L(\boldsymbol{\theta}) = (4\pi)^{-n/2} |\mathbf{K}|^{-1/2} \quad (\text{D3})$$

$$= (4\pi)^{-n/2} |\mathbf{K}(\mathbf{t}, \mathbf{t}) + \sigma^2 \mathbf{I}|^{-1/2} \quad (\text{D4})$$

$$\times |\mathbf{K}(\mathbf{t}, \mathbf{t}) + \sigma^2 \mathbf{I} - \mathbf{K}(\mathbf{t}, \mathbf{t})(\mathbf{K}(\mathbf{t}, \mathbf{t}) + \sigma^2 \mathbf{I})^{-1} \times \mathbf{K}(\mathbf{t}, \mathbf{t})|^{-1/2} \quad (\text{D5})$$

$$= (4\pi)^{-n/2} \sigma^{-n} |2\mathbf{K}(\mathbf{t}, \mathbf{t}) + \sigma^2 \mathbf{I}|^{-1/2}, \quad (\text{D6})$$

where we have used the formula for the determinant of a block matrix and the fact that $\mathbf{K}(\mathbf{t}, \mathbf{t}) + \sigma^2 \mathbf{I}$ and $\mathbf{K}(\mathbf{t}, \mathbf{t})$ commute. Similarly,

$$\mathbb{E}_{\mathbf{y}} L^2(\boldsymbol{\theta}) = (12\pi^2)^{-n/2} \sigma^{-2n} |2\mathbf{K}(\mathbf{t}, \mathbf{t}) + \sigma^2 \mathbf{I}|^{-1}. \quad (\text{D7})$$

Thus, the expectation and standard deviation both scale as $\mathcal{O}(\sigma^{-n})$ and diverge for small σ .

For Δt large,

$$\mathbf{K}_{\boldsymbol{\theta}} = \begin{bmatrix} \mathbf{K}(\mathbf{t}, \mathbf{t}) + \sigma^2 \mathbf{I} & \mathbf{0} \\ \mathbf{0} & \mathbf{K}(\mathbf{t}, \mathbf{t}) + \sigma^2 \mathbf{I} \end{bmatrix}. \quad (\text{D8})$$

To calculate $|\mathbf{K}_{\boldsymbol{\theta}} + \mathbf{K}|$, we again use the determinant of a block matrix. This requires the use of Jacobi's formula,

$$|\mathbf{K}(\mathbf{t}, \mathbf{t}) + \sigma^2 \mathbf{I}| \approx |\mathbf{K}(\mathbf{t}, \mathbf{t})| (1 + \sigma^2 \text{Tr}[\mathbf{K}(\mathbf{t}, \mathbf{t})^{-1}]), \quad (\text{D9})$$

from which we obtain $|\mathbf{K}_{\boldsymbol{\theta}} + \mathbf{K}| \approx 3^n |\mathbf{K}(\mathbf{t}, \mathbf{t})|^2$ to leading order in σ . This gives us

$$\mathbb{E}_{\mathbf{y}} L(\boldsymbol{\theta}) \approx (6\pi)^{-n/2} |\mathbf{K}(\mathbf{t}, \mathbf{t})|^{-1}. \quad (\text{D10})$$

By a similar procedure, we obtain

$$\mathbb{E}_{\mathbf{y}} L^2(\boldsymbol{\theta}) \approx (20\pi^2)^{-n/2} |\mathbf{K}(\mathbf{t}, \mathbf{t})|^{-2}. \quad (\text{D11})$$

Thus, the expectation and standard deviation remain finite for small σ and large Δt .

APPENDIX E: BAYES FACTOR

The Bayes factor between the true time delay $\Delta t = 0$ and large time delay, $\Delta t = \infty$, is given by

$$B = \frac{p(\mathbf{y} | \Delta t = 0)}{p(\mathbf{y} | \Delta t = \infty)} \quad (\text{E1})$$

$$= \sqrt{\frac{|\mathbf{K}_{\infty}|}{|\mathbf{K}_0|}} \exp\left(-\frac{1}{2} \mathbf{y}^T (\mathbf{K}_0^{-1} - \mathbf{K}_{\infty}^{-1}) \mathbf{y}\right), \quad (\text{E2})$$

where $\mathbf{K}_0$ and $\mathbf{K}_{\infty}$ are $n \times n$ matrices. To calculate the variance, $\mathbb{E}_{\mathbf{y}} B^2$ must be finite. We now show that this is not the case.

Since $\mathbf{y}$ is distributed as $\mathcal{N}(\mathbf{0}, \mathbf{K}_0)$, the Gaussian integral $\mathbb{E}_{\mathbf{y}} B^2$ is finite if the matrix $\mathbf{Q} = 3\mathbf{K}_0^{-1} - 2\mathbf{K}_{\infty}^{-1}$ is positive definite. This is equivalent to showing that, for every nonzero $\mathbf{v}$, $\mathbf{v}^T \mathbf{Q} \mathbf{v} > 0$. Noting that $\mathbf{K}_0$ is positive definite, we change basis to $\mathbf{w} = \mathbf{K}_0^{-1/2} \mathbf{v}$ so that the condition becomes $\mathbf{w}^T \mathbf{K}_0^{-1/2} \mathbf{S} \mathbf{K}_0^{1/2} \mathbf{w} > 0$, where we defined $\mathbf{S} = 3\mathbf{I} - 2\mathbf{K}_0 \mathbf{K}_{\infty}^{-1}$. This is a similarity transform of $\mathbf{S}$ by $\mathbf{K}_0^{-1/2}$. Since a similarity transform preserves positive definiteness, it remains to show that $\mathbf{S}$ is positive definite.

We now calculate the eigenvalues of $\mathbf{S}$. We have

$$\mathbf{K}_0 \mathbf{K}_\infty^{-1} = \begin{bmatrix} \mathbf{K}(\mathbf{t}, \mathbf{t}) + \sigma^2 \mathbf{I} & \mathbf{K}(\mathbf{t}, \mathbf{t}) \\ \mathbf{K}(\mathbf{t}, \mathbf{t}) & \mathbf{K}(\mathbf{t}, \mathbf{t}) + \sigma^2 \mathbf{I} \end{bmatrix} \quad (\text{E3})$$

$$\times \begin{bmatrix} (\mathbf{K}(\mathbf{t}, \mathbf{t}) + \sigma^2 \mathbf{I})^{-1} & \mathbf{0} \\ \mathbf{0} & (\mathbf{K}(\mathbf{t}, \mathbf{t}) + \sigma^2 \mathbf{I})^{-1} \end{bmatrix} \quad (\text{E4})$$

$$= \begin{bmatrix} \mathbf{I} & \mathbf{K}(\mathbf{t}, \mathbf{t})(\mathbf{K}(\mathbf{t}, \mathbf{t}) + \sigma^2 \mathbf{I})^{-1} \\ \mathbf{K}(\mathbf{t}, \mathbf{t})(\mathbf{K}(\mathbf{t}, \mathbf{t}) + \sigma^2 \mathbf{I})^{-1} & \mathbf{I} \end{bmatrix}. \quad (\text{E5})$$

Since for any block symmetric matrix, it holds that

$$\begin{vmatrix} \mathbf{A} & \mathbf{B} \\ \mathbf{B} & \mathbf{A} \end{vmatrix} = |\mathbf{A} + \mathbf{B}| |\mathbf{A} - \mathbf{B}|, \quad (\text{E6})$$

the eigenvalues of $\mathbf{K}_0 \mathbf{K}_\infty^{-1}$ are those of $\mathbf{I} + \mathbf{K}(\mathbf{t}, \mathbf{t})(\mathbf{K}(\mathbf{t}, \mathbf{t}) + \sigma^2 \mathbf{I})^{-1}$ and $\mathbf{I} - \mathbf{K}(\mathbf{t}, \mathbf{t})(\mathbf{K}(\mathbf{t}, \mathbf{t}) + \sigma^2 \mathbf{I})^{-1}$. Let κ_i , $i \in \{1, \dots, n/2\}$, be the eigenvalues of $\mathbf{K}(\mathbf{t}, \mathbf{t})$. Then, $\mathbf{K}_0 \mathbf{K}_\infty^{-1}$ has eigenvalues $\lambda_i = 1 + \frac{\kappa_i}{\kappa_i + \sigma^2}$ for $i \in \{1, \dots, n/2\}$ and $\lambda_i = 1 - \frac{\kappa_i}{\kappa_i + \sigma^2}$ for $i \in \{n/2 + 1, \dots, n\}$. Hence, $\mathbf{S}$ has eigenvalues $\rho_i = 3 - 2\lambda_i$. Expanding to leading order in σ , these are

$$\rho_i \approx \begin{cases} -1 + 2\frac{\sigma^2}{\kappa_i}, & \text{for } i \in \{1, \dots, n/2\}, \\ 3 - 2\frac{\sigma^2}{\kappa_i}, & \text{for } i \in \{n/2 + 1, \dots, n\}. \end{cases} \quad (\text{E7})$$

Following the argument in Appendix C, the eigenvalues of κ_i are of order $\mathcal{O}(n \ln A)$ and thus non-negligible. Therefore, there are $\rho_i < 0$ for small σ , i.e., $\mathbf{S}$ and hence $\mathbf{Q}$ is not positive definite so that the original integral under consideration diverges.

APPENDIX F: NUMERICAL CONDITIONING

We define the condition number of a covariance matrix as the ratio of the largest and smallest eigenvalues. Importantly, a larger condition number corresponds to a more singular matrix. The numerical errors accumulated in matrix operations,




FIG. 8. Condition number of the GP covariance matrix. The plot was created with 6×10^4 evenly spaced points along the time delay axis.

such as a Cholesky decomposition, are also correspondingly larger.

Figure 8 shows the condition number of the GP covariance matrix,

$$\begin{bmatrix} \mathbf{K}(\mathbf{t}, \mathbf{t}) + \sigma^2 \mathbf{I} & \mathbf{K}(\mathbf{t}, \mathbf{t} - \Delta \mathbf{t}) \\ \mathbf{K}(\mathbf{t} - \Delta \mathbf{t}, \mathbf{t}) & \mathbf{K}(\mathbf{t}, \mathbf{t}) + \sigma^2 \mathbf{I} \end{bmatrix}, \quad (\text{F1})$$

as a function of Δt , with fixed $\sigma = 0.01$ and 100 evenly spaced points between 0 and 10^3 in $\mathbf{t}$. The condition number spikes at regular intervals. This is expected since some of the values in $\mathbf{t}$ and $\mathbf{t} - \Delta \mathbf{t}$ are identical for certain values of Δt , which corresponds to two data points measured at the same time. At these values of Δt , the condition number is regularized by the $\sigma^2 \mathbf{I}$ terms. Indeed, increasing σ reduces the height of the spikes. Further, the figure shows that increasing ℓ also increases condition numbers. As we have more data points within the same correlation length, this is consistent with the previous result.

Notably, we do not observe exceptionally large condition numbers at the edges of the time delay prior. Hence, we conclude that the increase in the likelihood at the edges of the time delay prior does not stem from numerical ill-conditioning.

-
- [1] S. Refsdal, On the possibility of determining Hubble's parameter and the masses of galaxies from the gravitational lens effect, *Mon. Not. R. Astron. Soc.* **128**, 307 (1964).
 - [2] T. Treu and P. J. Marshall, Time delay cosmography, *Astron. Astrophys. Rev.* **24**, 11 (2016).
 - [3] S. Birrer, M. Millon, D. Sluse, A. Shajib, F. Courbin, L. Koopmans, S. Suyu, and T. Treu, Time-delay cosmography: Measuring the Hubble constant and other cosmological parameters with strong gravitational lensing, *Space Sci. Rev.* **220**, 48 (2024).
 - [4] T. Treu, S. H. Suyu, and P. J. Marshall, Strong lensing time-delay cosmography in the 2020s, *Astron. Astrophys. Rev.* **30**, 8 (2022).
 - [5] S. Birrer, G. P. Smith, A. J. Shajib, D. Rychanowski, and N. Arendse, Challenges and opportunities for time-delay cosmography with multi-messenger gravitational lensing, *Phil. Trans. A* **383**, 20240130 (2025).
 - [6] B. M. Peterson, Reverberation mapping of active galactic nuclei, *Publ. Astron. Soc. Pac.* **105**, 247 (1993).
 - [7] E. M. Cackett, M. C. Bentz, and E. Kara, Reverberation mapping of active galactic nuclei: From X-ray corona to dusty torus, *iScience* **24**, 102557 (2021).
 - [8] H. Netzer and B. M. Peterson, Reverberation mapping and the physics of active galactic nuclei, in *Astronomical Time Series. Astrophysics and Space Science Library*, edited by D. Maoz, A. Sternberg, and E. M. Leibowitz (Springer, Dordrecht, 1997), Vol. 218, pp. 85–108.
 - [9] P. A. Abell, J. Allison, S. F. Anderson, J. R. Andrew, J. R. P. Angel, L. Armus, D. Arnett, S. J. Asztalos, T. S. Axelrod,

- S. Bailey, *et al.*, LSST science book, version 2.0, Tech. Rep. (LSST Corporation, Tucson, AZ, 2009).
- [10] Ž. Ivezić, S. M. Kahn, J. A. Tyson, B. Abel, E. Acosta, R. Allsman, D. Alonso, Y. AlSayyad, S. F. Anderson, J. Andrew, *et al.*, LSST: From science drivers to reference design and anticipated data products, *Astrophys. J.* **873**, 111 (2019).
- [11] J. Kollmeier, S. Anderson, G. Blanc, M. Blanton, K. Covey, J. Crane, N. Drory, P. Frinchaboy, C. Froning, J. Johnson, *et al.*, SDSS-V: Pioneering panoptic spectroscopy, *Bull. Am. Astron. Soc.* **51**, 274 (2019).
- [12] A. Almeida, S. F. Anderson, M. Argudo-Fernández, C. Badenes, K. Barger, J. K. Barrera-Ballesteros, C. F. Bender, E. Benítez, F. Besser, J. C. Bird, *et al.*, The eighteenth data release of the Sloan Digital Sky Surveys: Targeting and first spectra from SDSS-V, *Astrophys. J. Suppl. Ser.* **267**, 44 (2023).
- [13] C. Grier, J. R. Trump, Y. Shen, K. Horne, K. Kinemuchi, I. D. McGreer, D. Starkey, W. Brandt, P. Hall, C. Kochanek, *et al.*, The Sloan Digital Sky Survey reverberation mapping project: H α and H β reverberation measurements from first-year spectroscopy and photometry, *Astrophys. J.* **851**, 21 (2017).
- [14] C. Grier, Y. Shen, K. Horne, W. Brandt, J. Trump, P. Hall, K. Kinemuchi, D. Starkey, D. Schneider, L. C. Ho, *et al.*, The Sloan Digital Sky Survey reverberation mapping project: Initial C IV lag results from four years of data, *Astrophys. J.* **887**, 38 (2019).
- [15] I. Jennifer, H. Li, Y. Shen, W. Brandt, C. Grier, P. Hall, L. Ho, Y. Homayouni, K. Horne, D. Schneider, *et al.*, The Sloan Digital Sky Survey reverberation mapping project: Comparison of lag measurement methods with simulated observations, *Astrophys. J.* **884**, 119 (2019).
- [16] Y. Homayouni, J. R. Trump, C. Grier, K. Horne, Y. Shen, W. Brandt, K. S. Dawson, G. F. Alvarez, P. J. Green, P. Hall, *et al.*, The Sloan Digital Sky Survey reverberation mapping project: Mg II lag results from four years of monitoring, *Astrophys. J.* **901**, 55 (2020).
- [17] H. McDougall, B. J. Pope, and T. M. Davis, LITMUS: Bayesian lag recovery in reverberation mapping with fast differentiable models, *Publ. Astron. Soc. Aust.* **43**, e018 (2026).
- [18] A. Hojjati, A. G. Kim, and E. V. Linder, Robust strong lensing time delay estimation, *Phys. Rev. D* **87**, 123512 (2013).
- [19] M. Tewes, F. Courbin, and G. Meylan, COSMOGRAIL: The cosmological monitoring of gravitational lenses-XI. Techniques for time delay measurement in presence of microlensing, *Astron. Astrophys.* **553**, A120 (2013).
- [20] Y. Zu, C. Kochanek, and B. M. Peterson, An alternative approach to measuring reverberation lags in active galactic nuclei, *Astrophys. J.* **735**, 80 (2011).
- [21] Y.-R. Li, J.-M. Wang, L. C. Ho, P. Du, and J.-M. Bai, A Bayesian approach to estimate the size and structure of the broad-line region in active galactic nuclei using reverberation mapping data, *Astrophys. J.* **779**, 110 (2013).
- [22] A. Pancoast, B. J. Brewer, and T. Treu, Modelling reverberation mapping data—I. Improved geometric and dynamical models and comparison with cross-correlation results, *Mon. Not. R. Astron. Soc.* **445**, 3055 (2014).
- [23] A. Pancoast, B. J. Brewer, T. Treu, D. Park, A. J. Barth, M. C. Bentz, and J.-H. Woo, Modelling reverberation mapping data—II. Dynamical modelling of the Lick AGN Monitoring Project 2008 data set, *Mon. Not. R. Astron. Soc.* **445**, 3073 (2014).
- [24] D. Starkey, K. Horne, M. Fausnaugh, B. Peterson, M. C. Bentz, C. Kochanek, K. Denney, R. Edelson, M. Goad, G. De Rosa, *et al.*, Space Telescope and Optical Reverberation Mapping Project. VI. Reverberating disk models for NGC 5548, *Astrophys. J.* **835**, 65 (2017).
- [25] C. Vanderriest, J. Schneider, G. Herpe, M. Chevreton, M. Moles, and G. Wlerick, The value of the time delay δt (A, B) for the “double” quasar 0957+561 from optical photometric monitoring, *Astron. Astrophys.* **215**, 1 (1989).
- [26] J. Lehar, J. N. Hewitt, D. H. Roberts, and B. F. Burke, The radio time delay in the double quasar 0957+561, *Astrophys. J.* **384**, 453 (1992).
- [27] W. H. Press, G. B. Rybicki, and J. N. Hewitt, The time delay of gravitational lens 0957+561. I. Methodology and analysis of optical photometric data. II. Analysis of radio data and combined optical-radio analysis, *Astrophys. J.* **385**, 404 (1992).
- [28] J. Pelt, R. Kayser, S. Refsdal, and T. Schramm, The light curve and the time delay of QSO 0957+561, [arXiv:astro-ph/9501036](https://arxiv.org/abs/astro-ph/9501036).
- [29] T. Kundic, E. L. Turner, W. N. Colley, J. R. Gott, J. E. Rhoads, Y. Wang, L. E. Bergeron, K. A. Gloria, D. C. Long, S. Malhotra, *et al.*, A robust determination of the time delay in 0957+561A, B and a measurement of the global value of Hubble’s constant, *Astrophys. J.* **482**, 75 (1997).
- [30] S. Suyu, P. Marshall, R. Blandford, C. Fassnacht, L. Koopmans, J. McKean, and T. Treu, Dissecting the gravitational lens B1608+656. I. Lens potential reconstruction, *Astrophys. J.* **691**, 277 (2009).
- [31] S. Suyu, P. Marshall, M. Auger, S. Hilbert, R. Blandford, L. Koopmans, C. Fassnacht, and T. Treu, Dissecting the gravitational lens B1608+656. II. Precision measurements of the Hubble constant, spatial curvature, and the dark energy equation of state, *Astrophys. J.* **711**, 201 (2010).
- [32] L. Koopmans, T. Treu, C. Fassnacht, R. Blandford, and G. Surpi, The Hubble constant from the gravitational lens B1608+656, *Astrophys. J.* **599**, 70 (2003).
- [33] G. Dobler, C. D. Fassnacht, T. Treu, P. Marshall, K. Liao, A. Hojjati, E. Linder, and N. Rumbaugh, Strong lens time delay challenge. I. Experimental design, *Astrophys. J.* **799**, 168 (2015).
- [34] K. Liao, T. Treu, P. Marshall, C. D. Fassnacht, N. Rumbaugh, G. Dobler, A. Aghamousa, V. Bonvin, F. Courbin, A. Hojjati, *et al.*, Strong lens time delay challenge. II. Results of TDC1, *Astrophys. J.* **800**, 11 (2015).
- [35] S. R. Kumar, M. Tewes, C. Stalin, F. Courbin, I. Asfandiyarov, G. Meylan, E. Eulaers, T. Prabhu, P. Magain, H. Van Winckel, *et al.*, COSMOGRAIL: The cosmological monitoring of gravitational lenses-XIV. Time delay of the doubly lensed quasar SDSS J1001+5027, *Astron. Astrophys.* **557**, A44 (2013).
- [36] F. Courbin, V. Bonvin, E. Buckley-Geer, C. Fassnacht, J. Frieman, H. Lin, P. Marshall, S. Suyu, T. Treu, T. Anguita, *et al.*, COSMOGRAIL XVI: Time delays for the quadruply imaged quasar DES J0408–5354 with high-cadence photometric monitoring, *Astron. Astrophys.* **609**, A71 (2018).
- [37] J. Skilling, Nested sampling for general Bayesian computation, *Bayesian Anal.* **1**, 833 (2006).
- [38] G. Ashton, N. Bernstein, J. Buchner, X. Chen, G. Csányi, A. Fowlie, F. Feroz, M. Griffiths, W. Handley, M. Habeck, *et al.*, Nested sampling for physical scientists, *Nat. Rev. Methods Primers* **2**, 39 (2022).

- [39] J. Buchner, Nested sampling methods, *Statist. Surv.* **17**, 169 (2023).
- [40] N. Kroupa, D. Yallup, W. Handley, and M. Hobson, Kernel-, mean-, and noise-marginalized Gaussian processes for exoplanet transits and H_0 inference, *Mon. Not. R. Astron. Soc.* **528**, 1232 (2024).
- [41] S. Chib, Marginal likelihood from the Gibbs output, *J. Am. Stat. Assoc.* **90**, 1313 (1995).
- [42] A. Gelman and X.-L. Meng, Simulating normalizing constants: From importance sampling to bridge sampling to path sampling, *Stat. Sci.* **13**, 163 (1998).
- [43] N. Lartillot and H. Philippe, Computing Bayes factors using thermodynamic integration, *Syst. Biol.* **55**, 195 (2006).
- [44] R. J. Urick, *Principles of Underwater Sound-2* (McGraw-Hill, New York, 1975).
- [45] G. L. Turin, Introduction to spread-spectrum antimultipath techniques and their application to urban digital radio, *Proc. IEEE* **68**, 328 (2005).
- [46] J. Chen, J. Benesty, and Y. Huang, Time delay estimation in room acoustic environments: An overview, *EURASIP J. Adv. Signal Process.* **2006**, 026503 (2006).
- [47] K. Gedalyahu and Y. C. Eldar, Time-delay estimation from low-rate samples: A union of subspaces approach, *IEEE Trans. Signal Process.* **58**, 3017 (2010).
- [48] A. Quazi, An overview on the time delay estimate in active and passive systems for target localization, *IEEE Trans. Acoust. Speech Signal Process.* **29**, 527 (1981).
- [49] S. Heerah, R. Molinari, S. Guerrier, and A. Marshall-Colon, Granger-causal testing for irregularly sampled time series with application to nitrogen signalling in Arabidopsis, *Bioinformatics* **37**, 2450 (2021).
- [50] D. P. Schaff and F. Waldhauser, Waveform cross-correlation-based differential travel-time measurements at the Northern California Seismic Network, *Bull. Seismol. Soc. Am.* **95**, 2446 (2005).
- [51] M. Hoffmann, M. Rosenbaum, and N. Yoshida, Estimation of the lead-lag parameter from non-synchronous data, *Bernoulli* **19**, 426 (2013).
- [52] S. Bjorklund and L. Ljung, A review of time-delay estimation techniques, in *Proceedings of the 42nd IEEE International Conference on Decision and Control (IEEE cat. no. 03CH37475)* (IEEE, Piscataway, NJ, 2003), Vol. 3, pp. 2502–2507.
- [53] L. B. Pártay, G. Csányi, and N. Bernstein, Nested sampling for materials, *Eur. Phys. J. B* **94**, 159 (2021).
- [54] P. Y. Chew and A. Reinhardt, Phase diagrams—why they matter and how to predict them, *J. Chem. Phys.* **158**, 030902 (2023).
- [55] C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning* (The MIT Press, Cambridge, MA, 2005).
- [56] C. L. MacLeod, Ž. Ivezić, C. Kochanek, S. Kozłowski, B. Kelly, E. Bullock, A. Kimball, B. Sesar, D. Westman, K. Brooks, *et al.*, Modeling the time variability of SDSS Stripe 82 quasars as a damped random walk, *Astrophys. J.* **721**, 1014 (2010).
- [57] Y. Zu, C. Kochanek, S. Kozłowski, and A. Udalski, Is quasar optical variability a damped random walk? *Astrophys. J.* **765**, 106 (2013).
- [58] N. Kroupa, D. Yallup, and W. Handley, Time delays and stationarity in quasar light curves, *Phys. Rev. D* **113**, 063549 (2026).
- [59] L. T. Hergt, W. J. Handley, M. P. Hobson, and A. N. Lasenby, Bayesian evidence for the tensor-to-scalar ratio r and neutrino masses m_ν : Effects of uniform versus logarithmic priors, *Phys. Rev. D* **103**, 123511 (2021).
- [60] H. Jeffreys, An invariant form for the prior probability in estimation problems, *Proc. R. Soc. London A* **186**, 453 (1946).
- [61] W. Handley, M. Hobson, and A. Lasenby, POLYCHORD: Nested sampling for cosmology, *Mon. Not. R. Astron. Soc.: Lett.* **450**, L61 (2015).
- [62] W. J. Handley, M. P. Hobson, and A. N. Lasenby, PolyChord: Next-generation nested sampling, *Mon. Not. R. Astron. Soc.* **453**, 4385 (2015).
- [63] W. Handley, anesthetic: Nested sampling visualisation, *J. Open Source Software* **4**, 1414 (2019).
- [64] D. Yallup, N. Kroupa, and W. Handley, Nested slice sampling, in *Frontiers in Probabilistic Inference: Learning meets Sampling* (ICLR 2025 Workshop, Singapore, 2025).
- [65] D. Yallup, N. Kroupa, and W. Handley, Nested slice sampling: Vectorized nested sampling for GPU-accelerated inference, *arXiv:2601.23252*.
- [66] R. Frostig, M. J. Johnson, and C. Leary, Compiling machine learning programs via high-level tracing, *Syst. Mach. Learn.* **4**, 9 (2018).
- [67] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, *Numerical Recipes: The Art of Scientific Computing*, 3rd ed. (Cambridge University Press, Cambridge, UK, 2007).
- [68] K. C. Wong, S. H. Suyu, M. W. Auger, V. Bonvin, F. Courbin, C. D. Fassnacht, A. Halkola, C. E. Rusu, D. Sluse, A. Sonnenfeld, *et al.*, H0LiCOW–IV. Lens mass model of HE 0435–1223 and blind measurement of its time-delay distance for cosmology, *Mon. Not. R. Astron. Soc.* **465**, 4895 (2017).
- [69] R. Barkana, FASTELL: Fast calculation of a family of elliptical mass gravitational lens models, Astrophysics Source Code Library, record ascl19910.003, 1999.
- [70] N. D. Morgan, C. S. Kochanek, O. Pevunova, and P. L. Schechter, The lens redshift and galaxy environment for HE 0435–1223, *Astron. J.* **129**, 2531 (2005).
- [71] D. Sluse, D. Hutsemékers, F. Courbin, G. Meylan, and J. Wambsganss, Microlensing of the broad line region in 17 lensed quasars, *Astron. Astrophys.* **544**, A62 (2012).
- [72] S. Birrer and A. Amara, lenstronomy: Multi-purpose gravitational lens modelling software package, *Phys. Dark Universe* **22**, 189 (2018).
- [73] S. Birrer, A. J. Shajib, D. Gilman, A. Galan, J. Aalbers, M. Millon, R. Morgan, G. Pagano, J. W. Park, L. Teodori, *et al.*, lenstronomy II: A gravitational lensing software ecosystem, *J. Open Source Software* **6**, 3283 (2021).
- [74] R. Narayan and M. Bartelmann, Lectures on gravitational lensing, *arXiv:astro-ph/9606001*.
- [75] K. Hukushima and K. Nemoto, Exchange Monte Carlo method and application to spin glass simulations, *J. Phys. Soc. Jpn.* **65**, 1604 (1996).
- [76] D. J. Earl and M. W. Deem, Parallel tempering: Theory, applications, and new perspectives, *Phys. Chem. Chem. Phys.* **7**, 3910 (2005).
- [77] N. Chopin, A sequential particle filter method for static models, *Biometrika* **89**, 539 (2002).
- [78] P. Del Moral, A. Doucet, and A. Jasra, Sequential Monte Carlo samplers, *J. R. Stat. Soc. B* **68**, 411 (2006).

- [79] B. J. Brewer, L. B. Pártay, and G. Csányi, Diffusive nested sampling, *Stat. Comput.* **21**, 649 (2011).
- [80] A. Cabezas, A. Corenflos, J. Lao, R. Louf, A. Carnec, K. Chaudhari, R. Cohn-Gordon, J. Coullon, W. Deng, S. Duffield, *et al.*, BlackJAX: Composable Bayesian inference in JAX, [arXiv:2402.10797](https://arxiv.org/abs/2402.10797).
- [81] S. Martiniani, J. D. Stevenson, D. J. Wales, and D. Frenkel, Superposition enhanced nested sampling, *Phys. Rev. X* **4**, 031034 (2014).
- [82] R. J. Baldock, L. B. Pártay, A. P. Bartók, M. C. Payne, and G. Csányi, Determining pressure-temperature phase diagrams of materials, *Phys. Rev. B* **93**, 174108 (2016).
- [83] N. Unglert, L. B. Pártay, and G. K. H. Madsen, Replica exchange nested sampling, *J. Chem. Theory Comput.* **21**, 7304 (2025).
- [84] N. Kroupa, G. Csányi, and W. Handley, Resonances in reflective Hamiltonian Monte Carlo, *Phys. Rev. E* **111**, 045308 (2025).
- [85] R. Salomone, L. F. South, C. Drovandi, D. P. Kroese, and A. M. Johansen, Unbiased and consistent nested sampling via sequential Monte Carlo, *J. R. Stat. Soc. Ser. B Stat. Methodol.* **87**, 1221 (2025).
- [86] <http://www.dirac.ac.uk/>.
- [87] A. M. Tulino and S. Verdú, *Random Matrix Theory and Wireless Communications* (Now Publishers, Hanover, MA, 2004).
- [88] T. Tao, *Topics in Random Matrix Theory* (American Mathematical Society, Providence, RI, 2012), Vol. 132.
- [89] M. Potters and J.-P. Bouchaud, *A First Course in Random Matrix Theory: For Physicists, Engineers and Data Scientists* (Cambridge University Press, Cambridge, 2020).