manuscript submitted to JGR: Solid Earth

# Bayesian Calibration of dynamic models of earthquake sequences using observations from past large earthquakes

Hojjat Kaveh[1,*], Oliver R. A. Dunbar [1], Jean Philippe Avouac[1] , Andrew M. Stuart[1]

[1] *California Institute of Technology*

[*] *corresponding author: hkaveh@caltech.edu*

## Key points

- We develop a Bayesian framework to infer earthquake-cycle parameters from statistics of past large earthquakes.
- The method exploits stationary statistics to avoid estimating unknown initial conditions in chaotic systems.
- Ensemble Kalman inversion combined with surrogate modeling enables efficient uncertainty quantification.

## ABSTRACT

Physics-based models of the earthquake cycle could be used for time-dependent hazard assessment. For such an application, their parameters must be calibrated so that simulated earthquake sequences reproduce the statistics of past earthquakes, including recurrence statistics and magnitudes. This is challenging because the dynamics are described by nonlinear partial differential equations, initial conditions are unknown, and records of past earthquakes are sparse, incomplete, and noisy. To address this challenge, we present a Bayesian inverse-problem framework to infer heterogeneous parameters in an earthquake-cycle model from observations of past large earthquakes. Rather than matching a time-series of individual events, we exploit ergodicity (long-term statistics independent of initial conditions) and invert for stationary, spatially distributed summary statistics of event timing and slip, including mean interevent times, coefficients of variation, burstiness, and average coseismic slip, which are quantities that can in principle be estimated from paleoseismic and historical records. Observational noise and event-detection thresholds are explicitly accounted for in the construction of these statistics. The inverse problem is solved using a derivative-free ensemble Kalman inversion method, and uncertainty quantification is enhanced by an emulate–sample strategy that leverages machine-learning for surrogate modeling, in which a Gaussian-process surrogate of the parameter-to-statistics map enables efficient MCMC sampling without further expensive forward simulations. Synthetic experiments with realistic data noise demonstrate that the model parameters can be inverted from statistical quantities observable from paleoseismology at a limited number of observation points. The framework provides a principled and computationally tractable method to calibrate earthquake-cycle models, providing a foundation for forecasting.



## Plain Language Summary

Physics-based models of the earthquake cycle simulate how stress builds up on faults and is released during earthquakes. These models could improve earthquake hazard assessment, but only if their physical parameters can be estimated from observations. This is challenging because records of past earthquakes are sparse, uncertain, and span only a small portion of the earthquake cycle. Paleoseismic data provide information about the timing and size of past earthquakes, but these observations are noisy and incomplete.

In this study, we develop a method to estimate earthquake model parameters using statistical properties of past earthquakes, such as average recurrence time and slip. Instead of matching individual earthquake sequences, we match long-term statistical behavior and quantify uncertainty using Bayesian methods and machine learning. Our results show that fault properties can be constrained even with sparse paleoseismic data, providing a step toward physics-based earthquake forecasting.

## 1 INTRODUCTION

A central challenge in earthquake-cycle modeling is the estimation of physical parameter values that appear in the governing equations of a physical simulation model and control long-term fault behavior. Dynamical modeling of earthquake sequences requires solving Partial Differential Equations (PDE) describing the motion resulting from frictional sliding on a fault embedded in an elastic medium, together with a frictional constitutive law (Rice, 1993; Lapusta et al., 2000; Chen and Lapusta, 2009). The predictive power and interpretation of these models depend critically on parameter calibration: simulated sequences must reproduce key observables of long-term fault behavior, including recurrence intervals and coseismic slip (Barbot et al., 2012). Achieving this requires tuning model parameters using observations that include a sufficient number of large events to yield meaningful statistics and therefore span decadal-to-millennial timescales. Because such timescales far exceed the instrumental and written historical record, the primary constraints on long-term earthquake recurrence come from paleoseismic observations. Paleoseismic investigations are generally carried out in trenches excavated across faults and provide information on the age and eventual slip produced by past earthquakes (Burbank and Anderson, 2013; Wang et al., 2024). These observations are sparse and noisy, yet provide informative constraints on event timing and slip. Further complicating this process is the chaotic nature of earthquake dynamics, as each event depends sensitively on the unknown initial condition of the dynamical model.

Several approaches have been developed to constrain parameters in physical models of earthquake occurrence and the earthquake cycle. The rate-and-state formulation derived from friction experiments provides a well-established framework linking the stress history to the earthquake

rate (Dieterich, 1994; Heimisson and Segall, 2018). This approach has been validated in particular using observational data in induced seismicity settings (Kaveh et al., 2023). However, this model only predicts the seismicity rate and not the occurrence of individual event so its application to time-dependent earthquake forecasting is limited. To overcome this limitation, approaches that explicitly model the spatiotemporal evolution of fault slip over the earthquake cycle are required.

Using a dynamical model of earthquake, Kaveh et al. (2025b) demonstrated that large events can be forecast as a consequence of the self-organization of the stress field resulting from prior ruptures. In a follow-up article Kaveh et al. (2025a) addressed the data assimilation problem: they inverted for the slip rate and stress on the fault and leveraged machine learning to learn the sequence of slow slip events, thereby accelerating the forward model. In that study, they also showed that future events can be predicted using noisy observations of slip rate. Both of these works, however, assumed full knowledge of the model parameters. In the present work, we remove this limitation by developing a methodology to constrain the model parameters using observations that can, in principle, be derived from historical and paleoseismological records.

Recent studies on finding parameters of earhtquake models have focused on the joint estimation of the model parameters and initial state. In particular, Stiernström et al. (2024) introduced an adjoint-based inversion framework that extends full waveform inversion concepts to earthquake-cycle models, enabling joint estimation of model parameters and the initial stress state through gradient-based optimization when event-specific observations and information about the system state are available. Fukushima et al. (2025) proposed a physics-informed deep learning framework to invert for spatially heterogeneous frictional parameters in slow slip event (SSE) regions using dense geodetic time series, demonstrating that heterogeneity can be constrained when continuous, high-resolution observations are available.

The approaches of (Stiernström et al., 2024; Fukushima et al., 2025; Riley et al., 2025) are well-suited for calibrating model parameters using observations of individual rupture dynamics or densely sampled system trajectories. However, these methods are less applicable when observations are sparse and have large uncertainties, as is typically the case with the time and slip of past large earthquake ruptures estimated from paleoseismology. In particular, when continuous high-

resolution observations are unavailable, data-driven or neural-network-based calibration strategies lack the trajectory-level information required for effective training. Moreover, in strongly chaotic earthquake-cycle models, constructing adjoint sensitivities over timescales that exceed the system's predictability horizon becomes uninformative. In addition, joint inversion for the initial condition and model parameters is a requirement for this process, and becomes infeasible for chaotic earthquake-cycle dynamics when the system trajectory has evolved over several thousand years. For physics-based earthquake forecasting, however, a model need not reproduce the detailed dynamics of any single rupture; rather, it must reproduce the long-term statistical behavior of the fault, such as recurrence intervals, variability, and average coseismic slip. Consistent with this perspective is the work of Corbi et al. (2024), who combined laboratory "foamquake" experiments with numerical inversions to constrain asperity properties and frictional parameters using statistical features of earthquake cycles. Their study highlights the importance of statistical observables for constraining recurrence behavior when trajectory-level information is unavailable. The use of statistical observables to constrain parameters of chaotic dynamical systems, rather than trajectory matching, has a long history in the climate dynamics and time-series communities (Lütkepohl, 1991; Neumaier and Schneider, 2001; Brockwell and Davis, 2016; Schneider et al., 2021; Dunbar et al., 2021a), but has been explored only to a limited extent in the context of earthquake-cycle models.

In this paper, we formalize this statistical perspective within a Bayesian inverse-problem framework for estimating heterogeneous parameters of earthquake-cycle models from paleoseismic observations. Rather than performing joint inversion for both model parameters and the unknown initial condition, we calibrate parameters using stationary statistical observables of the system, such as the mean interevent time, the coefficient of variation, burstiness, and the mean coseismic slip, which are quantities that can be inferred from paleoseismic records (Griffin et al., 2020). This shift from trajectory matching to stationary-statistics matching fundamentally changes the inverse problem: it removes the need to infer the initial condition and reframes chaos as a source of sampling variability rather than instability. Within this framework, we obtain rigorous uncertainty quantification, incorporate spatially distributed statistics from multiple observation sites,

representing paleoseismic trenching sites, and provide a principled treatment of observational uncertainty, including noise in event dating, slip-measurement noise, and incomplete detectability of small events.

The main contributions of this work are: (i) a Bayesian inverse-problem framework for calibrating heterogeneous parameters of earthquake-cycle models using stationary statistical observables derived from paleoseismic records; (ii) a formulation that exploits ergodicity to remove the need to estimate the unknown initial condition of the dynamical system; and (iii) a computational strategy that uses machine learning to enable efficient uncertainty quantification. The results demonstrate that paleoseismic-scale observations can meaningfully constrain heterogeneous frictional parameters despite noise and chaos, establishing a necessary calibration step for physics-based earthquake-cycle models and laying the groundwork for subsequent forecasting studies.

The remainder of the paper is organized as follows. In Section 2 we introduce the inverse-problem formulation and describe the earthquake-cycle forward model, the parameterization of heterogeneous frictional properties, and the construction of paleoseismic observables. We then present the Bayesian calibration methodology in Section 3, including the Ensemble Kalman Inversion framework and the machine learning strategy for uncertainty quantification. The inversion results and posterior analyses are presented in Section 4 and Section 5, followed by conclusions and perspectives for future work.

## 2 SETTING UP THE INVERSE PROBLEM

We work in a controlled synthetic setting, generating noisy observations of event timing and slip from numerical simulations, incorporating realistic dating and measurement noise, and accounting for the limited detectability of small events. This framework still enables a critical assessment of loss functions, parameter identifiability, and algorithmic feasibility, without the confounding challenges of model–data misspecification. From these records, we construct spatially distributed statistical observables at multiple trench-like locations along the fault, including mean interevent times, coefficients of variation, interevent correlations, and average coseismic slip.

We consider an earthquake-cycle model parameterized by a vector $u \in \mathbb{R}^{d_u}$, where $d_u$ de-

notes the number of unknown parameters. Although these parameters may vary spatially along the fault, we represent them through a finite-dimensional parameterization, so that the full spatial structure is described by a finite set of coefficients in $\mathbb{R}^{d_u}$. The observational data are collected in a vector $y \in \mathbb{R}^{d_y}$, which consists of summary statistics computed from paleoseismic records at trench-like locations along the fault. For intuition, one may think of these as quantities such as the average interevent time measured at $d_y$ trench sites. Because earthquake-cycle dynamics are often chaotic, trajectory-level predictions are highly sensitive to initial conditions; however, their long-term statistical properties are empirically stationary over sufficiently long evolutions. For large earthquakes, only finite and relatively short observation windows are available, typically limited to the Holocene ($\lesssim 10^4$ years). As a result, statistics estimated from paleoseismic records should be regarded as noisy approximations of their long-term stationary values. We exploit ergodicity to interpret these finite-time estimates as noisy samples of the underlying stationary statistics, with uncertainty arising from both observational noise and unknown initial conditions. In Section 2.1 we formulate the relationship between the parameters $u$ and the observable statistics $y$. This formulation combines the earthquake-cycle dynamical model, described in Section 2.2, with an observation operator introduced in Section 2.3, which defines the statistical quantities extracted from paleoseismic records. In Section 2.4 we explain why it is not necessary to invert for the unknown initial condition. Instead, by focusing on statistical observables, the uncertainty arising from the unknown initial condition can be treated as part of the observational uncertainty in the inverse problem.

## 2.1 Exploiting Stationary Statistics

The formulation of the inverse problem to estimate parameters relies on a structural assumption about the long-term behavior of the earthquake-cycle model. Specifically, we assume that the dynamical system exhibits stationary long-term statistics over long timescales. In other words, although individual earthquake sequences are sensitive to initial conditions (such as the initial stress distribution), the statistical properties of the sequence, such as the average interevent time

or the average coseismic slip at a given location, converge, over sufficiently long time horizons, to values that are independent of the initial state.

This assumption is consistent with the notion of ergodicity. For an ergodic dynamical system, time averages of suitable observables converge to deterministic values that depend only on the system parameters and not on the initial condition. To illustrate this idea, let $h_i$ denote a sequence of observations of a quantity of interest (for example, interevent times or coseismic slip measured at a trench site). Then the long-time average

$$\lim_{N\to\infty} \frac{1}{N} \sum_{i=1}^{N} h_i$$

is independent of where the sequence begins in the earthquake cycle and is unaffected by individual realizations of measurement noise. The limit is a deterministic quantity determined solely by the underlying model parameters.

Motivated by this perspective, we denote by $\mathcal{G}_\infty(u) \in \mathbb{R}^{d_y}$ the vector of infinite-time averaged summary statistics produced by the earthquake-cycle model with parameters $u$. In practice, however, neither nature nor numerical simulation provides access to infinite-time averages. Paleoseismic records extend over only a finite time window $T$, and therefore the statistics computed from such records are finite-time estimates of their infinite-time limits. Let $G(u;T) \in \mathbb{R}^{d_y}$ denote the summary statistics computed from a trajectory of the earthquake-cycle with parameters $u$, evolved over only a finite time horizon $T$ starting from a random initial condition. For notational convenience, when the dependence on the observation window is clear from context, we write $G(u)$ in place of $G(u;T)$. Unlike $\mathcal{G}_\infty(u)$, the quantity $G(u;T)$ is a random variable. Its randomness arises from two sources: (i) the unknown initial condition of the dynamical system (when we start gathering paleoseismic records in the earthquake cycle), and (ii) realizations of observational noise in the measurement of event timing and slip.

We assume that the observational noise in event dating and slip measurements are independent and identically distributed with zero mean across detected events. Because each component of $G(u;T)$ is constructed as an average over events occurring within the time window $[0,T]$, the contribution of measurement noise to $G(u;T)$ converges to zero as $T \to \infty$ by the law of large num-

bers. Consequently, the limiting map $\mathcal{G}_\infty(u)$ is deterministic and independent of both the initial condition (by ergodicity of the dynamical system) and the particular realizations of measurement noise (by the law of large numbers).

$$\lim_{T\to\infty} G(u;T) = \mathcal{G}_\infty(u).$$

We therefore interpret the statistics found from finite time observation ($y = G(u;T)$) as a noisy evaluation of $\mathcal{G}_\infty(u)$, and write

$$y = \mathcal{G}_\infty(u) + \eta_T, \tag{1}$$

where $\eta_T \in \mathbb{R}^{d_y}$ represents a noise with covariance matrix $\Gamma_T$ which encapsulates the contribution of the unknown initial condition and measurement noise. The magnitude of $\eta_T$ decreases as $T$ increases, consistent with ergodic averaging and the law of large numbers. The statistics (for example, interevent time or average coseismic slip) obtained from paleoseismic observations over a time window of length $T$ are a single realization of the random model described in Eq. 1. We denote this single realization by $y^\dagger$. Now the goal is to estimate the parameters ($u$) from $y^\dagger$. In particular, we want to characterize the posterior distribution of the parameters given the noisy observation $y^\dagger$, denoted by $\pi(u \mid y^\dagger)$. By Bayes' rule, the posterior is proportional to the likelihood times the prior,

$$\pi(u \mid y^\dagger) \propto \ell(y^\dagger \mid u)\, \rho(u), \tag{2}$$

where $\ell(y^\dagger \mid u)$ is the likelihood and $\rho(u)$ is the prior distribution on the model parameters. The likelihood takes the form

$$\ell(y^\dagger \mid u) \propto \exp\left(-\tfrac{1}{2}\|y^\dagger - \mathcal{G}_\infty(u)\|^2_{\Gamma_T}\right),$$

where $\|x\|^2_{\Gamma_T} := x^\top \Gamma_T^{-1} x$. The corresponding negative log-likelihood up to a normalization constant is

$$\Phi(u) = \tfrac{1}{2}\|y^\dagger - \mathcal{G}_\infty(u)\|^2_{\Gamma_T}. \tag{3}$$

Although inference is formally posed in terms of the idealized map $\mathcal{G}_\infty$, in practice we only access noisy evaluations $G(u;T)$ as discussed below. Direct Monte Carlo sampling of Eq. 2 is com-

putationally prohibitive, as a single evaluation of the forward map $\mathcal{G}_\infty(u)$ requires running the earthquake-cycle model over a very long time and can take days to months. To address this challenge, we adopt a recently developed strategy for approximating the posterior distribution, known as the Calibrate–Emulate–Sample (CES) framework (Cleary et al., 2021). This method has several desirable attributes. First, it does not require evaluate of$\mathcal{G}_\infty(u)$ and works only with noisy evaluation of it namely $G(u,T) = \mathcal{G}_\infty(u) + \eta_T$; second, it works with fewer evaluations of $G(u,T)$ compared to traditional Monte Carlo samplers. This approach enables efficient approximate sampling from the posterior and is described in Sections 3.1 and 3.2.

Nevertheless, evaluating $G(u,T)$ for some parameters $u$ to estimate the likelihood is necessary. In principle, to reduce computational cost, the model could be simulated over a shorter time horizon than $T$; however, this would require appropriately rescaling the noise term, since statistics computed from shorter time windows exhibit larger variability. In this paper we do not adopt this strategy and instead simulate the model up to time $T$. To get a realization of $G(u,T)$ one does not need to evaluate $\mathcal{G}_\infty(u)$ because $G(u,T)$ requires only the finite time statistics of a forward model with parameters $u$ and a random initial condition. In fact, it is instructive to think of $G$ as a composition of two functions the forward model map $\psi$ and an observation functional $H$, i.e.,

$$G = H \circ \psi.$$

The map $\psi$ which is an earthquake simulator takes the parameter vector $u$ and, starting from a random initial condition, evolves the earthquake-cycle model over a time horizon $T$ and produces the resulting time history. The functional $H$ maps the output of the earthquake simulator to paleoseismic-style observables by extracting informative summary statistics, including interevent time, coefficient of variation (COV), Pearson correlation coefficient (PCC), and mean coseismic slip at trench-like locations (Griffin et al., 2020). These statistics are described in Section 2.3. The mapping also incorporates randomness to represent measurement noise. In the following subsections, we discuss in detail each component of the inverse problem in Eq. 1: the parameters of interest $u$, the forward dynamical model $\psi$, and the observational model $H$.

### 2.2 Dynamical model ($\psi$) and parameters of interest ($u$)

To solve the inverse problem defined in Eq. 1, we require an earthquake simulator that contains the unknown model parameters. In this work, we model the spatiotemporal evolution of slip on a fault embedded in an elastic medium using the rate-and-state friction law with the aging evolution law Ruina (1983). Although we focus on this specific formulation, the inversion methodology is general and can be applied to other models of the earthquake cycle. The evolution of the fault system is assumed to be governed by

$$\partial_t v = \left[\kappa + \frac{a\,\bar{\sigma}}{v}\right]^{-1}\left[\mathcal{L}\big(v - v_{\mathrm{pl}}\big) - b\,\bar{\sigma}\Big(\frac{1}{\theta} - \frac{v}{d_{\mathrm{rs}}}\Big)\right], \qquad (z,t) \in \Gamma \times (0,\infty), \tag{4a}$$

$$\partial_t \theta = 1 - \frac{v\theta}{d_{\mathrm{rs}}}, \qquad (z,t) \in \Gamma \times (0,\infty), \tag{4b}$$

Here $v(z,t)$ denotes the slip rate and $\theta(z,t)$ is the state variable representing the memory of contact on the fault surface. The operator $\mathcal{L}$ denotes the static elastic stress transfer along the fault. To reduce computational cost, we use a quasidynamic approximation (Rice, 1993). Stress transfer are calculated with static elastic kernels rather than fully dynamic elastic kernel (Lapusta et al., 2000), which would be significantly more expensive to simulate. The parameter $\kappa = \mu/(2c_s)$ is the radiation damping coefficient, introduced to account for the energy radiated by seismic waves (Rice, 1993), where $\mu$ is the shear modulus and $c_s$ is the shear-wave speed; both are treated as known physical parameters. The coefficients $a$, $b$, and $d_{\mathrm{rs}}$ are frictional parameters that are constant in time but may vary spatially; in this work they are assumed to be piecewise constant. The effective normal stress $\bar{\sigma}$ is likewise assumed to be piecewise constant.

We consider a one-dimensional fault embedded in a two-dimensional elastic medium, with $\Gamma = [-L/2,\, L/2]$. The fault is assumed to be infinite in extent, but only a segment of length $L$ is governed by the rate-and-state friction law; the remainder slips steadily at the prescribed plate loading rate $v_{\mathrm{pl}}$. For certain parameter regimes, the system described by Eq. 4 exhibits multiscale, complex, and potentially chaotic behavior in both space and time.

In this study we perform a synthetic experimental set-up in which the data is derived from the model whose parameters we wish to learn. We prescribe the true spatial distribution of the

frictional parameters $(a\bar{\sigma},\ b\bar{\sigma},\ d_{\mathrm{rs}})$ and use synthetic, noisy paleoseismic data to infer their distribution. The true spatial profiles of these frictional properties are shown in Fig. 1. The model setup is inspired by (Kaneko et al., 2010), but with stronger spatial heterogeneity. Given these properties and a random initial condition, Eq. 4 is simulated forward in time to generate a long-term sequence of synthetic earthquake events. An example of the resulting spatiotemporal earthquake sequence produced by the model is shown in Fig. 2. From this simulated catalog we then construct synthetic historical (paleoseismic) observations subject to observational noise. The construction of these synthetic observations is described in the next subsection.

When inverting for the parameters $(a\bar{\sigma}, b\bar{\sigma}, d_{\mathrm{rs}})$ of model Eq. 4, one typically does not know the spatial structure of the frictional properties: they could take many forms—piecewise constant, polynomial, or parameterized through polynomial or GP models, and so on. In this work, however, we make a simplifying assumption, assuming that they vary in a piecewise constant manner. Specifically, we assume that the fault consists of five distinct segments and that the sign of $a - b$ is known in each segment. Such information may be obtained from interseismic coupling maps derived from geodetic measurements (Avouac, 2015).

The unknown parameters $u$ consist of the values of $(a\bar{\sigma}, b\bar{\sigma}, d_{rs})$ in the five segments of the fault, together with the lengths of the velocity-weakening (VW) segments. The parameter vector $u$ is illustrated in Fig. 3. We assume that the lengths of the two outer velocity-strengthening (VS) segments on the left and right, denoted $L_{VS}^{(1)}$ and $L_{VS}^{(3)}$, are known. These segments exhibit no paleoseismic slip, so their extents can be inferred from historical observations. In contrast, we take the lengths of the remaining segments as unknowns: the left velocity-weakening (VW) segment $L_{VW}^{(1)}$, the central VS segment $L_{VS}^{(2)}$, and the right VW segment $L_{VW}^{(2)}$. The fault is therefore partitioned into five piecewise-constant regions: three velocity-strengthening segments ($VS1$, $VS2$, $VS3$, shown in green in Fig. 3) and two velocity-weakening segments ($VW1$ and $VW2$, shown in blue). Each segment is assigned distinct frictional properties.

In fact, there are only two unknown length segments because of known overall length constraint $L$. We treat $L_{VW}^{(1)}$ and $L_{VW}^{(2)}$ as unknown parameters and determine the remaining central VS

segment length by

$$L_{VS}^{(2)} = L - \left(L_{VS}^{(1)} + L_{VS}^{(3)} + L_{VW}^{(1)} + L_{VW}^{(2)}\right).$$

In total, the number of unknown parameters in our synthetic test is $d_u = 17$: two segment-length parameters and fifteen frictional parameters $(a\bar{\sigma}, b\bar{\sigma}, d_{\mathrm{rs}})$ across the five segments. The vertical red ticks in Fig. 3 are the locations of observational points (trenches). It is worth noting that in the outer VS segments ($VS1$ and $VS3$), only a single observational trench exists, and these locations exhibit the longest interevent times. Moving farther into these segments, the recurrence intervals increase until no events are observed in the historical data. This behavior supports the practical identifiability of $L_{VS}^{(1)}$ and $L_{VS}^{(3)}$ from historical data, as their boundaries correspond to the onset of zero recorded coseismic slip.

We denote by $\psi$ the dynamical model operator. Given model parameters $u$, the map $\psi(u; T)$ evolves Eq. 4 from a random initial condition and returns the system trajectory over a finite time horizon $[0, T]$. In practice, the random initial condition is not used directly. Instead, the model is first integrated over a spin-up period, after which the resulting state is taken as the effective initial condition for the simulation over $[0, T]$. This procedure ensures that the trajectory begins from a dynamically consistent state representative of the long-term behavior of the system; formally, it approximates a sample from the invariant (stationary) measure of the earthquake-cycle dynamics.

### 2.3 Observational model ($H$)

This part describes how the observable vector $y$ is constructed from the trajectory of the system over a finite time horizon $[0, T]$. We also describe how observational noise in event timing and slip measurements, as well as the limited detectability of events, is incorporated into the construction of these observables. As stated earlier, we use only statistical summaries of the observations rather than the full sequence of individual event times or slip values at specific locations. This choice allows us to invert solely for the model parameters, without needing to infer the unknown initial condition of the dynamical system. This comes at the cost of introducing an additional source of uncertainty arising from the unknown initial condition, which is explained in the next subsection.

The statistical quantities we use at each observational point (trench) include the mean in-

terevent time, the coefficient of variation, the Pearson correlation coefficient, and the average coseismic slip. To define these observables precisely, we first introduce the underlying time-series quantities from which they are derived. We also explicitly account for observational uncertainties, since both event-timing data (due to dating errors) and slip measurements (due to stratigraphic and instrumental noise) are inherently noisy.

We first define the cumulative slip function $c(z,t) : \Gamma \times [0,T] \mapsto \mathbb{R}^+$ as:

$$c(z,t) := \int_0^t v(z,t')\, dt', \tag{5}$$

where $v(z,t')$ represents the slip rate at location $z$ on the fault at time $t'$.

We assume that there are $n_{\mathrm{t}}^{\mathrm{o}}$ observational points on the fault where we observe only the *timing* of past events. We denote the set of these locations by $Z_{\mathrm{t}}^{\mathrm{o}}$. We further assume we have access to slip measurements at $n_{\mathrm{s}}^{\mathrm{o}}$ points on the fault, and we denote the set of all slip observational points by $Z_{\mathrm{s}}^{\mathrm{o}}$. Typically, $n_{\mathrm{t}}^{\mathrm{o}} \geq n_{\mathrm{s}}^{\mathrm{o}}$, and the temporal information tends to be less noisy than the slip measurements and extend further in time. In this paper, for simplicity, we take $n_{\mathrm{t}}^{\mathrm{o}} = n_{\mathrm{s}}^{\mathrm{o}} = 20$ and $Z_{\mathrm{s}}^{\mathrm{o}} = Z_{\mathrm{t}}^{\mathrm{o}}$. The formulation, however, applies to the general case in which $Z_{\mathrm{s}}^{\mathrm{o}} \neq Z_{\mathrm{t}}^{\mathrm{o}}$.

For each $z \in Z_{\mathrm{t}}^{\mathrm{o}} \cup Z_{\mathrm{s}}^{\mathrm{o}}$ we define the sets $T^s(z)$, $T^e(z)$, and $S(z)$ containing start, end and amount of slip of each event that point $z$ has experienced. Importantly, elements of $T^s(z)$ and $T^e(z)$ come in pairs: each $t_i \in T^s(z)$ corresponds to a unique $t_i' \in T^e(z)$, and coseismic slip values in $S(z)$ are constructed from these matched pairs. We define these quantities by using a threshold on the slip rate $v_{\text{thresh}}$ at that location.

$$T^s(z) := \{t \mid v(z,t) = v_{\text{thresh}}, v(t+\epsilon) > v(t) \text{ for } \epsilon \text{ small } (v(z,t) \text{ increasing})\} \tag{6}$$

$T^s(z)$ is the set of starting times of earthquakes that a point $z$ has experienced.

$$T^e(z) := \{t \mid v(z,t) = v_{\text{thresh}}, v(t+\epsilon) < v(t) \text{ for } \epsilon \text{ small } (v(z,t) \text{ decreasing})\}, \tag{7}$$

$T^e(z)$ is the set of the ending times of earthquakes at a point $z$.

$$S(z) = \{c(z,t_i') - c(z,t_i) \mid i \in I(z), t_i' \in T^e(z), t_i \in T^s(z)\}, \tag{8}$$

where $I(z)$ is the indexing set for $T^s(z)$, and $S(z)$ is the set of coseismic slips that a point $z$

has experienced. Not all of the events in these sets are detectable using paleoseismic data. To account for this, we introduce two distinct slip thresholds: $s_{\text{thresh}}$ denotes the minimum coseismic slip required for an event to appear in timing records, and $s'_{\text{thresh}}$ denotes the minimum coseismic slip required for an event to appear in slip records. In this paper, we take $s_{\text{thresh}} = s'_{\text{thresh}}$, but we present the problem for the general case when they are different.

Now, for the points $z \in Z_{\text{t}}^{\text{o}}$ we define:

$$T^o(z) := \{t_i + \zeta_i \mid i \in I, t_i \in T^s(z), s_i > s_{\text{thresh}}, s_i \in S(z), \zeta_i \sim \mathcal{N}(0, \sigma^2_{dating})\}, \tag{9}$$

where $\sigma^2_{dating}$ is the variance of the dating noise of the historical record (e.g., carbon dating). We assume independence between dating noise and slip-measurement noise, and also independence across events. This assumption can, however, be relaxed, as older paleoseismic observations are often noisier than more recent ones. $T^o(z)$ is the noisy observed dates for each point $z \in Z_{\text{t}}^{\text{o}}$ on the fault and for events with coseismic slip bigger than $s_{\text{thresh}}$.

We define the set of interevents $(\Delta T^o(z))$ for the point $z \in Z_t^o$ as follows:

$$\Delta T^o(z) := \{t_{i+1} - t_i | 1 \leq i \leq m, t_i \in T^o(z)\}, \tag{10}$$

where $m$ is defined as the number of elements in $T^o(z)$ minus one. For each $z$, we define the mean empirical interevent time $\overline{\Delta t}(z)$ as the average of the elements in the set $\Delta T^{\text{o}}(z)$, and the empirical standard deviation $\sigma_{\Delta t}(z)$ as the standard deviation of the same elements. From these, we define the coefficient of variation at location $z$ as

$$\text{COV}(z) = \left(\frac{\sigma_{\Delta t}(z)}{\overline{\Delta t}(z)}\right)\left(1 + \frac{1}{4m}\right), \tag{11}$$

This adjusted formula provides an approximately unbiased estimator of the population coefficient of variation $\sigma/\mu$, assuming the elements in $\Delta T^{\text{o}}(z)$ are drawn from a normally distributed population (Subrahmanya Nairy and Aruna Rao, 2003). If $\text{COV} = 0$, the signal is perfectly periodic. If $\text{COV} = 1$, the events are memoryless, as in a Poisson process. Values of $\text{COV} > 1$ indicate clustering of events.

We further wish to quantify the burstiness of the events. Define

$$\tau(z) := \{\Delta t_i \in \Delta T^o(z) \mid i \neq m \text{ (all elements of } \Delta T^{obs}(z) \text{ except the last one)}\}, \tag{12}$$

$$\tau_+(z) := \{\Delta t_i \in \Delta T^o(z) \mid i \neq 1 \text{ (all elements of } \Delta T^{obs}(z) \text{ except the first one)}\}, \tag{13}$$

as the lagged vectors of interevent times. Then, the burstiness is defined as the *Pearson sample Correlation Coefficient* (PCC) between consecutive interevent times at location $z$

$$\mathrm{PCC}(z) := \frac{\sum_{i=1}^{m-1}\left(\tau_i - \overline{\tau}\right)\left(\tau_i^+ - \overline{\tau}_+\right)}{\sqrt{\sum_{i=1}^{m-1}\left(\tau_i - \overline{\tau}\right)^2}\sqrt{\sum_{i=1}^{m-1}\left(\tau_i^+ - \overline{\tau}_+\right)^2}}, \tag{14}$$

where $\overline{\tau}$ and $\overline{\tau}_+$ are the sample means of $\tau(z)$ and $\tau_+(z)$, respectively.

Now that we have defined the statistical quantities of interest (average interevent time, coefficient of variation, and burstiness) for timing of events and accounted for the observational noise in those statistics, we next define the statistics of coseismic slip observation. Using the definition of $S(z)$, at each point $z \in Z_s^o$, we define the observed slip:

$$S^o(z) := \{s_i(z) + \gamma_i \mid s_i(z) \in S(z), s_i(z) > s'_{thresh}, \gamma_i \sim \mathcal{N}(0, \sigma^2_{slip})\}, \tag{15}$$

where $s'_{thresh}$ is a slip threshold above which we can observe an effect in the slip paleoseismic record and $\sigma_{slip}$ is the standard deviation of slip measurement noise. For each $z \in Z_s^o$ we take the average of elements in $S^o(z)$ and construct the $\overline{S^o(z)}$. Paleoseismic slip measurements are often sparse and noisy. For this reason, we make use only of the average coseismic slip in the inversion and do not attempt to incorporate higher-order statistics of the slip measurements.

We concatenate the observations from all trench locations to form the full observational vector. Namely, we define $\overline{\Delta t}(Z_t^o) \in \mathbb{R}^{n_t^o}$ as the vector of mean interevent times, $\mathrm{COV}(Z_t^o) \in \mathbb{R}^{n_t^o}$ as the coefficients of variation, $\mathrm{PCC}(Z_t^o) \in \mathbb{R}^{n_t^o}$ as the Pearson sample correlation coefficients of interevent times, and $\overline{S^o}(Z_s^o) \in \mathbb{R}^{n_s^o}$ as the average observed coseismic slip at the corresponding slip-observation locations. We define the observational vector to be:

$$y = \left[\overline{\Delta t}(Z_{\mathrm{t}}^{\mathrm{o}}),\ \mathrm{COV}(Z_t^o), PCC(Z_t^o), \overline{S^o}(Z_s^o)\right] \in \mathbb{R}^{3n_{\mathrm{t}}^{\mathrm{o}}+n_{\mathrm{s}}^{\mathrm{o}}}. \tag{16}$$

At observation locations where three or fewer paleoseismic event times are available (i.e., fewer than two interevent intervals), we consider the number of events insufficient to compute meaning-

ful interevent statistics. In these cases, we define the mean interevent time as the total duration of the simulation, and we set $\mathrm{COV} = 1$ and $\mathrm{PCC} = 0$.

We illustrate the construction of the synthetic paleoseismic observation in Fig 4 from which we extract desired statistics. Panels (a)-(b) correspond to noiseless records, while panels (c)-(d) include dating noise in event times and slip-measurement noise in coseismic slips. In each row, the left panel shows only the occurrence times of events (black markers), whereas the right panel shows the same events colored by the associated coseismic slip. This comparison highlights the impact of measurement noise on both the timing and slip components of paleoseismic observations. In this paper, we use statistics derived from noisy event timing and slip data, of the type illustrated in Fig. 4(d).

We conduct a synthetic experiment, in which we define a $u^{true}$, and we make a synthetic noisy historical record ($y^{\dagger}$) where the noise comes from not knowing the initial condition and also the observational noise. In this study, using true parameters plotted in Fig 1, and presented in Tables 1 and 2, we simulate dynamical model (Eq. 4) for $5500$ years and remove the first $500$ years ($T = 5000$ years). We set the paleoseismic event detection thresholds $s_{thresh} = s'_{thresh} = 0.8(m)$; events that have coseismic slip greater than $0.8(m)$ at the locations of a trench are observed in the synthetic data set. We use $n_t^o = n_s^o = 20$, the synthetic trenches are denoted by vertical red ticks in Fig 3. We use $\sigma_{dating} = 10(year)$ as the standard deviation of timing measurement noise and $\sigma_{slip} = 0.4(m)$ as the standard deviation of the slip measure noise. Fig 5 plot the noisy observation that is used in this study to invert for parameters.

With this, the observation functional $H$ is defined as a map that takes the full slip-rate trajectory $v(t, z)$, for $t \in [0, T]$ and $z \in \Gamma$, and computes the statistics defined in Eqs. 5–16. For systems with stationary statistics as $T \to \infty$, the output of the map $G(u; T) = H \circ \psi(u; T)$ is independent of observational noise and the unknown initial condition. For finite-time statistics, however, the map $G(u; T)$ can be viewed as a random map whose expectation equals the infinite-time limit of the operator ($\mathcal{G}_\infty(u)$).

## 2.4 Why the Initial Condition Need Not Be Inferred

A common approach to parameter estimation in dynamical systems is *state augmentation*, in which unknown parameters are appended to the state vector and estimated jointly with the evolving system trajectory. This strategy has a long history in data assimilation (Anderson, 2001; Doucet et al., 2001; Julier et al., 2000; Evensen, 2018) and can be effective when dense, time-resolved observations of the system state are available.

For earthquake-cycle models over paleoseismic timescales, however, state augmentation is structurally ill-suited. The governing dynamics are nonlinear and often lead to chaotic dynamics (Barbot, 2019), so trajectory-level predictions are highly sensitive to initial conditions. Over time horizons of thousands of years, small perturbations in the initial stress state grow exponentially, rendering the trajectory matching unstable and effectively non-identifiable. When observations are limited to sparse and noisy event times and slip measurements, the data contain insufficient information to jointly recover both the initial condition and heterogeneous model parameters. In such settings, state-augmentation methods are prone to degeneracy, ensemble collapse, or spurious parameter compensation.

Rather than attempting to reproduce a specific earthquake sequence, we instead target *stationary statistical properties* of the system. As described in Subection 2.1, we assume that the earthquake-cycle model admits stationary long-term statistics for the observables of interest. Under this assumption, there exists a deterministic mapping

$$\mathcal{G}_\infty(u) = \mathbb{E}\big[G(u;T)\big],$$

where $G(u;T)$ is the finite time average of statistics, and the expectation is taken with respect to the stationary process induced by the dynamics and observational noise. The map $\mathcal{G}_\infty(u)$ depends only on the model parameters and not on the initial condition and realizations of noise. Appendix A provides a pedagogical illustration of finite-time averaging in a chaotic system using the Lorenz–63 equations.

The assumption of stationary statistics is supported by the dissipative nature of earthquake-cycle models. Dissipative dynamical systems possess a compact global attractor that contains the

long-term dynamics of the system (Temam, 1997). Regardless of the initial stress state, trajectories eventually converge toward this attracting set and subsequently evolve within it. On this attractor, the dynamics admit invariant probability measures that describe the long-term statistical behavior of the system. Consequently, statistical observables computed over sufficiently long time intervals become independent of the initial condition and depend only on the model parameters. This provides the dynamical justification for calibrating model parameters using stationary statistics rather than attempting to reproduce a specific earthquake sequence.

In practice, paleoseismic records span a finite time window $[0, T]$. Statistics computed from such records therefore fluctuate around their infinite-time values due to finite sampling and observational noise. For fixed parameters $u$, the finite-time quantity $G(u; T)$ is thus a random variable whose variability reflects both the unknown initial condition and measurement noise. Within this framework, the initial condition enters only through the sampling variability of $G(u; T)$ and is absorbed into the effective noise term of the inverse problem (Eq. 1).

The key modeling decision of this work is therefore to invert for parameters using stationary statistics rather than trajectories. Under the stationarity assumption, the inverse problem can be posed without specifying or estimating the initial condition explicitly. The unknown initial state contributes to uncertainty ($\eta$ in Eq. 1), but it does not constitute an additional parameter to be inferred.

To illustrate the stochastic nature of the finite-time map $G(u; T)$, we fix $u = u^{\mathrm{true}}$ and generate multiple realizations of the random variable $G(u^{\mathrm{true}}; T)$ in Fig. 6. Recall that

$$G(u^{true}; T) = H \circ \psi(u^{true}; T),$$

where $\psi$ denotes the dynamical simulator and $H$ the observational functional. To sample $G(u^{\mathrm{true}}; T)$, we proceed as follows. Within $\psi$, the model is initialized from a random initial condition and evolved for a spin-up period of $500$ years to ensure convergence toward its stationary (ergodic) regime. From the resulting state, the system is integrated forward to $T = 5000$ years. The functional $H$ is then applied to extract paleoseismic-style summary statistics. Randomness enters both

through the random initial condition in $\psi$ and through observational noise in $H$, including dating uncertainty and slip-measurement noise.

We generate $120$ independent realizations of $G(u^{\mathrm{true}}; T)$. These are shown as black curves in Fig. 6. The red curve denotes their empirical mean, which serves as an empirical approximation of the infinite-time deterministic map $\mathcal{G}_\infty(u^{\mathrm{true}})$. In the inverse problem, however, we observe only a *single* realization of this random variable. Denoting this realization by $y^\dagger$ (one of the black curves in Fig. 6), our objective is to characterize the conditional distribution $\pi(u \mid y^\dagger)$. The specific realization used in the inversion is shown in Fig. 5.

## 3 METHOD

We solve the inverse problem in Eq. 1 using Ensemble Kalman inversion (EKI) (Chen and Oliver, 2012; Iglesias et al., 2013; Schillings and Stuart, 2016; Chada et al., 2021; Huang et al., 2022; Dunbar et al., 2022b; Gjini et al., 2025), a family of derivative-free, ensemble-based methods that have demonstrated success in a wide range of inverse problems involving complex and noisy forward models. Ensemble methods are robust to noisy evaluations of the parameter-to-observable map (Dunbar et al., 2022a), naturally parallelizable, and scale well to high-dimensional parameter spaces, making them well suited for earthquake-cycle applications. Out of the box, the classical variants of EKI provide only pointwise estimates of the Bayesian solution (i.e. the posterior mean, or a maximum-likelihood solution) in nonlinear problems, and so we use a strategy known as Calibrate, Emulate, Sample, (Cleary et al., 2021; Dunbar et al., 2021b, 2024), which uses EKI samples to build a targeted machine-learning-based emulator that is then used for sampling the posterior. Specifically, we train a Gaussian process surrogate to approximate the mapping from model parameters to observable statistics, leveraging the fact that Gaussian processes perform well in small-data regimes (as each training pair arises from a computationally expensive seismic record simulation). This surrogate enables efficient Bayesian sampling of the posterior distribution using standard techniques such as Markov Chain Monte Carlo. The method is organized into two stages, described in the following subsections. Section 3.1 presents the calibration step based on

Ensemble Kalman Inversion, and Section 3.2 describes the emulate-sample stage used for posterior uncertainty quantification.

### 3.1 Calibrate: Ensemble Kalman Inversion

To solve the inverse problem defined in Eq. 1, we employ the continuous-time formulation of Ensemble Kalman Inversion (EKI) (Iglesias et al., 2013; Iglesias and Yang, 2021). Referring to the inverse problem defined in Eq. 1, let $y^\dagger$ denote the observed data. The calibration stage seeks parameters $u$ for which the model outputs $G(u)$ reproduce $y^\dagger$ with a small misfit.

Let $\{u^{(j)}(t)\}_{j=1}^J$ denote an ensemble of parameters at final algorithmic time $t$. Define the ensemble means

$$\bar{u}(t) = \frac{1}{J}\sum_{j=1}^{J} u^{(j)}(t), \qquad \bar{G}(t) = \frac{1}{J}\sum_{j=1}^{J} G(u^{(j)}(t)),$$

and the empirical cross-covariance between parameters and model outputs

$$C_{uG}(t) = \frac{1}{J}\sum_{j=1}^{J}\big(u^{(j)}(t) - \bar{u}(t)\big)\big(G(u^{(j)}(t)) - \bar{G}(t)\big)^\top. \tag{17}$$

Following Iglesias and Yang (2021), each ensemble member evolves according to the continuous-time EKI dynamics

$$\frac{du^{(j)}}{dt} = -\,C_{uG}(t)\,\Gamma_y^{-1}\big(G(u^{(j)}(t)) - y^\dagger\big), \tag{18}$$

where $\Gamma_y$ is the observational noise covariance. At algorithmic time $t = 1$, each ensemble member represents a sample from the Gaussian approximation of the Bayesian posterior distribution ($\pi(u|y^\dagger)$) (Iglesias and Yang, 2021). This formulation is *derivative-free*: it does not require gradients or adjoints of the forward map, but instead relies only on forward evaluations at the ensemble members. Nevertheless, the evolution of the ensemble has a close connection to gradient-based optimization methods. In linear inverse problems this link can be made precise, while in nonlinear settings it holds approximately, with the ensemble dynamics behaving similarly to a projected gradient descent in parameter space (Schillings and Stuart, 2016; Garbuno-Inigo et al., 2020). In addition, the use of ensemble differences introduces an intrinsic averaging effect that improves robustness to noisy forward evaluations (Dunbar et al., 2022a). In practice, the method typically

moves the ensemble toward regions of small data misfit within relatively few iterations, making EKI well suited to problems with expensive forward models.

Empirically, Ensemble Kalman Inversion (EKI) is observed to robustly produce accurate point estimates in nonlinear, derivative-free inverse problems such as Eq. 1. In practice, the ensemble at the final iteration is often used as an approximation to the posterior distribution, and in linear–Gaussian settings this correspondence can be justified rigorously Calvello et al. (2025). Even in nonlinear problems, the resulting ensemble typically provides a highly informative characterization of parameter uncertainty, in particular for unimodal Bayesian inverse problems. Nevertheless, the ensemble spread does not capture the full posterior variability, and additional refinement can improve uncertainty quantification. In this work, we further refine the approximate posterior obtained from EKI using a Gaussian process–based machine learning surrogate to achieve a more accurate characterization of the model parameters; this procedure is described in the following subsection.

### 3.2 Emulate and Sample: Surrogate-Based Uncertainty Quantification

The EKI calibration stage yields a set of parameter–observation $\{(u_i, G(u_i))\}_i$ pairs concentrated near regions of high posterior probability. Using these pairs, it is possible to construct a surrogate of the forward map and perform posterior sampling with the surrogate in place of the expensive forward model, resulting in a computationally efficient two-stage approach to approximate Bayesian inference. This strategy of training an input–output emulator from EKI samples and subsequently sampling the posterior has been developed in previous work (Cleary et al., 2021), primarily within the climate dynamics community, and we adapt it here to the earthquake-cycle setting. It should be noted that the emulator that we train here is different from the emulators designed to simulate the full earthquake cycles models (Kaveh et al., 2025a; Magen et al., 2025). Here, the emulator only models the parameter ($u \in \mathbb{R}^{d_u}$) to observable statistics $y \in \mathbb{R}^{d_y}$; the underlying PDE dynamics is not directly emulated.

**Emulator construction.** Training an accurate surrogate for the map $u \mapsto G(u)$ is typically the main computational bottleneck in Bayesian inversion for two reasons: (i) the parameter space

is moderately high dimensional while the number of available training samples is limited, and (ii) each training datum consists of a pair $(u, G(u))$, where evaluating $G(u)$ requires running the forward model for a long time. In contrast to the reduced-order modelling strategy of Kaveh et al. (2025a), where a neural network emulator is trained to approximate the forward model in a low-dimensional latent space, the goal in the present work is to directly approximate the map from model parameters to observable statistics. This map is considerably smoother and lower-dimensional than the full dynamical trajectory, and therefore better suited to statistical emulation methods. For this reason, we employ Gaussian process (GP) regression rather than neural networks: GPs provide reliable uncertainty estimates and perform well in regimes in which data are sparse and additional samples are expensive to obtain.

A further difficulty is that the corresponding model outputs may be strongly correlated, which degrades the conditioning of the regression problem and increases the effective dimensionality of the surrogate. To address this, we apply linear dimension reduction to the output space. Let $\{u^{(j)}, G(u^{(j)})\}_{j=1}^{M}$ denote the EKI-adapted training set. We perform principal component analysis (PCA) on the outputs; in particular, we use the estimated observational covariance matrix $\Gamma_y$ to define the correlation of the output space. Writing

$$G(u^{(j)}) = \bar{y} + U_{\mathrm{out}} w^{(j)},$$

the matrix $U_{\mathrm{out}}$ contains the leading principal directions of the output ensembles, and the latent variables $w^{(j)}$ contain the associated PCA coordinates. We retain only the principal components that capture a prescribed percentage (here $95\%$) of the empirical variance. Applying this threshold to the training data shows that the $80$-dimensional output vector has an intrinsic dimension of $23$. These reductions reflect the presence of strong correlations in the spatially distributed observational data.

We use independent Gaussian processes for each latent output dimension:

$$w_k = f_k(u) + \varepsilon_k, \qquad k = 1, \ldots, r_{\mathrm{out}},$$

where $r_{\mathrm{out}}$ is the retained latent dimensions of the output spaces, and $\varepsilon_k$ represents observational

noise or modelling error. For each latent component $k$, we employ a Gaussian Process (GP) regression model with a squared-exponential (radial basis function) kernel augmented by a diagonal noise term. The kernel hyperparameters consist of a signal variance and one length-scale parameter per input dimension. These parameters are learned by maximizing the marginal likelihood of the training data. Further details of the GP formulation and training procedure are provided in Appendix B. Let $\widehat{f}_k$ denote the GP posterior mean for each latent direction. The resulting surrogate in latent space is

$$\widehat{w}_k(u) = \widehat{f}_k(u), \qquad k = 1, \ldots, r_{\mathrm{out}}.$$

Finally, lifting back to the physical output space is achieved via the PCA decoder,

$$\widehat{G}(u) = \bar{y} + U_{\mathrm{out}}\, \widehat{w}(u).$$

Thus, the emulator is trained entirely in low-dimensional latent coordinates and reconstructed in the full output space through linear maps. This procedure whitens the data, removes redundant degrees of freedom, improves numerical conditioning, and reduces the effective complexity of the surrogate-learning problem.

The resulting surrogate $\widehat{G}(u)$ provides a smooth and accurate approximation of the forward map in the region of interest,

$$G(u) \approx \widehat{G}(u),$$

and the smoothing inherent to GP regression mitigates small-scale numerical irregularities in $G(u)$. This regularity substantially improves the robustness and efficiency of the MCMC sampling carried out in the subsequent stage, using the GP in place of the exact parameter to observables map. Alongside the Gaussian process posterior mean providing a surrogate model prediction at $u$, it also gives us an estimate of the uncertainty $\widehat{\Gamma}(u)$, this uncertainty combines both the observational noise, and an estimate of emulators own uncertainty (related to the lengthscale-weighted distance of $u$ to the emulator's training data set).

**Posterior sampling with an emulator:** Given the surrogate, we define, as in previous studies

(Cleary et al., 2021; Dunbar et al., 2021b), an approximate negative log-likelihood

$$\Phi_{\text{emu}}(u) = \tfrac{1}{2}\|y^\dagger - \widehat{G}(u)\|^2_{\widehat{\Gamma}(u)} + \log\det\widehat{\Gamma}(u),$$

and use this expression for posterior exploration. The log-determinant term is typically part of the Bayes normalization constant, but as $\widehat{\Gamma}(u)$ now depends on $u$, we explicitly write it out as part of the likelihood. All sampling steps rely solely on the surrogate, so no further evaluations of the forward model are required. We use the preconditioned Crank–Nicolson (pCN) algorithm to perform MCMC, which is dimension-robust under Gaussian priors (Cotter et al., 2013). However, the framework is modular, and other MCMC or SMC algorithms could be used in place of pCN.

This emulate–sample strategy offers several advantages. First, posterior sampling is extremely fast, as it uses only inexpensive surrogate evaluations. Second, the smoothing effect of the statistical emulator leads to a well-behaved objective landscape, enabling efficient sampling without the need for advanced proposals or annealing strategies. Third, because the emulator is trained on EKI-adapted samples and on dimension-reduced data, it achieves high accuracy using only $\mathcal{O}(10^3)$ forward-model evaluations-several orders of magnitude fewer than required in most classical or ML-accelerated Bayesian inversion frameworks. The resulting method provides scalable approximate uncertainty quantification in settings where direct Bayesian sampling would be computationally infeasible. As the posterior represents a very small volume relative to the prior, the Calibration stage is very important to tailor the training point selection to this region of space; in this way we do not waste forward model runs in sampling regions for training that are very unlikely to be consistent with observations.

## 4 INVERSION RESULTS

We now present the results of the inverse problem defined in Eq. 1. The calibration stage of solving the inverse problem is initialized with $90$ ensemble members drawn from the uniform prior distributions listed in Table 2. To incorporate prior knowledge about the seismogenic fault segments, we assume the sign of $(a - b)$ in each segment is known, and rather than inverting for $a\bar{\sigma}$ and $b\bar{\sigma}$

independently, we invert for $a\bar{\sigma}$ and the ratio $x = b/a$ in velocity-strengthening (VS) segments and $x = a/b$ in velocity-weakening (VW) segments, imposing $0 \leq x \leq 1$.

Fig 7 shows the evolution of the ensemble members during the EKI iterations. Grey curves denote individual ensemble trajectories, the black dashed curve shows the ensemble mean, and the red dashed line marks the true parameter value. The algorithm rapidly drives the ensemble toward regions of low misfit, and convergence is achieved in $\mathcal{O}(10^3)$ forward simulations. This demonstrates that the method is computationally tractable despite the chaotic dynamics of the forward model and the noisy, finite-length statistics used in the inversion.

The first row of Fig. 7 shows that the lengths of the two VW segments are well constrained, with the ensemble collapsing tightly around the true values. In contrast, parameters associated with the outer VS segments exhibit weaker contraction and noticeable bias, especially on the right-hand side of the fault. This behavior is physically consistent: each of these outer VS segments contains only one trench, and the number of detectable events decreases rapidly with distance into the VS regions. Although the properties of these VS regions influence neighboring VW segments, the resulting imprint on the paleoseismic statistics is too weak to robustly constrain $a\bar{\sigma}$ and $b/a$ in those segments. Nevertheless, the parameter $d_{rs}$ in the VS regions is moderately constrained, whereas $a\bar{\sigma}$ and the friction ratio remain biased, with the bias more pronounced in VS3 than VS1.

Within the two VW segments, the inversion performance is significantly stronger. The parameters $a\bar{\sigma}$, $d_{rs}$, and the friction ratio are all well recovered, with the ratio parameter exhibiting the smallest uncertainty among all inferred quantities. The middle VS segment (VS2) is also reasonably constrained, although the posterior mean shows slightly larger bias and variance than the VW segments. This is likely due to the strong dynamical role of VS2 in controlling rupture segmentation and hence its stronger connection to the paleoseismic statistics.

Following EKI, we train a Gaussian-process surrogate for the parameter-to-statistics map using the adapted ensemble samples and perform Bayesian sampling using Markov chain Monte Carlo. We run $200{,}000$ MCMC iterations, achieving an average acceptance rate of approximately $0.206$, which is well within the desirable range for high-dimensional posteriors. The chains exhibit stable behavior and visually good mixing.

Fig 8 shows the resulting marginal posterior distributions (black) together with the priors (blue). The true parameter values are marked by vertical dashed lines. For parameters in the VW segments and the central VS segment, the posteriors are substantially narrower than the priors and typically peak close to the truth, confirming strong data constraint. In contrast, parameters in the outer VS regions retain a much broader spread and bias, consistent with the qualitative EKI behavior discussed above.

Fig 9 presents the reconstructed spatial profiles of $a\bar{\sigma}$, the friction ratio, and $d_{rs}$ along strike. The posterior mean profiles closely follow the true distributions in the VW segments and in the central VS region, with credible intervals that reflect the expected level of uncertainty. Deviations occur primarily in VS1 and VS3, where observational information is sparse and finite-time statistics provide limited discriminatory power. Nonetheless, the inversion successfully recovers the large-scale structure of the heterogeneous frictional properties and the segmentation geometry.

These results demonstrate that paleoseismic-scale observations, despite sparsity, noise, and limited duration, can constrain key heterogeneous earthquake-cycle parameters within a Bayesian framework built on stationary statistics. Parameters that directly influence recurrence, variability, and slip amplitude are best constrained, while those associated with poorly observed VS regions remain less identifiable. The approach, therefore, provides a statistically grounded and computationally feasible framework route toward calibrating physics-based earthquake-cycle models from paleoseismic data.

## 5 DISCUSSION

This section discusses several aspects of the inversion framework and its results. Section 5.1 summarizes the main modeling, statistical, and computational assumptions and discusses their limitations. Section 5.2 examines the statistical consistency of the inversion results and the identifiability of the parameters. Section 5.3 interprets the inferred values of $(a-b)\bar{\sigma}$ in terms of physical scaling relations from rate-and-state friction theory. Section 5.4 explains why the parameter $d_{rs}$ is comparatively well constrained in the velocity-strengthening segments. Section 5.5 presents posterior

predictive validation using the full forward model. Finally, Section 5.6 evaluates the robustness of the surrogate model outside the calibration domain.

## 5.1 Limitations & Assumptions

Our formulation relies on several modeling, statistical, and computational assumptions, which we recall and discuss below.

*FORWARD-MODEL APPROXIMATION.*

We employ a quasi-dynamic earthquake-cycle model with static stress transfer (Eq. 4). This represents a first-order approximation to fully dynamic rupture processes. The choice is motivated by computational considerations: solving the inverse problem requires repeated forward simulations, and fully dynamic models are computationally prohibitive. Nevertheless, the proposed approach is computationally feasible: as demonstrated later, the EKI stage converges using approximately 2,500 forward simulations of the 5,500-year dynamical model. Because the framework is designed to require relatively few forward-model evaluations, it can readily be extended to more complex forward models, including fully dynamic PDE models with memory (Lapusta et al., 2000; Lapusta and Liu, 2009).

*Ergodicity and stationarity of statistics.*

A central assumption of this work is that the earthquake-cycle system is ergodic, so that long-term statistics of observable quantities are stationary and independent of initial conditions. This property is readily observed in long numerical simulations of the dynamical model (see Section 2.4), but it has not yet been rigorously established for natural tectonic settings. Whether real earthquake sequences are ergodic remains an open question (Brodsky and Farge, 2026), and we do not attempt to address it here. Moreover, many paleoseismic records are too sparse to directly test ergodicity. Our framework should therefore be interpreted as demonstrating what is achievable if ergodicity holds.

*Fault segmentation and frictional regime.*

To aid identifiability and reduce parameter dimensionality, we assume a priori that the fault is divided into five segments and that the sign of $(a-b)$ is known in each. For the outer segments (VS1 and VS3), this assumption is supported by the absence of large historical earthquakes, suggesting velocity-strengthening behavior. Inferring that the central segment is velocity-weakening is more challenging, but is consistent with geodetic constraints and estimates of interseismic coupling. While relaxing this assumption is possible, it would enlarge the parameter space and may introduce additional uncertainties in recovering the parameters of the central segment.

*Resolution and nucleation-scale constraint.*

For computational tractability, we restrict attention to parameter sets that do not lead to excessively small nucleation sizes, since these require very fine spatial discretization. The nucleation length scale (Rubin and Ampuero, 2005) is estimated along the strike using

$$h_{\mathrm{ra}} \;=\; \min_{VW\ \text{segments}} \frac{2\mu\, d_{rs}\, b\bar{\sigma}}{\pi\big((b-a)\bar{\sigma}\big)^2},$$

and the grid spacing is chosen such that $\Delta x \le h_{\mathrm{ra}}/\kappa$ with $\kappa = 15$ points per nucleation length. This yields the grid size

$$N \;=\; 2^{\left\lceil \log_2\left(\frac{\kappa L}{h_{\mathrm{ra}}}\right)\right\rceil},$$

where $L$ is the fault length. During Calibration using, simulations are only performed if $N \le 2^{13}$; otherwise, the model is skipped to accelerate computation. For reference, the true model used in this study requires $N = 2^{11}$. We emphasize that this resolution constraint is imposed only during the calibration stage, where full forward simulations are required. The emulator is trained on input–output pairs generated from parameter sets satisfying this constraint; however, once trained, the emulator can be evaluated for parameter values outside this region.

*Observation-noise covariance.*

We assume the noise $\eta$ in Eq. 1 is Gaussian and we have access to an estimate of the observational covariance matrix $\Gamma$ obtained from long-model simulations using the true parameters. The true

state $u^{\mathrm{true}}$ and the corresponding covariance $\Gamma_{\mathrm{true}}$ are, of course, unknown in realistic settings, making this an approximation. Similar assumptions are commonplace in settings such as weather forecasting, (Cleary et al., 2021) as system variability can be estimated from observable data even without knowledge of the true state. In our context, with only a single paleoseismic record available, obtaining such an estimate is considerably more challenging, and an important direction for future work is to develop strategies for approximating $\Gamma$ from limited data—for example by exploring the sensitivity of $\Gamma$ to model parameters or by combining observations from multiple sites. In what follows, we proceed with $\Gamma$ computed from the true model, while acknowledging this limitation.

*Perfect model setting.*

We create the synthetic data using the forward model defined in Eq. 4. As we later use this framework for inversion, we are not able to make quantitative statements about how suitable our choice of forward model is for representing existing paleoseismic data sets. This is because we neglect that there is any discrepancy between the optimally configured model, and the true paleoseismic record, that cannot be described by additive Gaussian noise. We adopt this perfect-model assumption to isolate methodological aspects of the inversion framework, allowing quantitative investigation of the choice of loss functions and observables, parameter identifiability, and algorithms for uncertainty quantification in a setting free from confounding factors present in real data. For application to specific paleoseismic data sets, however, model misspecification and structural error must also be addressed. In principle, such effects could be incorporated by augmenting the parameter vector $u$ with additional parameters describing model discrepancy or structural corrections. In practice, however, these parameters may be only weakly identifiable from the available observations, and careful experimental design and the inclusion of additional data sources would likely be required.

These assumptions define the scope of the present work: our aim is to assess the feasibility and performance of a statistically grounded inversion framework under controlled conditions, thereby providing a foundation for more realistic future applications.

## 5.2 Misfit, statistical consistency, and identifiability

The EKI calibration successfully drives the ensemble toward parameters whose forward simulations reproduce the observed statistics within uncertainty (Fig. 10). The ensemble-mean prediction lies almost entirely within the $2\sigma$ observational bands for mean recurrence time, COV, PCC, and average slip, confirming that the inversion finds parameter sets that are statistically consistent with the paleoseismic data. This demonstrates that the statistics used here contain sufficient information to constrain the dominant features of long-term earthquake-cycle behavior, even though no trajectory-level matching is performed.

At the same time, the results make clear that identifiability is not uniform. Parameters in the central VW segments are well constrained because these regions repeatedly host large earthquakes that directly shape the observed statistics. In contrast, parameters in the outer VS segments influence the data mainly through their effect on rupture penetration and segmentation. Consequently, these parameters may remain biased even while the statistical misfit is well within the $2\sigma$ tolerance.

## 5.3 Discussion of $(a-b)\bar{\sigma}$

In classical rate-and-state theory, the recurrence time of earthquakes on a velocity-weakening patch is controlled to first order by $(b-a)\sigma$ and the rate at which the fault is reloaded. A scaling estimate gives $T \sim \Delta\tau/\dot{\tau}$ with $\Delta\tau \sim (b-a)\sigma$ and $\dot{\tau} \sim kV_{\mathrm{pl}}$. More explicit formulations of this scaling lead to recurrence-time estimates of the form (Barbot et al., 2012)

$$T_r \;=\; \frac{(b-a)\bar{\sigma}}{\mu}\,\frac{W}{v_{\mathrm{pl}}}\ln\!\left(\frac{v_{\mathrm{co}}}{v_{\mathrm{int}}}\right). \tag{19}$$

Here, $W$ denotes the characteristic width of the seismogenic velocity-weakening patch, $v_{\mathrm{int}}$ represents the typical interseismic slip velocity within that patch , and $v_{\mathrm{co}}$ denotes the coseismic slip rate. These expressions are consistent with nucleation-controlled models (Dieterich, 1992). Although our system is chaotic and the recurrence process is more complex, the inversion results are consistent with this physical picture: $(a-b)\bar{\sigma}$ is well constrained and convergenes quickly in the VW segments (Fig. 11).

This indicates that the interevent time statistics encode the effective $(a-b)\bar{\sigma}$ in the segments

that host repeated ruptures. Although $v_{\rm co}$ and $v_{\rm int}$ are not strictly constant in our model and depend on the evolving fault state, their influence enters primarily through the logarithmic factor in the recurrence-time expression. Over the parameter range explored in the inversion, this contribution varies more weakly than the leading multiplicative term. To first order, the dominant control on the recurrence scale is therefore the factor $(b-a)\bar{\sigma}$, while variations in $v_{\rm co}$, $v_{\rm int}$ are effectively absorbed into a secondary correction. This explains why $(a-b)\bar{\sigma}$ emerges as one of the most tightly constrained and identifiable parameter combinations in the velocity-weakening segments.

### 5.4 Discussion of $d_{rs}$ in velocity–strengthening regions

A notable feature of the inversion is that in the VS segments, $d_{rs}$ is better constrained than $a\bar{\sigma}$ and $b/a$. This follows directly from rate-and-state friction physics, and we briefly summarize the argument. Strength of the fault is given by:

$$\tau = \bar{\sigma}\Big[\mu^* + a\ln\Big(\frac{v}{v^*}\Big) + b\ln\Big(\frac{v^*\theta}{d_{rs}}\Big)\Big]. \tag{20}$$

Far inside a VS regions (segment 1 and 5) the fault creeps at a rate close to $v_{\rm pl}$, so $\theta^{ss} \approx d_{rs}/v_{\rm pl}$, giving

$$\tau_i = \bar{\sigma}\Big[\mu^* + a\ln\Big(\frac{v_{\rm pl}}{v^*}\Big) + b\ln\Big(\frac{v^*}{v_{\rm pl}}\Big)\Big]. \tag{21}$$

When rupture arrives, $v \to v_{\rm dyn}$ while $\theta$ initially remains $d_{rs}/v_{\rm pl}$, producing

$$\tau_{\rm peak} = \bar{\sigma}\Big[\mu^* + a\ln\Big(\frac{v_{\rm dyn}}{v^*}\Big) + b\ln\Big(\frac{v^*}{v_{\rm pl}}\Big)\Big], \tag{22}$$

and after state relaxation toward $\theta = d_{rs}/v_{\rm dyn}$,

$$\tau_d = \bar{\sigma}\Big[\mu^* + a\ln\Big(\frac{v_{\rm dyn}}{v^*}\Big) + b\ln\Big(\frac{v^*}{v_{\rm dyn}}\Big)\Big]. \tag{23}$$

Thus the stress drop is

$$\Delta\tau = \bar{\sigma} b\,\ln\Big(\frac{v_{\rm dyn}}{v_{\rm pl}}\Big). \tag{24}$$

For a single rupture, the weakening may be approximated as linear slip weakening with $d\tau/dS = b\bar{\sigma}/d_{rs}$, giving

$$\Delta\tau = \frac{b\bar{\sigma}}{d_{rs}}\,S. \tag{25}$$

Equating the two expressions yields

$$d_{rs} = \frac{S}{\ln(v_{\mathrm{dyn}}/v_{\mathrm{pl}})}, \tag{26}$$

where $S$ denotes the slip accumulated over the weakening phase of rupture. This expression shows that, in velocity–strengthening segments where the slip rate remains close to $v_{\mathrm{pl}}$, the parameter $d_{rs}$ is primarily controlled by the local weakening slip $S$ and by the ratio between dynamic and interseismic slip rates. We emphasize that $v_{\mathrm{dyn}}$ is not an independent quantity: it is itself determined by other frictional parameters, fault geometry, and rupture dynamics. Nevertheless, this relation provides the closest practical decoupling of $d_{rs}$ from the remaining parameters, and it explains why $d_{rs}$ is more identifiable than $a\bar{\sigma}$ or the friction ratio in the VS regions.

### 5.5 High-Fidelity Posterior Predictive Validation

The posterior distribution is obtained using the approximate (machine-learned) forward operator $\widehat{G}(u)$ together with pCN sampling. Because the true forward model $G(u)$ is computationally expensive, it is not used directly within the Bayesian inference procedure. Instead, it is employed only during the calibration stage to generate training samples, which are subsequently used to construct the emulator $\widehat{G}(u)$. Consequently, we do not have access to the exact posterior associated with the true model, and a direct comparison between the approximate and true posteriors is not possible.

Instead, we assess the physical consistency of the inferred parameters through a *posterior predictive* check performed in the true model. We draw samples $u^{(j)} \sim \pi(u \mid y^{\dagger})$ from the approximate posterior and propagate them through the high-fidelity forward model,

$$G\big(u^{(j)}\big), \qquad j = 1, \ldots, N_s.$$

This pushforward ensemble provides an estimate of the predictive distribution of paleo-observables under posterior parameter uncertainty. The resulting distribution is then compared against the observations $y^{\dagger}$.

For computational tractability, full-resolution simulations were performed only for posterior samples satisfying the numerical resolution constraint $N \leq 2^{13}$. From 150 posterior samples

drawn at random, 93 satisfied this condition and were used to construct the posterior predictive ensemble. The median and percentile ranges of the pushforward distribution are shown in Fig. 12. The median of the pushforward ensemble closely follows the observed paleo-statistics, while the spread of the ensemble quantifies the uncertainty induced by the posterior distribution.

### 5.6 Surrogate Robustness Outside the Calibration Domain

During the calibration stage, forward simulations were restricted to parameter values requiring spatial resolution $N \leq 2^{13}$. Consequently, the emulator $\widehat{G}(u)$ was trained only on this subset of parameter space.

The posterior distribution $\pi(u \mid y^\dagger)$ obtained using the surrogate model is not explicitly constrained to this regime and assigns nonzero probability to parameter values requiring higher resolution. Among 150 posterior samples drawn at random, 57 required $N > 2^{13}$, and 16 required $N = 2^{14}$. The empirical frequency $57/150$ therefore provides an estimate of the posterior probability mass associated with this higher-resolution region.

To assess the surrogate's extrapolation behavior, we evaluate the full forward model $G(u)$ for these 16 samples. This analysis is conditional on the subset of posterior draws lying outside the calibration regime and does not constitute an independent posterior predictive sample from the full posterior.

The resulting pushforward distribution is shown in Fig. 13. Compared to Fig. 12, the median prediction follows the observations $y^\dagger$ with reduced accuracy. Nevertheless, for most components of $y^\dagger$, the observed values remain within the 16th–84th percentile range of the conditional predictive distribution. These results indicate that the surrogate model retains partial predictive capability beyond the calibration domain.

## 6 CONCLUSIONS

We have developed and tested a Bayesian framework for constraining heterogeneous frictional parameters and fault-segmentation geometry in earthquake-cycle models using paleoseismic-style observations. A key contribution of this work is the rigorous formulation of parameter estima-

tion as an inverse problem based on observational statistics that mimic realistic paleoseismic data (noisy event times, noisy slip, and detection thresholds) in a controlled synthetic environment. Rather than attempting to reproduce individual event trajectories or to infer the full system state, we focus on spatially distributed statistics of interevent times, variability, temporal clustering, and mean coseismic slip, interpreted through an ergodic perspective. In chaotic earthquake-cycle models with observational noise, joint inversion for both model parameters and the initial condition over paleoseismic timescales is effectively ill-posed, as small errors in the initial state grow exponentially and render trajectory-level matching unstable and non-unique. This shift in perspective-from reproducing specific historical earthquake sequences to matching the stationary statistical behavior of the fault-targets the quantities most relevant for long-term forecasting and hazard assessment. It therefore becomes possible to invert for model parameters without specifying or estimating the initial condition, while still accounting for both observational noise and finite-time sampling effects.

We showed that Ensemble Kalman Inversion can efficiently identify parameter sets whose simulated statistics match the synthetic paleoseismic observations within estimated uncertainty, despite chaotic dynamics and noisy, finite-length records. The subsequent emulate–sample stage uses a machine-learning surrogate, specifically a PCA-reduced Gaussian process emulator, to refine the posterior samples obtained from the EKI algorithm and to improve the reliability of the resulting uncertainty quantification. The inversion recovers the frictional parameters in the velocity-weakening and central velocity-strengthening segments with relatively tight posteriors, while highlighting limited identifiability in poorly observed outer velocity-strengthening regions. The analysis of $(a-b)\bar{\sigma}$ and $d_{rs}$ further clarifies which combinations of rate-and-state parameters are most strongly encoded in paleoseismic-scale statistics and why some are better constrained than others.

Taken together, these results demonstrate that long-term paleoseismic observations, when interpreted through a statistically grounded, ergodic framework, can provide meaningful constraints on heterogeneous earthquake-cycle parameters and segmentation geometry, even in the presence of chaos and realistic observational noise. The study is conducted in a controlled synthetic set-

ting, with the same model class used for data generation and inversion, so that methodological performance can be clearly assessed. Future work should relax these idealizations by incorporating model error, exploring alternative forward models and additional data types (e.g., geodetic time series), and developing practical strategies to estimate observational covariances from limited real-world records. Nonetheless, the present results establish a feasible, physics-based calibration pathway from paleoseismic-scale data to parameter-consistent earthquake-cycle models, providing a foundation on which more realistic forecasting and hazard-assessment studies can be built.

## CONFLICT OF INTEREST DECLARATION

The authors declare there are no conflicts of interest for this manuscript.

## Acknowledgements

AMS is supported by a Department of Defense (DoD) Vannevar Bush Faculty Fellowship (award N00014-22-1-2790), which also partially supports HK. JPA and HK acknowledge partial support from the Center for Geomechanics and Mitigation of Geohazards. ORAD was supported by Schmidt Sciences, LLC, the U.S. National Science Foundation (Grant No. 2530746), and the Office of Naval Research (Grant No. N00014-23-1-2654).

**Table 1.** Known parameters

| Category | Symbol | Description | Value / Unit |
|---|---|---|---|
| **(a) Physical Parameters** | | | |
| | $\mu$ | Shear modulus | $30\ \mathrm{GPa}$ |
| | $c_s$ | Shear-wave speed | $3.3\ \mathrm{km\ s^{-1}}$ |
| | $v_{\mathrm{pl}}$ | Plate loading rate | $50\ \mathrm{mm\ yr^{-1}}$ |
| **(b) Geometrical Parameters** | | | |
| | $L_{\mathrm{VS}}^{(1)}$ | Length of the first VS segment (see Fig. 3) | $40\ \mathrm{km}$ |
| | $L_{\mathrm{VS}}^{(3)}$ | Length of the third VS segment (see Fig. 3) | $40\ \mathrm{km}$ |
| | $L$ | Total fault length subjected to rate-and-state friction (see Fig. 3) | $240\ \mathrm{km}$ |
| **(c) Observational Parameters** | | | |
| | $\sigma_{\mathrm{dating}}$ | Standard deviation of dating noise | $10\ \mathrm{yr}$ |
| | $\sigma_{\mathrm{slip}}$ | Standard deviation of slip measurement noise | $0.4\ \mathrm{m}$ |
| | $s_{\mathrm{thresh}} = s'_{thresh}$ | Coseismic slip threshold for an event to be detectable | $0.8\ \mathrm{m}$ |
| | $T$ | Total duration of the synthetic paleoseismic record | $5000\ \mathrm{yr}$ |
| | $T_{\mathrm{filter}}$ | Time window removed before event detection | $500\ \mathrm{yr}$ |

**Table 2.** Unknown parameters. Range of the priors and the True values

| **Segment** | **Parameter** | **Lower Bound** | **Upper Bound** | **True** |
|---|---|---|---|---|
| VW1 | $L_{VW1}^{(1)}$(km) | 65 | 77 | 72.5 |
| VW2 | $L_{VW1}^{(2)}$(km) | 65 | 77 | 72.5 |
| VS1 | $a\bar{\sigma}$(MPa) | 0.2 | 0.8 | 0.5 |
| VS1 | $b\bar{\sigma}$(MPa) | 0 | $a\bar{\sigma}$ | 0.05 |
| VS1 | $d_{\mathrm{rs}}$(mm) | 2 | 20 | 8 |
| VW1 | $a\bar{\sigma}$ (MPa) | 0.2 | 0.8 | 0.525 |
| VW1 | $b\bar{\sigma}$(MPa) | $a\bar{\sigma}$ | 1 | 0.7 |
| VW1 | $d_{\mathrm{rs}}$(mm) | 2 | 20 | 12 |
| VS2 | $a\bar{\sigma}$ (MPa) | 0.2 | 0.8 | 0.6 |
| VS2 | $b\bar{\sigma}$ (MPa) | 0 | $a\bar{\sigma}$ | 0.4 |
| VS2 | $d_{\mathrm{rs}}$(mm) | 2 | 20 | 8.8 |
| VW2 | $a\bar{\sigma}$ (MPa) | 0.2 | 0.8 | 0.5 |
| VW2 | $b\bar{\sigma}$ (MPa) | $a\bar{\sigma}$ | 1 | 0.75 |
| VW2 | $d_{\mathrm{rs}}$(mm) | 2 | 20 | 14.4 |
| VS3 | $a\bar{\sigma}$ (MPa) | 0.2 | 0.8 | 0.5 |
| VS3 | $b\bar{\sigma}$ (MPa) | 0 | $a\bar{\sigma}$ | 0.05 |
| VS3 | $d_{\mathrm{rs}}$(mm) | 2 | 20 | 8 |

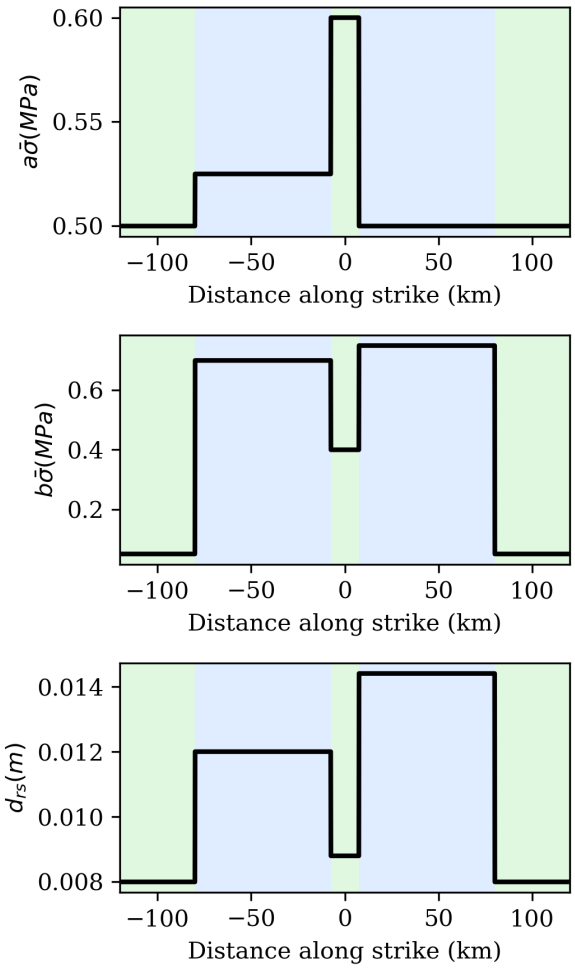


**Figure 1.** True parameter values used in this study ($u^{true}$), which are the targets of the inversion. Shaded bands denote alternating velocity-strengthening (green) and velocity-weakening (blue) fault segments.

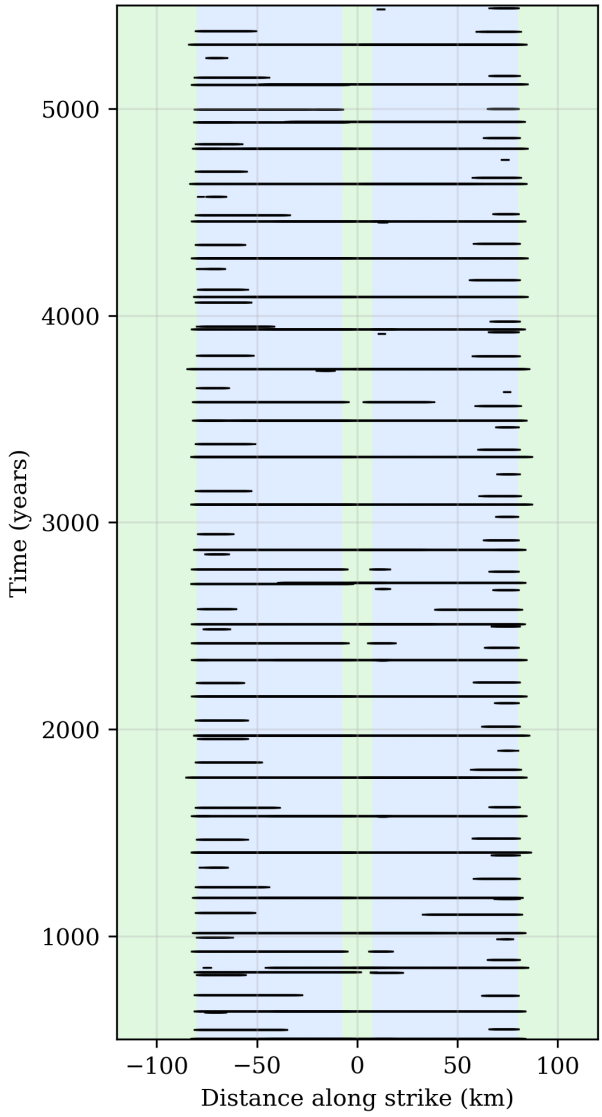


**Figure 2.** An example of the earthquake sequence generated by model Eq. 4 using the frictional parameter distributions shown in Fig. 1 and the parameters listed in Table 1. Horizontal line segments indicate the rupture extent of each earthquake. Shaded bands follow the convention of Fig. 1.

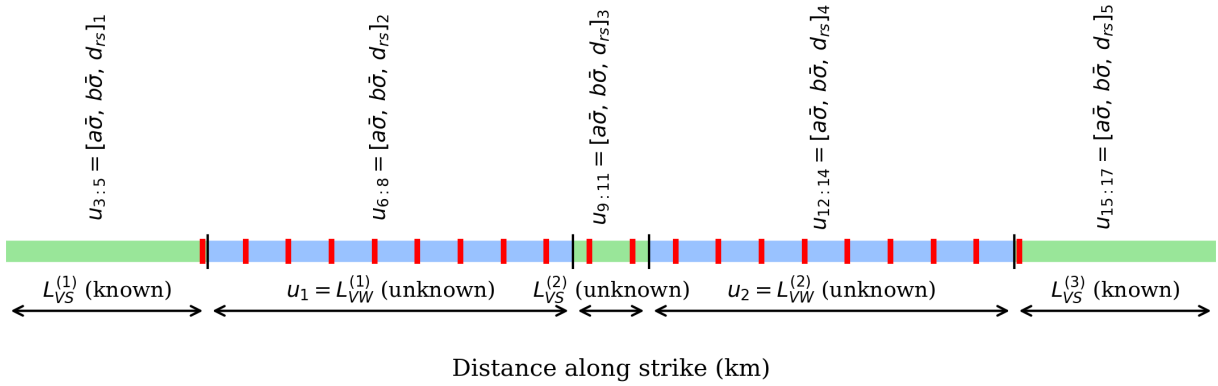


**Figure 3.** unknown parameters $(u)$ that we invert for in our framework. $u_1$ and $u_2$ are the lengths of VW segments, while the others $(u_{3:17})$ are the physical model parameters in different segments that we invert for. Red vertical lines are the locations in which we assume we have paleoseismic data for $(Z_t^o)$ (same as $Z_s^o$). Black vertical lines mark the boundaries between adjacent fault segments. Shaded bands follow the convention of Fig. 1.

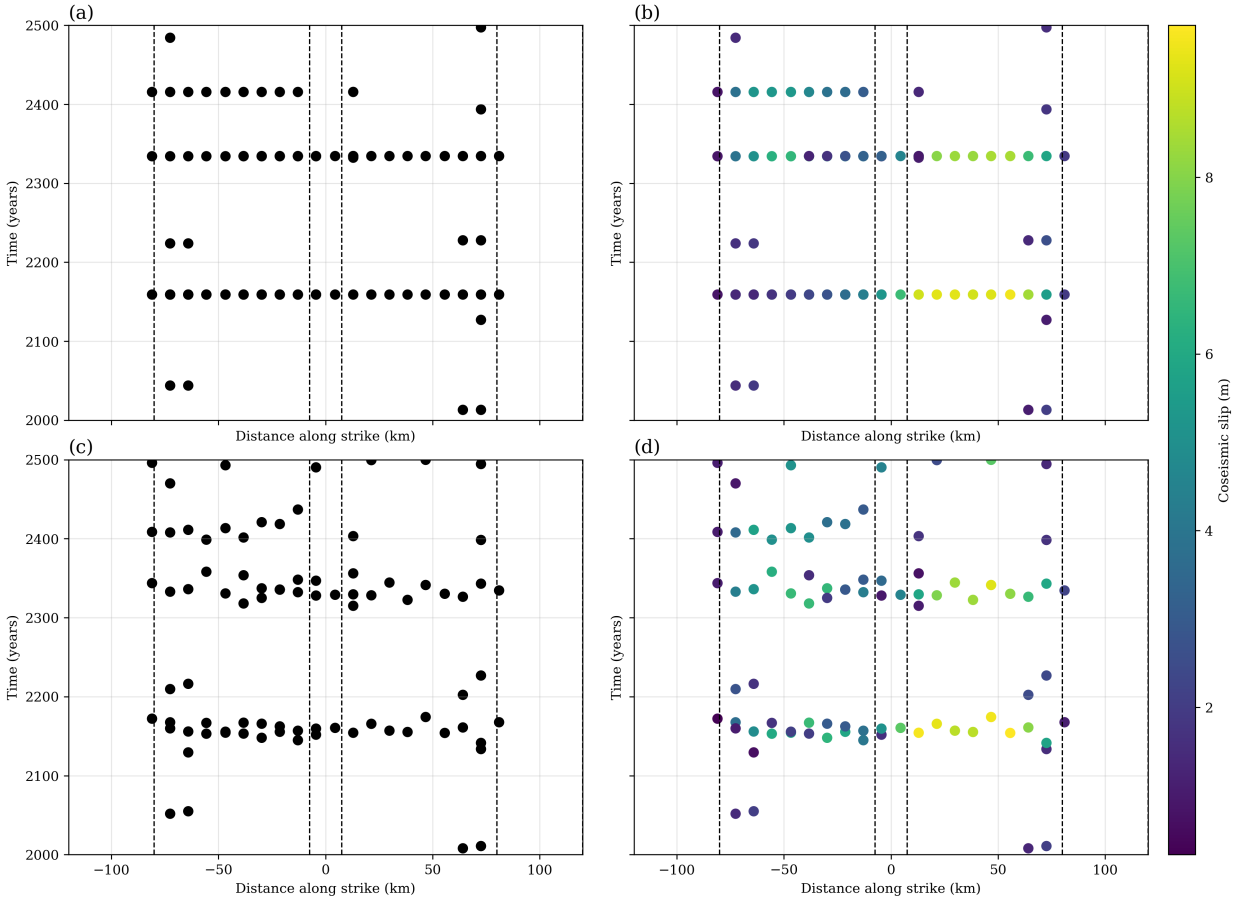


**Figure 4.** Synthetic paleoseismic observations extracted from the simulation at locations $Z_{\mathrm{t}}^{\mathrm{o}}$, plotted versus fault position (km) and time (years). Events are observed only if the coseismic slip exceeds $s'_{\mathrm{thresh}}$ and fall within the plotting window. (a) Noiseless *event times only* (black markers); slip values are not displayed. (b) Same noiseless events as in (a), colored by coseismic slip magnitude (m). (c) *Event times only* with dating noise added (black markers); slip values are not displayed. (d) Same noisy events as in (c), colored by coseismic slip magnitude after adding both dating noise (to times) and slip-measurement noise (to slips). Vertical dashed lines indicate boundaries between velocity-strengthening and velocity-weakening segments.

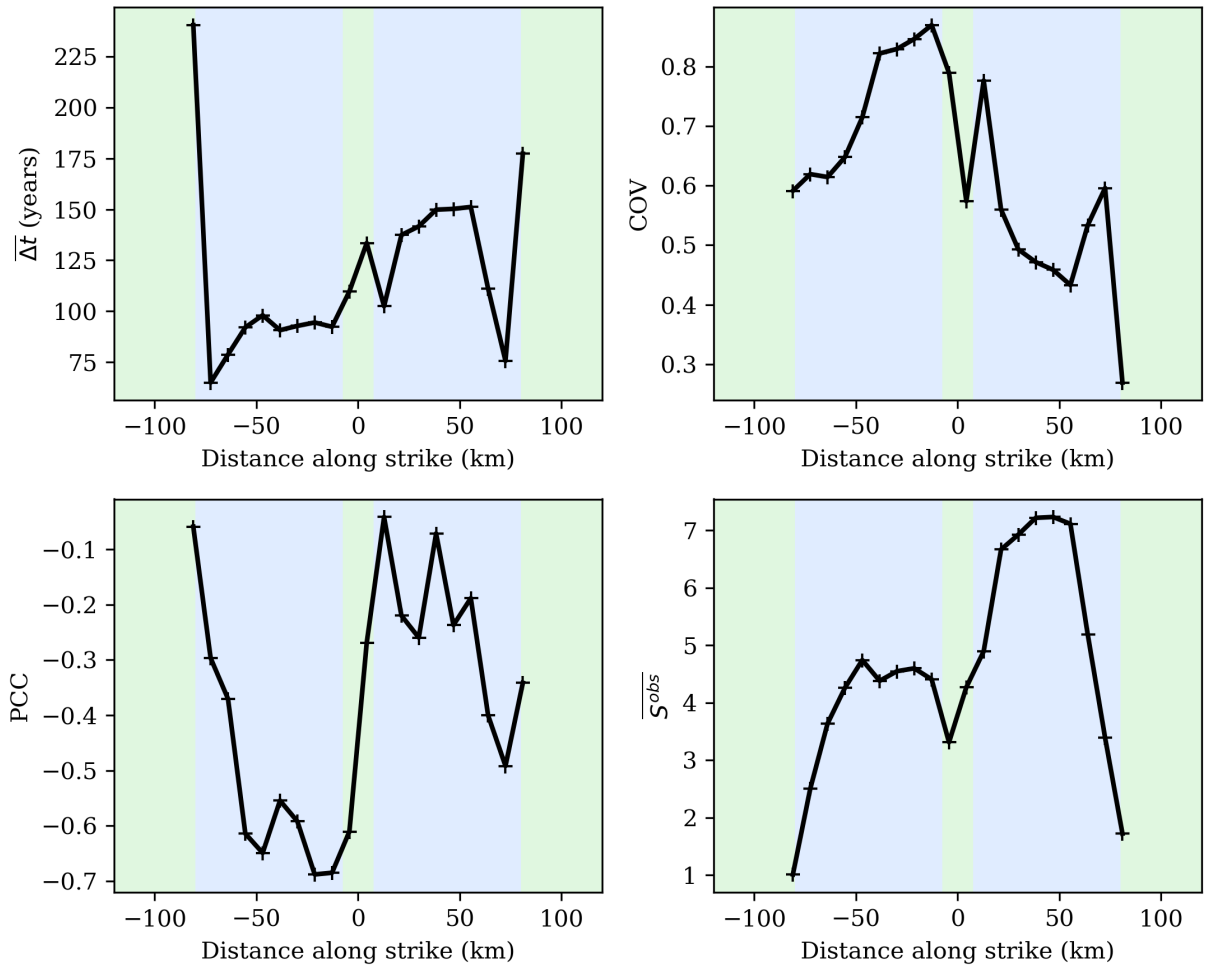


**Figure 5.** Noisy paleo-style summary of earthquake-cycle statistics ($y^\dagger$) along the fault for the true model used to estimate $u$. Each panel shows the spatial variation of a diagnostic quantity at the trench observation locations (indicated by "+" symbols), connected by black lines. Panels show: (a) mean interevent time, (b) coefficient of variation (COV) of interevent times, (c) Pearson correlation coefficient of successive recurrence intervals, and (d) mean coseismic slip. Shaded bands follow the convention of Fig. 1.

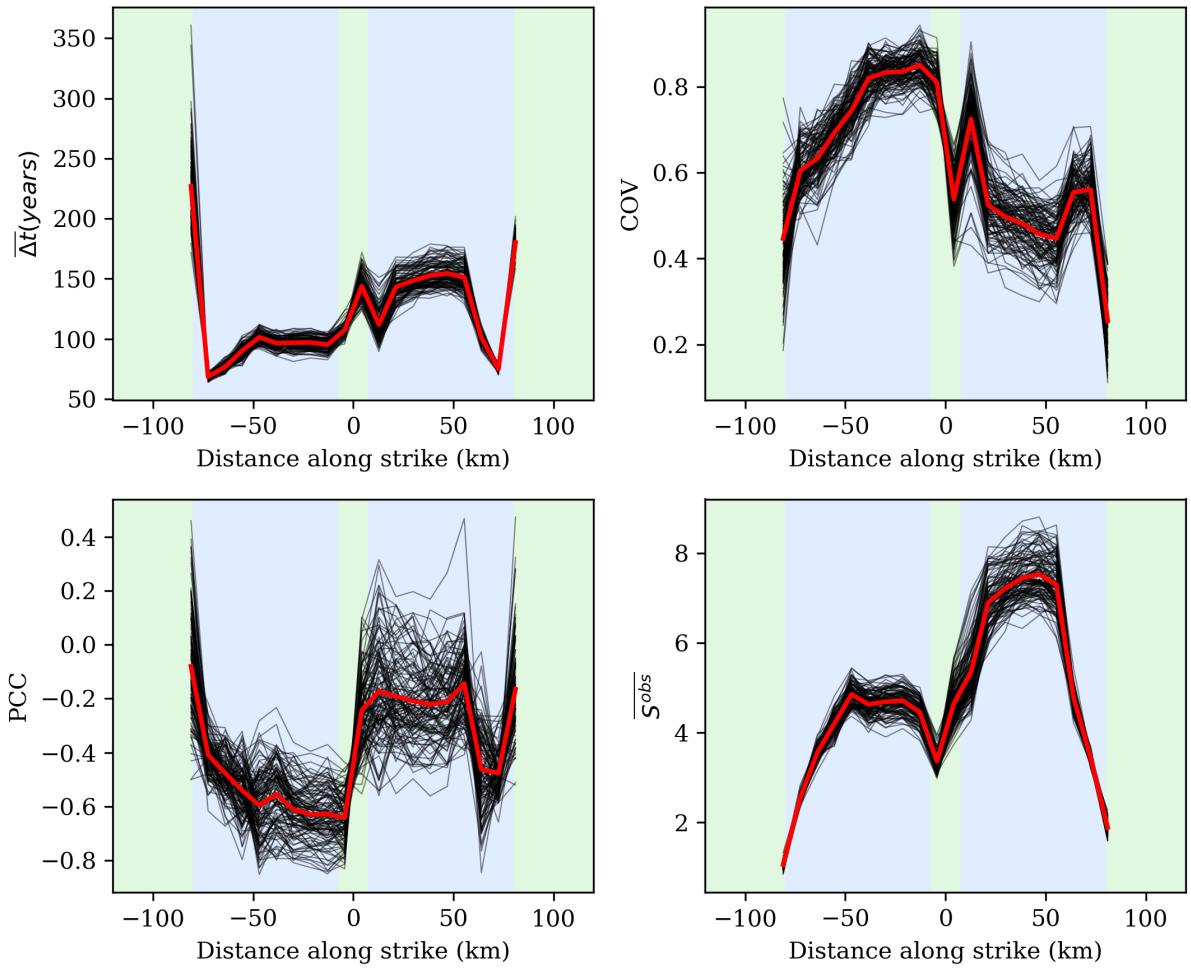


**Figure 6.** Realizations of $G(u^{true};T)$ and estimate of $\mathcal{G}_\infty(u^{true})$ plotted at the observational points. Average interevent time, coefficient of variation of the events, Pearson sample correlation coefficient, and the average slip for the observational points. The black lines are the random realizations of $G(u^{true};T)$, and the red line is the empirical average of those realizations. Shaded bands follow the convention of Fig. 1.

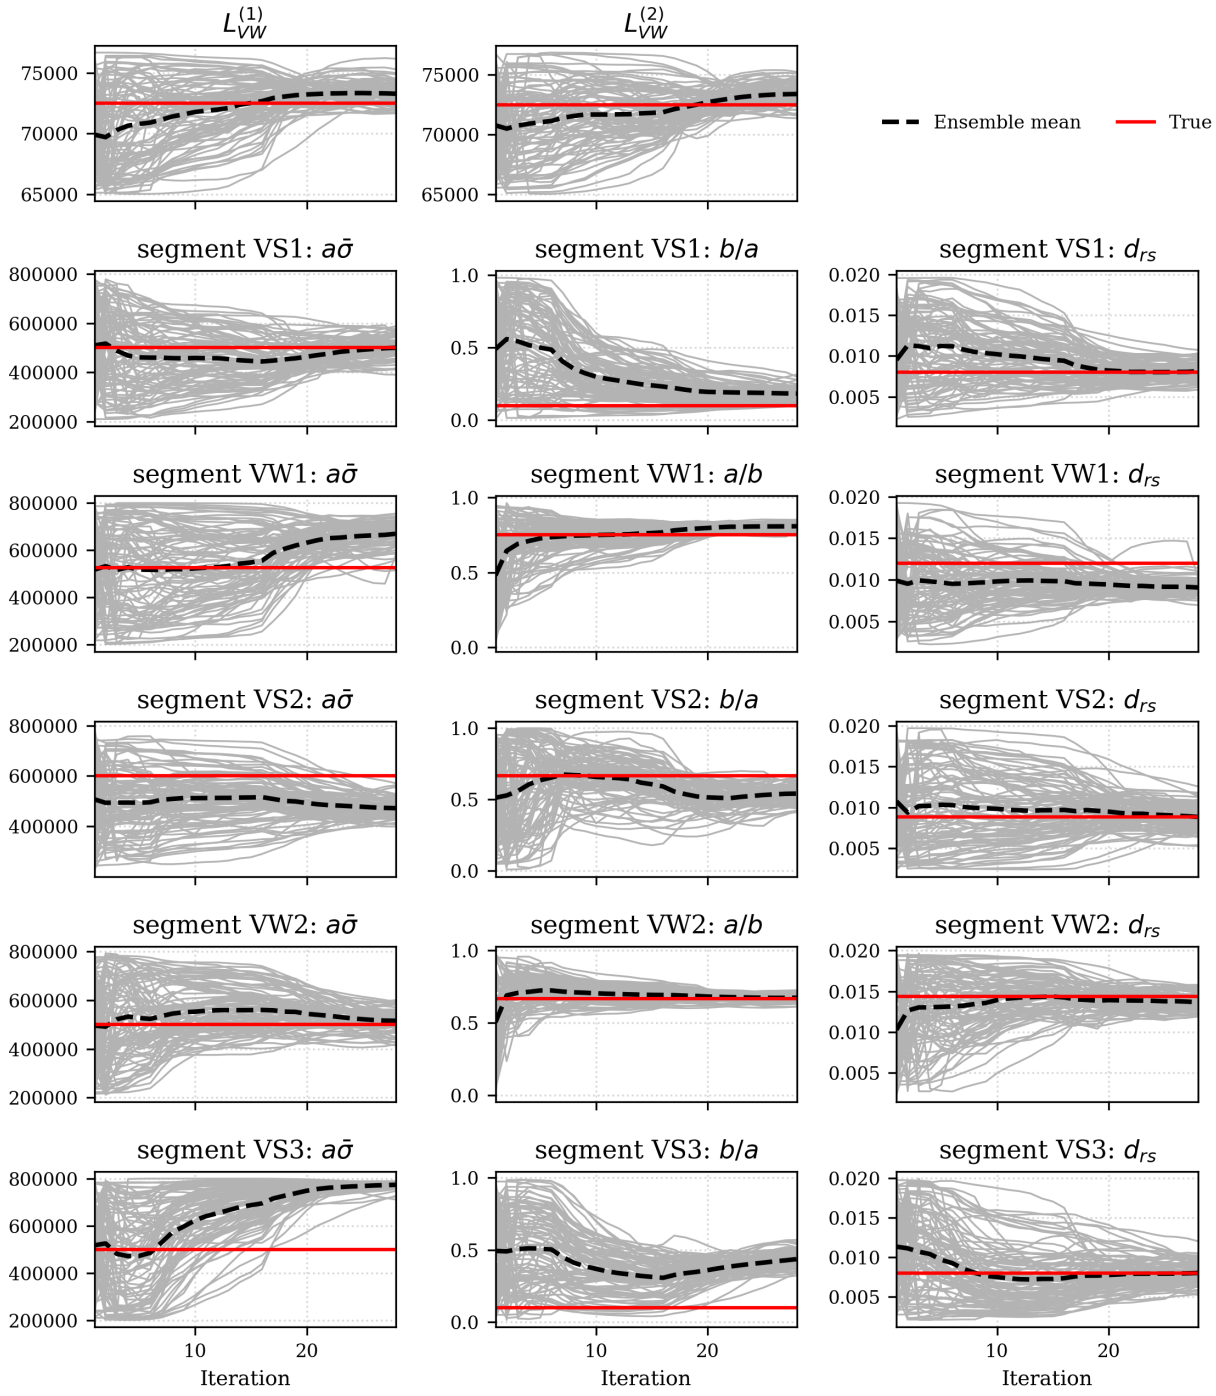


**Figure 7.** EKI sample paths. Each panel corresponds to one model parameter. Grey curves show the evolution of individual ensemble members across iterations, the dashed black curve denotes the ensemble mean, and the red line marks the true parameter value. The first row shows the inversion of the lengths of the velocity-weakening (VW) segments. Subsequent rows correspond to parameters within each fault segment, ordered from left to right along strike.

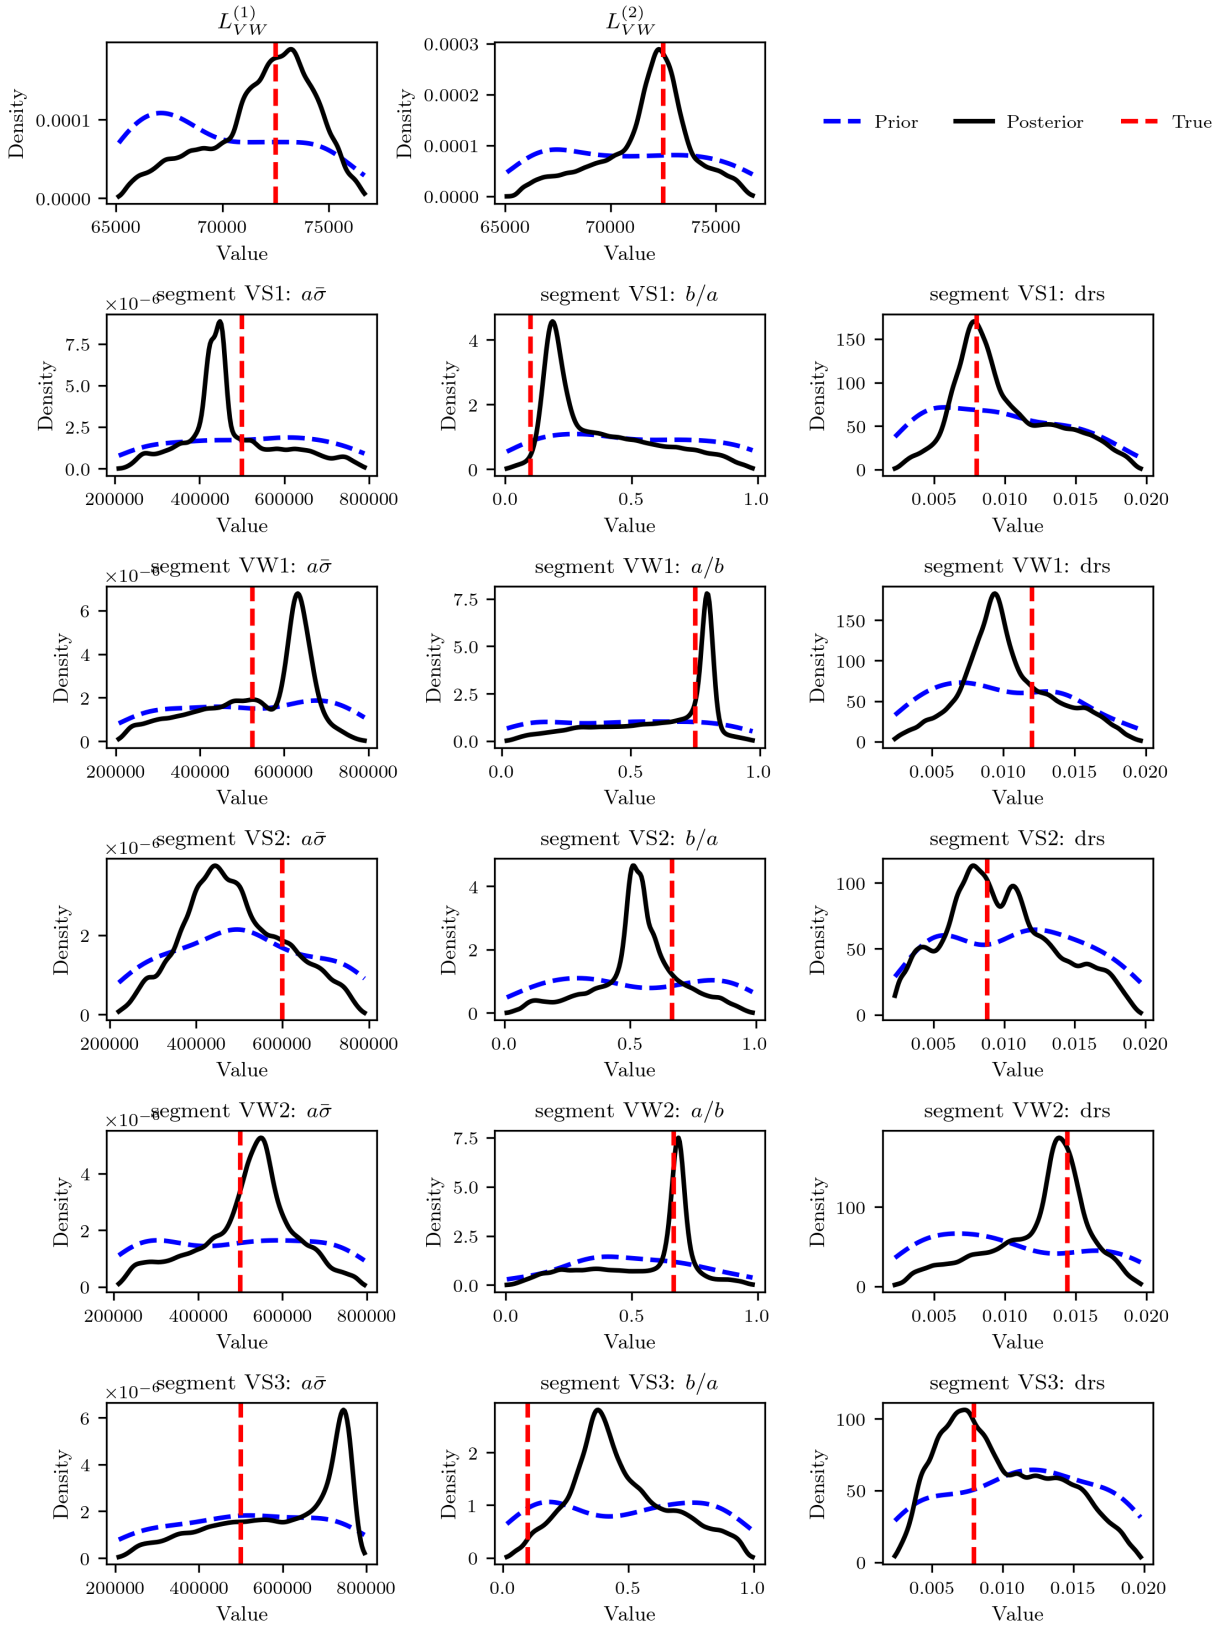


**Figure 8.** Posterior distributions of the model parameters (black) and the prior distributions (blue). The vertical dashed line indicates the true parameter value used to generate the synthetic data.

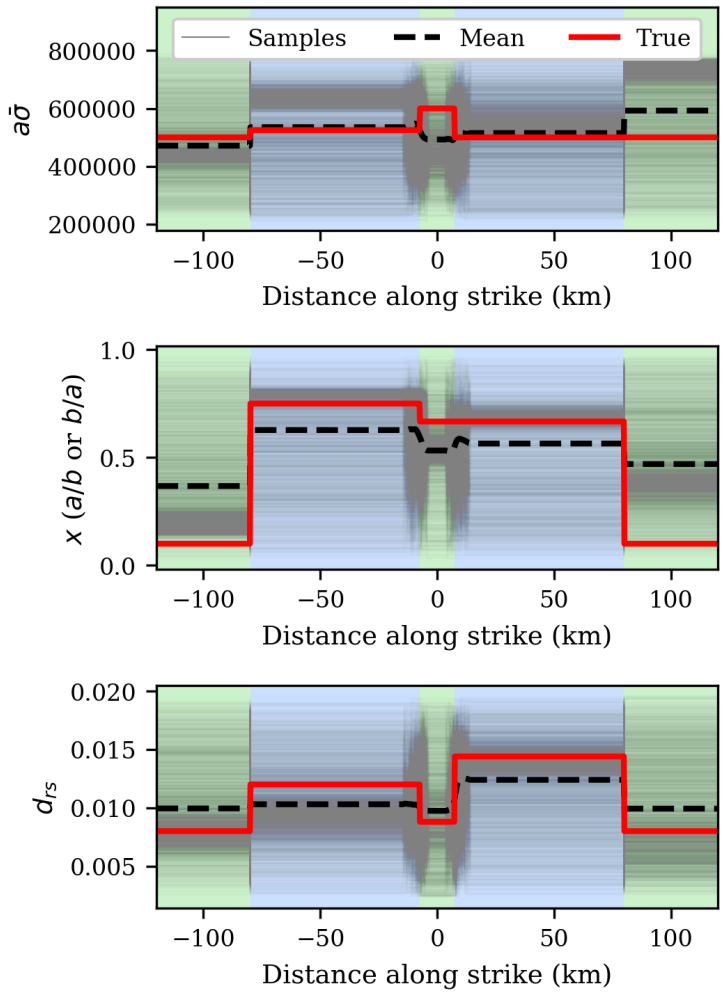


**Figure 9.** Posterior reconstructions of along-strike fault properties. Gray lines show $5{,}000$ posterior samples of the inferred piecewise-constant profiles for (top) $a\bar{\sigma}$, (middle) the friction ratio $x$ ($a/b$ or $b/a$, depending on segment type), and (bottom) $d_{rs}$. The dashed black curves denote the posterior mean profiles, and the red curves show the true profiles used to generate the synthetic data. Each posterior sample uses its own inferred lengths of VW segments ($L_{VW}^{(1)}$ and $L_{VW}^{(2)}$), while the outer boundaries are fixed. Shaded bands follow the convention of Fig. 1.

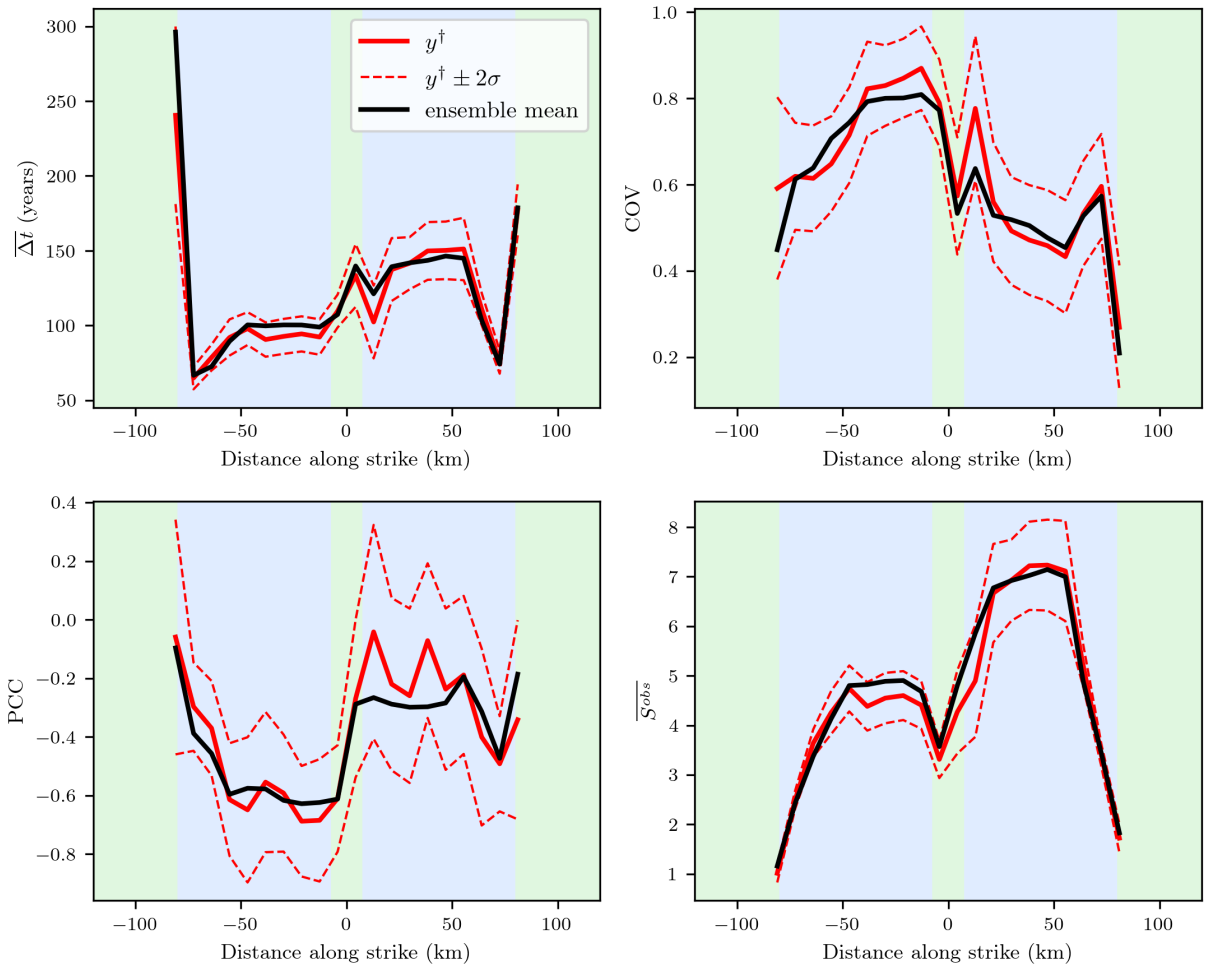


**Figure 10.** Comparison between the observed paleoseismic statistics ($y^\dagger$) and the ensemble-mean prediction from the final EKI iteration. Red curves show observations with $2\sigma$ uncertainty; Black curves denote the ensemble mean. Panels show mean interevent time, COV, PCC, and mean coseismic slip. Shaded bands follow the convention of Fig. 1.

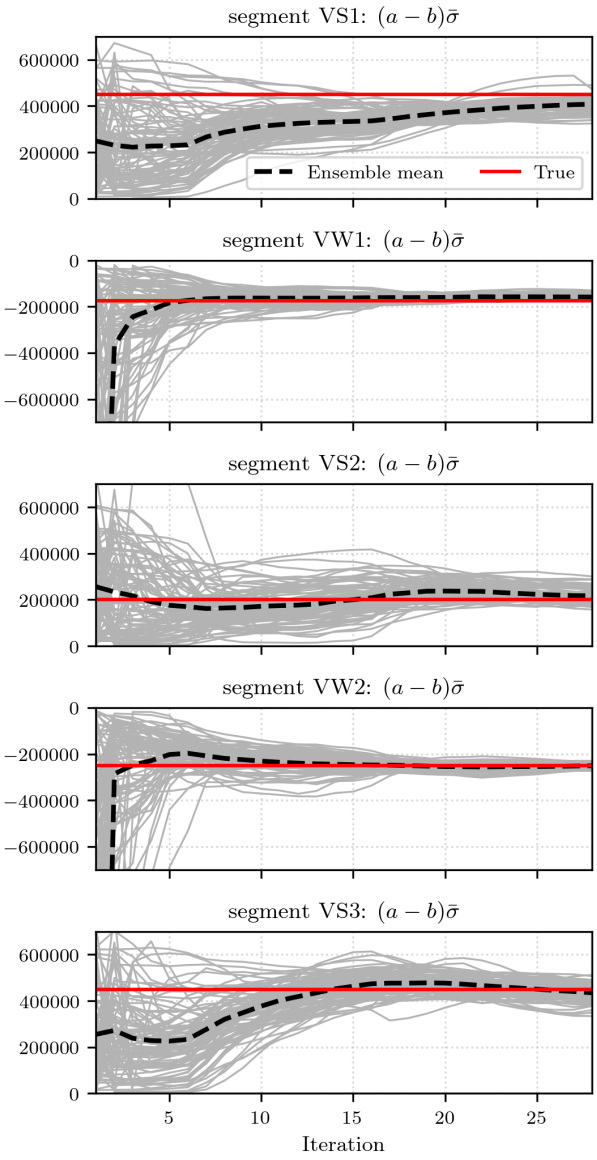


**Figure 11.** Evolution of $(a-b)\bar{\sigma}$ during Ensemble Kalman Inversion. Grey lines show ensemble members, the dashed black curve is the ensemble mean, and the red line indicates the true value for each segment.

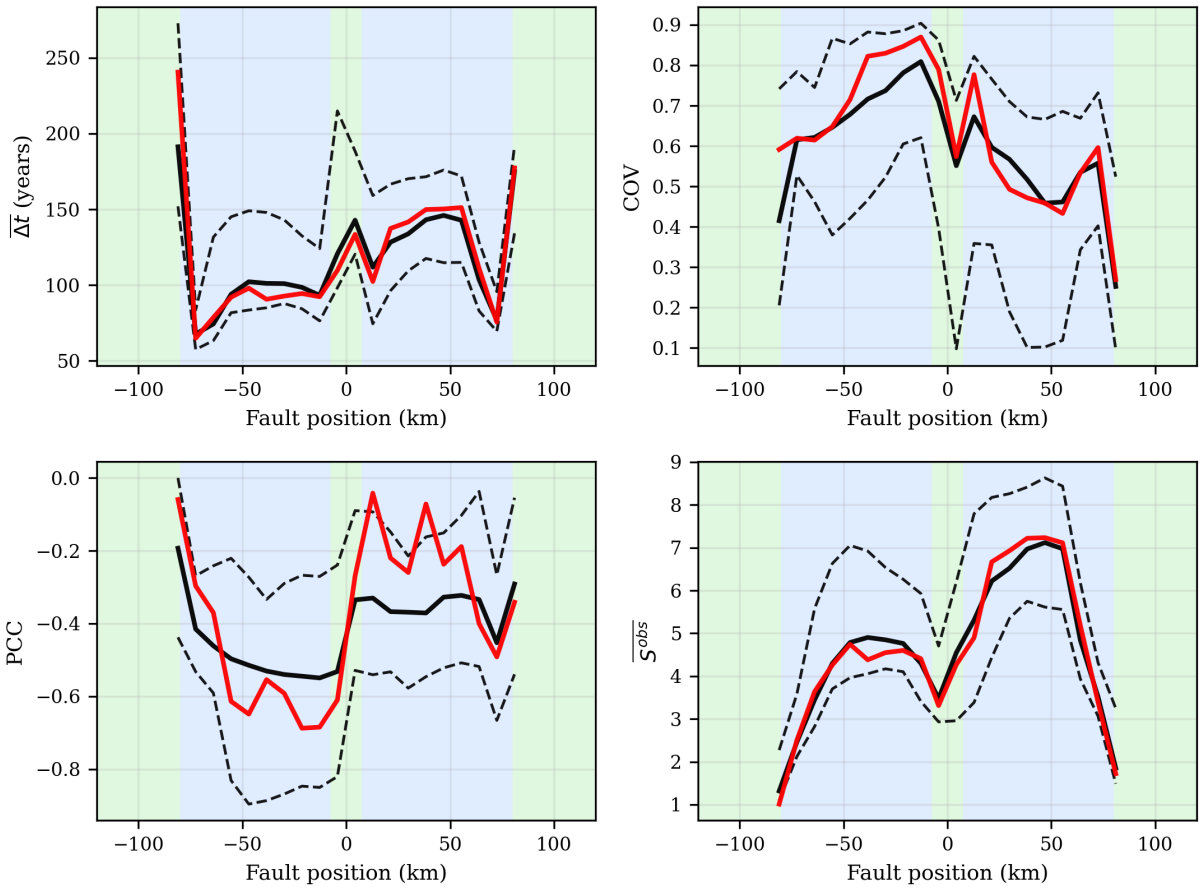


**Figure 12.** Spatial profiles of paleo-observables along the fault obtained via posterior predictive evaluation in the high-fidelity forward model. Parameter samples $u^{(j)} \sim \pi(u \mid y^{\dagger})$ that satisfy the numerical resolution constraint $N \leq 2^{13}$ are propagated through the true model to compute noisy summary statistics $G(u^{(j)}; T)$. Black solid curves show the median of the resulting pushforward ensemble, with dashed curves indicating the 16th–84th percentile range. Red curves denote the observed paleo-statistics $y^{\dagger}$. Panels display mean interevent time, COV, PCC, and mean coseismic slip. Shaded bands follow the convention of Fig. 1.

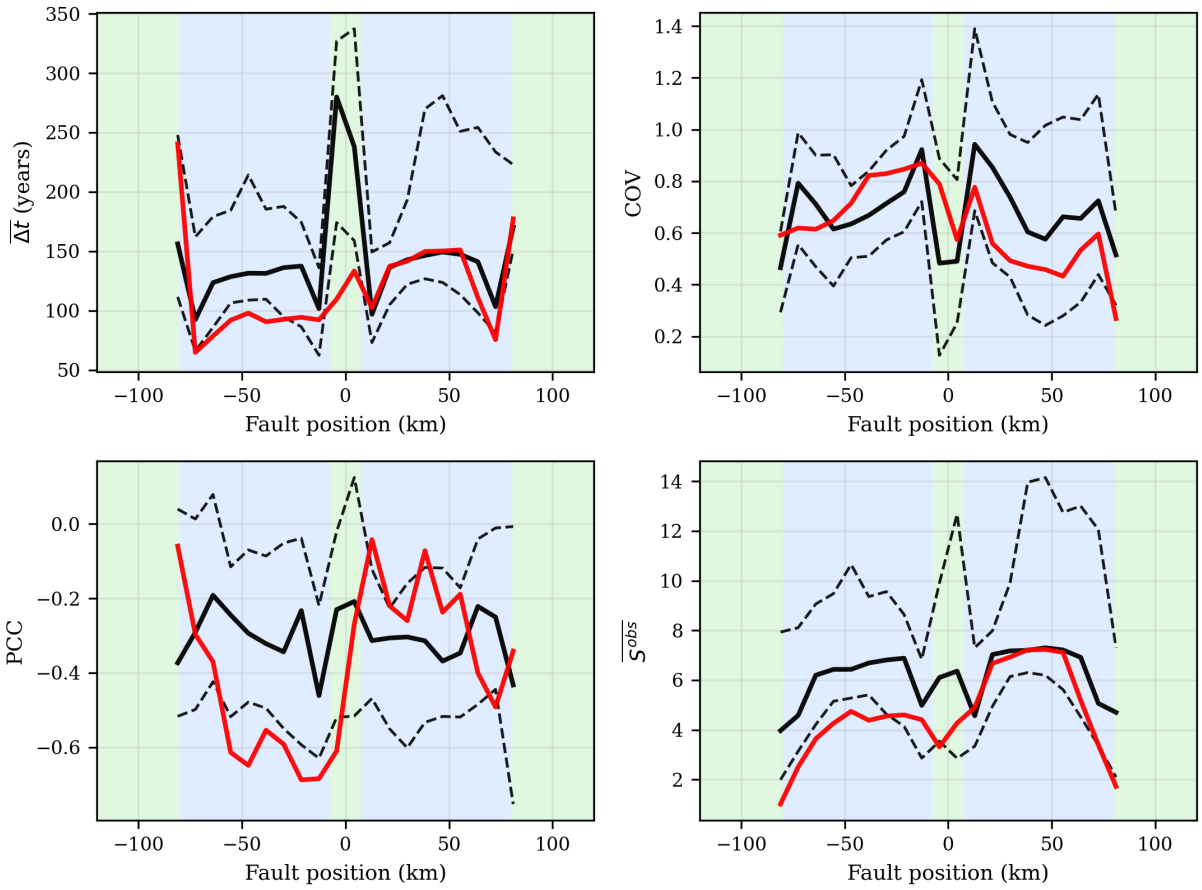


**Figure 13.** Spatial profiles of paleo-observables along the fault for the 16 posterior samples requiring resolution $N = 2^{14}$, i.e., outside the calibration regime ($N \leq 2^{13}$). These samples were drawn from $\pi(u \mid y^{\dagger})$ and represent an empirical posterior probability mass of $16/150$. The high-fidelity forward model $G(u)$ is evaluated for this subset to assess surrogate extrapolation. Black solid curves show the median of the resulting conditional pushforward distribution, with dashed curves indicating the 16th–84th percentile range. Red curves denote the observations. Panels show mean interevent time, COV, PCC, and mean coseismic slip. Shaded bands follow the convention of Fig. 1.

## APPENDIX A: FINITE-TIME AVERAGES IN ERGODIC SYSTEMS: AN EXAMPLE

**Example (Lorenz system).** Consider the Lorenz '63 model (Lorenz, 1963):

$$\dot{x}_1 = c_1(x_2 - x_1), \tag{A.1a}$$

$$\dot{x}_2 = x_1(c_2 - x_3) - x_2, \tag{A.1b}$$

$$\dot{x}_3 = x_1 x_2 - c_3 x_3. \tag{A.1c}$$

If parameters $[c_1, c_2, c_3] = [10, 28, 8/3]$ are chosen then the dynamics are chaotic and admit a natural invariant probability measure $\mu$ supported on the global attractor (Tucker, 1999). We denote by $x_1(\cdot) \in C(\mathbb{R}^+, \mathbb{R})$ the first coordinate of a trajectory initialized according to this invariant measure, i.e. $x(0) \sim \mu$.

Define the finite-time average

$$G(T) := \frac{1}{T}\int_0^T x_1(t)\, dt, \qquad \mathcal{G}_\infty := \mathbb{E}_\mu[G(T)]\,.$$

By symmetry of the Lorenz attractor in the $x_1$-direction, the true expectation satisfies $\mathcal{G}_\infty = 0$. For finite $T$, however, $G(T)$ is a random variable whose variability arises from sensitivity to initial conditions. Its variance decreases as the averaging window $T$ increases, and in the limit $T \to \infty$ the time average converges to the invariant expectation (Holland and Melbourne, 2007), zero in this case. This behavior is illustrated in Fig. A1.

## APPENDIX B: GAUSSIAN PROCESS EMULATOR

In this section we briefly describe the Gaussian Process (GP) regression framework used to construct the surrogate forward operator $\widehat{G}(u)$ from calibration data.

Suppose we are given a training data set

$$\{(u_i, w_i)\}_{i=1}^n, \qquad u_i \in \mathbb{R}^{d_u}, \quad w_i \in \mathbb{R},$$

where

$$w_i = f(u_i) + \epsilon_i, \qquad \epsilon_i \sim \mathcal{N}(0, \gamma^2)$$

represents observational noise, and $n$ denotes the number of calibration samples used in training

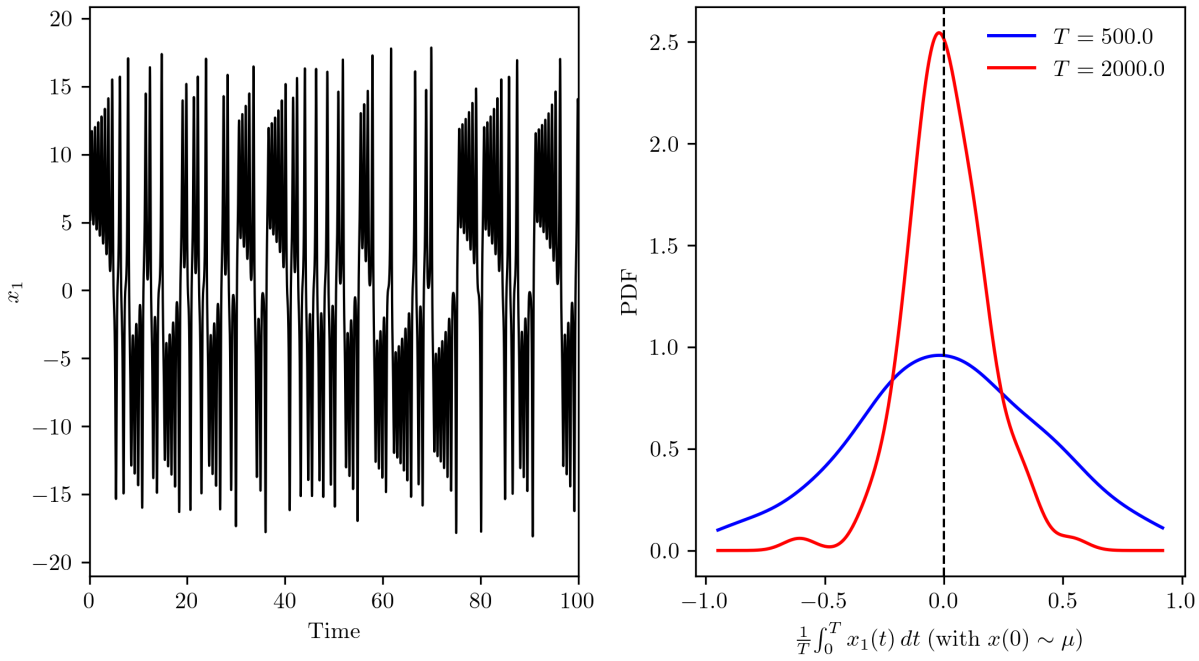


**Figure A1.** Finite-time averaging in the Lorenz system. Left: Chaotic time series of $x_1(t)$ in Eq. A.1. Right: Empirical distribution of the finite-time average $\frac{1}{T}\int_0^T x_1(t)\,dt$ computed from random initial conditions for two different values of $T$. The finite-time average is a random variable due to sensitivity to initial conditions; its variance decreases as the averaging window $T$ increases. The vertical dashed line indicates the infinite-time expectation $\lim_{T\to\infty} \frac{1}{T}\int_0^T x_1(t)\,dt$.

the emulator. A Gaussian Process prior places a probability distribution over functions $f : \mathbb{R}^{d_u} \to \mathbb{R}$ such that for any finite collection of inputs, the corresponding function values are jointly Gaussian. We write

$$f \sim \mathcal{GP}(0, k_\lambda),$$

where $k_\lambda(u, u')$ is a covariance kernel parameterized by hyperparameters $\lambda$. In this work we employ the squared-exponential (radial basis) kernel,

$$k_\lambda(u, u') = \sigma^2 \exp\left(-\frac{1}{2}\sum_{j=1}^{d_u} \frac{(u_j - u'_j)^2}{\ell_j^2}\right),$$

where $\sigma^2$ denotes the signal variance and $\{\ell_j\}_{j=1}^{d_u}$ are length-scale parameters controlling smoothness in each coordinate direction.

Let $K_\lambda(U, U) \in \mathbb{R}^{n\times n}$ denote the kernel matrix with entries

$$[K_\lambda(U, U)]_{ij} = k_\lambda(u_i, u_j).$$

Under the GP prior and Gaussian noise model, the data vector $W = (w_1, \ldots, w_n)^\top$ follows

$$W \sim \mathcal{N}(0,\, K_\lambda(U,U) + \gamma^2 I).$$

The kernel parameters $\lambda$ are determined by maximizing the marginal likelihood of the data, equivalently minimizing the negative log-marginal likelihood

$$\mathcal{L}(\lambda) = \frac{1}{2} W^\top K_w^{-1} W + \frac{1}{2} \log\det K_w + \frac{n}{2}\log(2\pi),$$

where

$$K_w = K_\lambda(U,U) + \gamma^2 I.$$

The first term measures data misfit while the log-determinant term penalizes overly complex covariance structures, providing automatic regularization. For a new parameter value $\tilde{u}$, define

$$k_* = K_\lambda(\tilde{u}, U), \qquad k_{**} = k_\lambda(\tilde{u}, \tilde{u}).$$

The predictive distribution for the latent function value $f(\tilde{u})$ is Gaussian with mean and variance

$$m(\tilde{u}) = k_* K_w^{-1} W,$$
$$s^2(\tilde{u}) = k_{**} - k_* K_w^{-1} k_*^\top.$$

We define the emulator prediction as

$$\widehat{f}(\tilde{u}) = m(\tilde{u}).$$

## APPENDIX C: GAUSSIAN PROCESS EMULATOR VALIDATION

The Bayesian inversion framework in this study uses a Gaussian process (GP) emulator to approximate the forward mapping from model parameters to the observation vector of summary statistics. Since the resulting posterior may be sensitive to emulator error, we provide here a concise validation of the GP accuracy and uncertainty quantification.

### C1 Test error versus training set size

The training data used to construct the Gaussian process (GP) emulator are obtained from forward-model evaluations performed during the calibration step described in Section 3.1. In the ensem-

ble Kalman inversion procedure, the forward earthquake-cycle model is evaluated repeatedly at ensemble parameter values $u$. For each forward simulation, we compute the summary-statistics vector $G(u, T)$ from the simulated earthquake sequence over a finite time horizon $T$. The resulting input–output pairs $(u, G(u, T))$ collected during the calibration stage are then used as the training data for the GP emulator. No additional forward simulations are required beyond those already performed for calibration. We trained a sequence of GP emulators using increasing numbers of training simulations (subsampled from available model evaluations) while keeping a fixed, held-out test set for evaluation. For each emulator, we computed the mean squared prediction error on the test set,

$$\mathrm{MSE}_{\mathrm{test}} = \frac{1}{N_{\mathrm{test}}} \sum_{j=1}^{N_{\mathrm{test}}} \left\| y^{(j)} - \hat{y}^{(j)} \right\|_2^2, \tag{C.1}$$

where $y^{(j)}$ denotes the forward-model output (summary-statistics vector) and $\hat{y}^{(j)}$ is the GP predictive mean at the same input. Fig A2 shows that the test error decreases as the training set grows and then approaches a plateau. We therefore select the emulator corresponding to a training set size at which the test error curve begins to plateau, indicating diminishing returns from adding additional training simulations. balancing improved accuracy with diminishing returns from additional training simulations.

## C2 Consistency with observations and predictive uncertainty

As an additional check, we evaluate the selected emulator at the true parameter $u_{true}$ and compare the GP prediction to the observed summary-statistics vector. The GP provides a predictive mean $\mu_{\mathrm{GP}}$ and predictive covariance $\Sigma_{\mathrm{GP}}$ for the $80$-dimensional observation vector. We visualize the comparison by plotting the observations against $\mu_{\mathrm{GP}}$ and indicating pointwise predictive uncertainty using $\pm 1\,\sigma$, where

$$\sigma_i = \sqrt{(\Sigma_{\mathrm{GP}})_{ii}}, \qquad i = 1, \ldots, 80. \tag{C.2}$$

Figure A3 shows close agreement between the observation vector and the GP predictive mean across all components, with deviations that are small relative to the indicated predictive uncer-

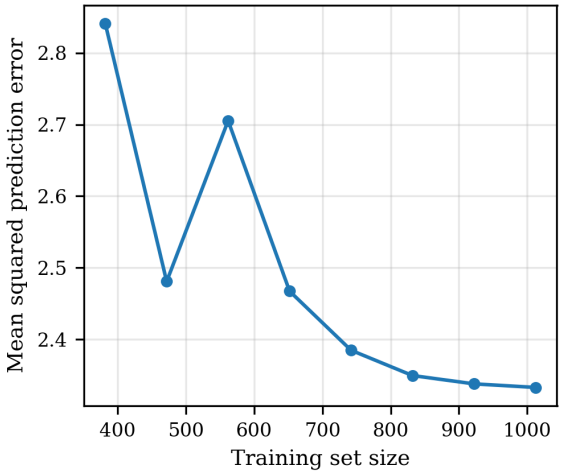


**Figure A2.** GP emulator test error versus training set size. The plotted quantity is the mean squared prediction error on a test set. The error decreases with increasing training simulations and then plateaus; we therefore select an emulator trained with 922 simulations, which lies near this plateau, for subsequent inference.

tainty. Together with the plateau behavior in Figure A2, this supports that the GP emulator is sufficiently accurate for the inference results reported in the main text.

We observe that the predictive standard deviation produced by the GP emulator is larger than the empirical variability obtained when the forward model is evaluated at the true parameter $u_{\mathrm{true}}$. This discrepancy can be explained by the manner in which the training data are generated.

Each training pair $(u, G(u))$ requires a full forward-model simulation. Because these simulations are computationally expensive, we restrict each run to a maximum wall-clock time of 8 CPU hours and record the summary statistics obtained within this time window. Ideally, the forward model would be integrated over the full simulation horizon ( 5500 years) to obtain statistically converged summary statistics. Under the imposed time constraint, however, some parameter configurations are simulated over shorter effective time horizons, leading to increased variability in the estimated observables.

As a result, the training outputs used to construct the emulator contain additional numerical and sampling noise relative to the fully converged model outputs. The GP regression therefore learns a mapping with an inflated effective noise level, which manifests as larger predictive variance.

From a Bayesian inversion perspective, this effect leads to a conservative posterior: the uncer-

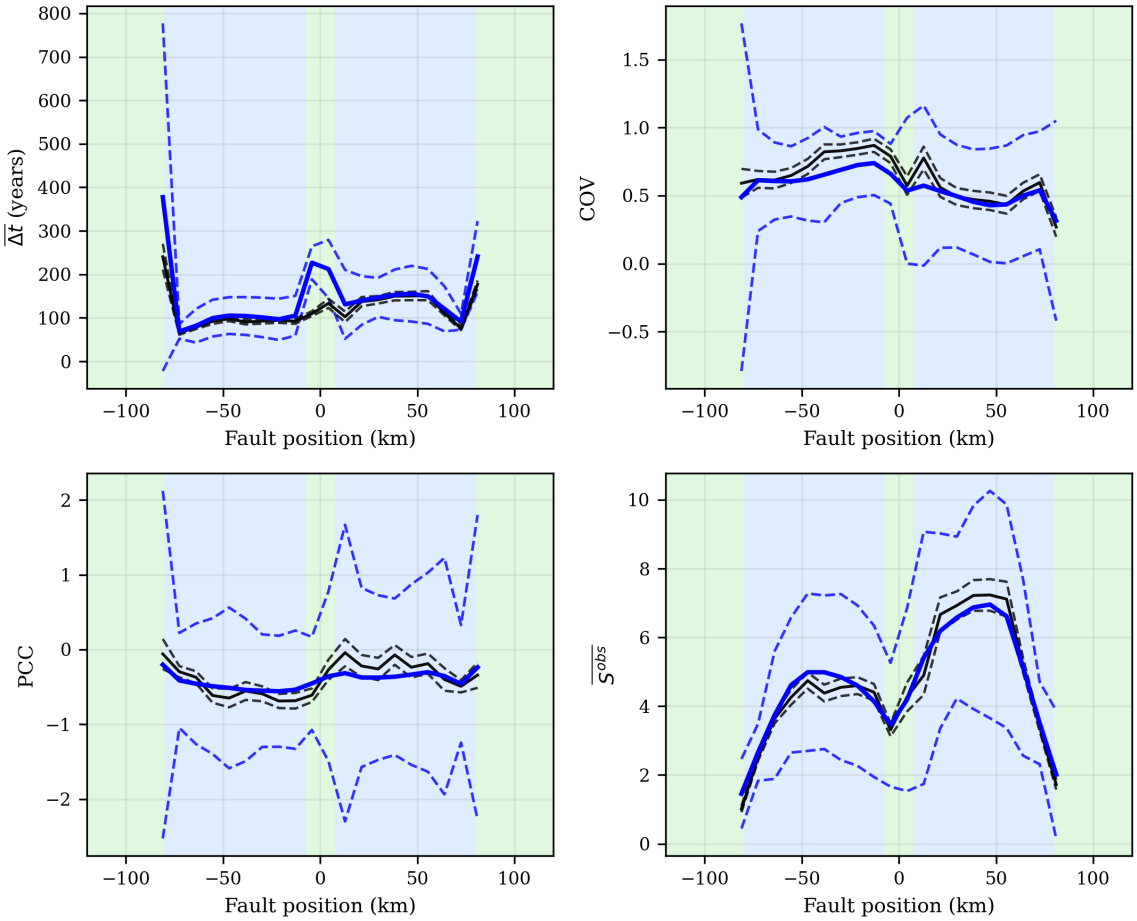


**Figure A3.** Comparison of the observed summary-statistics vector (black) with the Gaussian process (GP) predictive mean of the selected emulator (blue), evaluated at the true parameter $u^{\text{true}}$. Dashed blue curves indicate pointwise predictive uncertainty, $\mu_{\text{GP}} \pm 1\sigma_{\text{GP}}$, where $\sigma_{\text{GP}} = \sqrt{\text{diag}(\Sigma_{\text{GP}})}$ is computed from the diagonal of the GP predictive covariance. Dashed black curves indicate pointwise observational uncertainty, $y^{\dagger} \pm 1\sigma_{\text{obs}}$, where $\sigma_{\text{obs}}$ denotes the standard deviation of the forward-model noise in Eq. 1, estimated empirically. Panels show mean interevent time, COV, PCC, and mean coseismic slip. Shaded bands follow the convention of Fig. 1.

tainty in parameter space is slightly overestimated due to the inflated emulator noise. In practical applications where forward simulations can be run to the full designed time horizon $T$, the resulting emulator would exhibit reduced predictive variance and yield a correspondingly tighter posterior distribution.

## References

Anderson, J. L. (2001). An Ensemble Adjustment Kalman Filter for Data Assimilation.

Avouac, J.-P. (2015). From Geodetic Imaging of Seismic and Aseismic Fault Slip to Dynamic Modeling of the Seismic Cycle. *Annual Review of Earth and Planetary Sciences*, 43(1):233–271. _eprint: https://doi.org/10.1146/annurev-earth-060614-105302.

Barbot, S. (2019). Slow-slip, slow earthquakes, period-two cycles, full and partial ruptures, and deterministic chaos in a single asperity fault. *Tectonophysics*, 768:228171.

Barbot, S., Lapusta, N., and Avouac, J.-P. (2012). Under the Hood of the Earthquake Machine: Toward Predictive Modeling of the Seismic Cycle. *Science*, 336(6082):707–710.

Brockwell, P. J. and Davis, R. A. (2016). *Introduction to Time Series and Forecasting*. Springer Texts in Statistics. Springer International Publishing, Cham.

Brodsky, E. E. and Farge, G. (2026). How earthquakes organize stress. *Proceedings of the National Academy of Sciences*, 123(6):e2530754123.

Burbank, D. W. and Anderson, R. S. (2013). *Tectonic geomorphology*. EEGS 1720 South Bellaire, Suite 110, Denver, CO 80222-4303, USA.

Calvello, E., Reich, S., and Stuart, A. M. (2025). Ensemble kalman methods: A mean-field perspective. *Acta Numerica*, 34:123–291.

Chada, N. K., Chen, Y., and Sanz-Alonso, D. (2021). Iterative ensemble Kalman methods: A unified perspective with some new variants. *Foundations of Data Science*, 3(3):331–369.

Chen, T. and Lapusta, N. (2009). Scaling of small repeating earthquakes explained by interaction of seismic and aseismic slip in a rate and state fault model. *Journal of Geophysical Research: Solid Earth*, 114(B1). _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1029/2008JB005749.

Chen, Y. and Oliver, D. S. (2012). Ensemble Randomized Maximum Likelihood Method as an Iterative Ensemble Smoother. *Mathematical Geosciences*, 44(1):1–26.

Cleary, E., Garbuno-Inigo, A., Lan, S., Schneider, T., and Stuart, A. M. (2021). Calibrate, emulate, sample. *Journal of Computational Physics*, 424:109716.

Corbi, F., Mastella, G., Tinti, E., Rosenau, M., Sandri, L., Pardo, S., and Funiciello, F. (2024). Asperity Size and Neighboring Segments Can Change the Frictional Response

985 and Fault Slip Behavior: Insights From Laboratory Experiments and Numerical Simula-
986 tions. *Journal of Geophysical Research: Solid Earth*, 129(1):e2023JB026594. _eprint:
987 https://onlinelibrary.wiley.com/doi/pdf/10.1029/2023JB026594.

988 Cotter, S. L., Roberts, G. O., Stuart, A. M., and White, D. (2013). MCMC Methods for Functions:
989 Modifying Old Algorithms to Make Them Faster. *Statistical Science*, 28(3):424–446.

990 Dieterich, J. (1994). A constitutive law for rate of earthquake production and its application
991 to earthquake clustering. *Journal of Geophysical Research: Solid Earth*, 99(B2):2601–2618.
992 _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1029/93JB02581.

993 Dieterich, J. H. (1992). Earthquake nucleation on faults with rate-and state-dependent strength.
994 *Tectonophysics*, 211(1):115–134.

995 Doucet, A., Gordon, N., and Krishnamurthy, V. (2001). Particle filters for state estimation of
996 jump Markov linear systems. *IEEE Transactions on Signal Processing*, 49(3):613–624.

997 Dunbar, O. R. A., Bieli, M., Garbuno-I nigo, A., Howland, M., de Souza, A. N., Mansfield,
998 L. A., Wagner, G. L., Efrat-Henrici, N., et al. (2024). Calibrateemulatesample.jl: Accelerated
999 parametric uncertainty quantification. *Journal of Open Source Software*, 9(97):6372.

1000 Dunbar, O. R. A., Duncan, A. B., Stuart, A. M., and Wolfram, M.-T. (2022a). Ensemble Inference
1001 Methods for Models With Noisy and Expensive Likelihoods. arXiv:2104.03384 [math].

1002 Dunbar, O. R. A., Garbuno-Inigo, A., Schneider, T., and Stuart, A. M. (2021a). Cal-
1003 ibration and Uncertainty Quantification of Convective Parameters in an Idealized GCM.
1004 *Journal of Advances in Modeling Earth Systems*, 13(9):e2020MS002454. _eprint:
1005 https://agupubs.onlinelibrary.wiley.com/doi/pdf/10.1029/2020MS002454.

1006 Dunbar, O. R. A., Garbuno-Inigo, A., Schneider, T., and Stuart, A. M. (2021b). Calibration and
1007 uncertainty quantification of convective parameters in an idealized gcm. *Journal of Advances*
1008 *in Modeling Earth Systems*, 13(9):e2020MS002454.

1009 Dunbar, O. R. A., Lopez-Gomez, I., Garbuno-Iñigo, A., Huang, D. Z., Bach, E., and long Wu,
1010 J. (2022b). EnsembleKalmanprocesses.jl: Derivative-free ensemble-based model calibration.
1011 *Journal of Open Source Software*, 7(80):4869.

1012 Evensen, G. (2018). Accounting for model errors in iterative ensemble smoothers.

arXiv:1806.00237 [physics].

Fukushima, R., Kano, M., Hirahara, K., Ohtani, M., Im, K., and Avouac, J.-P. (2025). Physics-Informed Deep Learning for Estimating the Spatial Distribution of Frictional Parameters in Slow Slip Regions. *Journal of Geophysical Research: Solid Earth*, 130(5):e2024JB030256. _eprint: https://agupubs.onlinelibrary.wiley.com/doi/pdf/10.1029/2024JB030256.

Garbuno-Inigo, A., Hoffmann, F., Li, W., and Stuart, A. M. (2020). Interacting Langevin Diffusions: Gradient Structure and Ensemble Kalman Sampler. *SIAM Journal on Applied Dynamical Systems*, 19(1):412–441.

Gjini, R., Morzfeld, M., Dunbar, O. R. A., and Schneider, T. (2025). The Ensemble Kalman Inversion Race. arXiv:2511.15853 [physics].

Griffin, J. D., Stirling, M. W., and Wang, T. (2020). Periodicity and Clustering in the Long-Term Earthquake Record. *Geophysical Research Letters*, 47(22):e2020GL089272. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1029/2020GL089272.

Heimisson, E. R. and Segall, P. (2018). Constitutive Law for Earthquake Production Based on Rate-and-State Friction: Dieterich 1994 Revisited. *Journal of Geophysical Research: Solid Earth*, 123(5):4141–4156. _eprint: https://agupubs.onlinelibrary.wiley.com/doi/pdf/10.1029/2018JB015656.

Holland, M. and Melbourne, I. (2007). Central limit theorems and invariance principles for Lorenz attractors. *Journal of the London Mathematical Society*, 76(2):345–364. _eprint: https://londmathsoc.onlinelibrary.wiley.com/doi/pdf/10.1112/jlms/jdm060.

Huang, D. Z., Huang, J., Reich, S., and Stuart, A. M. (2022). Efficient Derivative-free Bayesian Inference for Large-Scale Inverse Problems. arXiv:2204.04386 [math].

Iglesias, M. and Yang, Y. (2021). Adaptive regularisation for ensemble Kalman inversion. *Inverse Problems*, 37(2):025008.

Iglesias, M. A., Law, K. J. H., and Stuart, A. M. (2013). Ensemble Kalman methods for inverse problems. *Inverse Problems*, 29(4):045001.

Julier, S., Uhlmann, J., and Durrant-Whyte, H. (2000). A new method for the nonlinear transformation of means and covariances in filters and estimators. *IEEE Transactions on Automatic*

*Control*, 45(3):477–482.

Kaneko, Y., Avouac, J.-P., and Lapusta, N. (2010). Towards inferring earthquake patterns from geodetic observations of interseismic coupling. *Nature Geoscience*, 3(5):363–369. Number: 5.

Kaveh, H., Avouac, J. P., and Stuart, A. M. (2025a). Data Assimilation in Machine-Learned Reduced-order Model of Chaotic Earthquake Sequences. *Geophysical Journal International*, page ggaf518.

Kaveh, H., Avouac, J. P., and Stuart, A. M. (2025b). Spatiotemporal forecast of extreme events in a chaotic model of slow slip events. *Geophysical Journal International*, 240(2):870–885.

Kaveh, H., Batlle, P., Acosta, M., Kulkarni, P., Bourne, S. J., and Avouac, J. P. (2023). Induced Seismicity Forecasting with Uncertainty Quantification: Application to the Groningen Gas Field. *Seismological Research Letters*, 95(2A):773–790.

Lapusta, N. and Liu, Y. (2009). Three-dimensional boundary integral modeling of spontaneous earthquake sequences and aseismic slip. *Journal of Geophysical Research: Solid Earth*, 114(B9). _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1029/2008JB005934.

Lapusta, N., Rice, J. R., Ben-Zion, Y., and Zheng, G. (2000). Elastodynamic analysis for slow tectonic loading with spontaneous rupture episodes on faults with rate- and state-dependent friction. *Journal of Geophysical Research: Solid Earth*, 105(B10):23765–23789. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1029/2000JB900250.

Lorenz, E. N. (1963). Deterministic Nonperiodic Flow.

Lütkepohl, H. (1991). *Introduction to Multiple Time Series Analysis*. Springer, Berlin, Heidelberg.

Magen, Y., May, D. A., and Gabriel, A.-A. (2025). Reduced-order modelling of Cascadia's slow slip cycles.

Neumaier, A. and Schneider, T. (2001). Estimation of parameters and eigenmodes of multivariate autoregressive models. *ACM Trans. Math. Softw.*, 27(1):27–57.

Rice, J. R. (1993). Spatio-temporal complexity of slip on a fault. *Journal of Geophysical Research: Solid Earth*, 98(B6):9885–9907. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1029/93JB00191.

Riley, D. M., Gutiérrez-Oribio, D., and Stefanou, I. (2025). Slip, Differentiate, Observe: State and Parameter Estimation for Rate and State Friction from Noisy Data. arXiv:2512.01705 [physics].

Rubin, A. M. and Ampuero, J.-P. (2005). Earthquake nucleation on (aging) rate and state faults. *Journal of Geophysical Research: Solid Earth*, 110(B11). _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1029/2005JB003686.

Ruina, A. (1983). Slip instability and state variable friction laws. *Journal of Geophysical Research: Solid Earth*, 88(B12):10359–10370. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1029/JB088iB12p10359.

Schillings, C. and Stuart, A. M. (2016). Analysis of the ensemble Kalman filter for inverse problems. arXiv:1602.02020 [math].

Schneider, T., Stuart, A. M., and Wu, J.-L. (2021). Learning stochastic closures using ensemble Kalman inversion. *Transactions of Mathematics and Its Applications*, 5(1):tnab003.

Stiernström, V., Almquist, M., and Dunham, E. M. (2024). Adjoint-based inversion for stress and frictional parameters in earthquake modeling. arXiv:2310.12279 [cs, math].

Subrahmanya Nairy, K. and Aruna Rao, K. (2003). Tests of Coefficients of Variation of Normal Population. *Communications in Statistics - Simulation and Computation*, 32(3):641–661. _eprint: https://doi.org/10.1081/SAC-120017854.

Temam, R. (1997). *Infinite-Dimensional Dynamical Systems in Mechanics and Physics*, volume 68 of *Applied Mathematical Sciences*. Springer, New York, NY.

Tucker, W. (1999). The Lorenz attractor exists. *Comptes Rendus de l'Académie des Sciences - Series I - Mathematics*, 328(12):1197–1202.

Wang, T., Griffin, J. D., Brenna, M., Fletcher, D., Zeng, J., Stirling, M., Dillingham, P. W., and Kang, J. (2024). Earthquake forecasting from paleoseismic records. *Nature Communications*, 15(1):1944.