Fast Propagation of Epistemic Uncertainty in Seismic Hazard via Adaptive Importance Sampling

Soung Eil Houng*[1], Luis Ceferino[1], and Norman Abrahamson[1]

ABSTRACT

The conventional logic-tree approach used to represent epistemic uncertainty in probabilistic seismic hazard analysis (PSHA) is both computationally intensive and prone to introducing bias in hazard estimates. We propose a Gaussian Population Monte Carlo Adaptive Importance Sampling (G-PMC AIS) framework for PSHA computation to efficiently propagate epistemic uncertainty which uses the continuous distribution of uncertain model parameters. First, we demonstrate that the logic-tree approach introduces significant bias in mean and fractile hazards due to the ad hoc discretization and becomes computationally infeasible as the number of variables grows. We then present a mathematical formulation of PSHA using importance sampling (IS), illustrating how the IS optimal density enables efficient estimation of mean hazard and that its marginalization supports the fast computation of fractile hazards and sensitivity analysis. To efficiently approximate the optimal density and its marginalization, we adopt the G-PMC AIS framework within the IS formulation for PSHA. Numerical experiments show that our approach accelerates the computation of mean and fractile hazards and sensitivity metrics by factors of 13–224 compared to the logic-tree method, and 123–3,775 compared to Monte Carlo methods with continuous distributions. These gains are achieved while keeping the coefficient of variation (COV) of mean hazard below 1% and the Kolmogorov–Smirnov distance (K-S D) of fractile estimates from the true fractile below 5%. The proposed framework significantly reduces computational costs without compromising accuracy and is broadly applicable to PSHA projects.

KEY POINTS

- The logic-tree approach for epistemic uncertainty in PSHA is biased and computationally expensive.
- We propose a framework to efficiently compute accurate mean and fractile hazards and perform sensitivity analysis.
- The method is broadly applicable to PSHA projects with numerous epistemic uncertainty variables.

Supplemental Material

INTRODUCTION

Probabilistic seismic hazard analysis (PSHA) estimates the annual probability of ground motion exceedance at various intensity levels and serves as a foundational tool for applications such as the seismic design of new structures, evaluation of existing buildings, and regional risk analyses (Cornell, 1968; U. S. Nuclear Regulatory Commission, 2007; McGuire, 2008; ASCE, 2022; Kennedy et al., 1980; Ceferino et al., 2020; Silva et al., 2020; Baker et al., 2021; Papadopoulos and Bazzurro, 2021).

In PSHA, the annual exceedance rates (or probabilities) at various ground motion levels are calculated. This is achieved by coupling two primary input models: a seismic source model and a ground motion model. The seismic source model defines potential earthquake scenarios, including occurrence rates and source-to-site geometries, while the ground motion model estimates the intensity at the site for a specified scenario (Cornell, 1968; Esteva and Ordaz, 1963; Blume and Kiremidjian, 1979; Baker et al., 2021). Given a model, the randomness in earthquake magnitude, location, and resulting ground motion, is integrated to determine the annual exceedance rate, the aleatory uncertainty of PSHA.

However, the input models in PSHA are inherently uncertain due to limitations in current scientific understanding (Marzocchi et al., 2015; Abrahamson and Bommer, 2005; Bommer and Scherbaum, 2008; Baker et al., 2021; Hanks et al., 2009). These model-related uncertainties are termed epistemic uncertainties. To account for them, parameter distributions are defined and multiple alternative hazard curves are generated. For instance, hazard curves are evaluated across a range of seismic source parameters such as the maximum magnitude, Gutenberg-Richter b-value, earthquake rate, and across alternative ground motion models. The mean hazard curve is then computed by averaging these alternatives, while fractile hazard curves (e.g., 16th, 50th, and 84th percentiles) are used to represent the uncertainty range. These mean and fractile hazard curves are critical outputs in modern PSHA practice.

Numerical evaluation of the PSHA mean and fractiles is traditionally achieved via double-nested integration (Baker et al., 2021). The inner integration evaluates aleatory uncertainty, the individual hazard curves, typically using Riemann sums or Monte Carlo (MC) simulations. The outer integration propagates epistemic uncertainty in model parameters such as the maximum magnitude, b-value, and ground motion model parameters. The latter is commonly implemented via a logic-tree framework, where branches represent alternative models or parameters and are weighted based on expert judgment (Assatourians and Atkinson, 2013; Musson, 2000; Ordaz et al., 2021; Hale et al., 2018).

Despite its widespread adoption, the logic-tree approach presents limitations in both computational efficiency and accuracy. From a computational standpoint, the nested integration process is highly demanding. Two primary factors contribute to the computational burden: 1) the time-intensive evaluation of individual hazard curves, especially at low exceedance probabilities (Hale et al., 2018; Houng and Ceferino, 2025), and 2) the exponential growth in the number of individual hazard evaluations required as the number of epistemic variables increases (Kwong and Jaiswal, 2023). This "curse of dimensionality" often forces analysts to simplify the logic tree by reducing the number of branches, which, in turn, necessitates additional sensitivity analyses, further increasing computational costs.

In terms of accuracy, the logic-tree method approximates continuous epistemic uncertainty distributions using discrete representations. This discretization can introduce bias in the resulting hazard curves, especially when the number of discrete points per variable is limited, typically to three to five (Bommer, 2012). Furthermore, the process is subjective, as different experts may produce different discretizations for the same underlying distribution. Although strategies exist to reduce discretization bias (Keefer and Bodily, 1983; Miller and Rice, 1983), they cannot eliminate it entirely. Increasing the number of discretization points could improve accuracy, but this significantly increases computational cost. Specifically, the complexity scales exponentially as $O(n^d)$, where n is the number of discretization points per variable and d is the number of independent epistemic variables. For instance, increasing the number of points from three to five across ten variables results in a 165-fold increase in computational cost ($5^{10}/3^{10}$), underscoring the interdependence between accuracy and efficiency in PSHA.

To mitigate these challenges, several studies have explored Monte Carlo-based importance sampling (IS) methods (Crowley and Bommer, 2006; Kiremidjian et al., 2007; Jayaram and Baker, 2010; Han and Davidson, 2012; Manzour et al., 2016; Christou et al., 2018; Rahimi and Mahsuli, 2019). IS improves efficiency by focusing sampling efforts on regions that contribute most to the integration result, thereby reducing the number of required simulations. However, identifying these important regions in PSHA is nontrivial, motivating the use of adaptive importance sampling algorithms that iteratively refine the proposal distribution (Houng and Ceferino, 2025). Houng and Ceferino (2025) proposed an adaptive IS framework that closely approximates the optimal sampling density and significantly accelerates individual hazard and disaggregation computations. However, this approach assumes independence among variables, an assumption invalidated in the presence of epistemic dependencies, rendering it unsuitable for computing mean and fractile hazards.

Concurrently, Lacour and Abrahamson (2019, 2021, 2025) proposed the use of polynomial chaos expansion (PCE) to propagate epistemic uncertainty in ground motion models. Their method effectively reduced computational costs, particularly for non-ergodic models, by providing surrogate models of ground motion medians. However, its applicability is limited, as extending PCE to other epistemic variables—such as seismic source parameters and ground motion standard deviations—requires closed-form expansions, which are not always attainable.

In this paper, we present a generalized computational framework for computing the mean and fractile hazard curves along with sensitivity analysis in PSHA using adaptive importance sampling. Our method builds upon the adaptive IS framework developed

by Houng and Ceferino (2025), extending its applicability beyond individual hazard and disaggregation computations. We begin by mathematically formulating PSHA to establish the groundwork for our approach. We then analyze the bias and inefficiency introduced by the logic-tree framework, particularly due to its discretization of continuous distributions. Next, we introduce a generalized adaptive IS method and incorporate it into the PSHA formulation to develop a novel framework based on Gaussian Population Monte Carlo Adaptive Importance Sampling (G-PMC AIS). The proposed approach efficiently estimates both mean and fractile hazards, while enabling sensitivity analysis. We demonstrate its computational advantages through numerical examples that incorporate commonly used epistemic uncertainties, including the uncertainties in both seismic source and ground motion models.

PSHA MATHEMATICAL FORMULATION

Individual Hazard

Given the seismic source and ground motion models, the individual hazard for a single source can be expressed as:

$$\lambda(a|\theta) = \nu(\theta) \iiint_{m,r,\varepsilon} I(u(m,r,\varepsilon;\theta) > a) f_{M,R,\mathscr{E}|\Theta}(m,r,\varepsilon|\theta) dm dr d\varepsilon \tag{1}$$

, where θ denotes the model parameters associated with epistemic uncertainty, such as the Gutenberg–Richter b-value, $m_{\max}$, and alternative ground motion models. The variables (m, r, ε) represent magnitude, distance, and the ground motion residual, respectively, encapsulating the aleatory uncertainty in hazard. The function $f_{M,R,\mathscr{E}|\Theta}(m,r,\varepsilon|\theta)$ denotes the conditional probability density function (PDF) of the aleatory variables given θ. The term $\nu(\theta)$ represents the earthquake occurrence rate under θ, and $u(m,r,\varepsilon;\theta)$ is the simulated ground motion intensity. The indicator function $I(\cdot)$ evaluates to 1 if $u(\cdot)$ exceeds the ground motion level a.

The aleatory variables (m, r, ε) can be collectively denoted as ϕ, leading to the simplified expression:

$$\lambda(a|\theta) = \nu(\theta) \int_{\phi} I(u(\phi;\theta) > a) f_{\Phi|\Theta}(\phi|\theta) d\phi \tag{2}$$

Mean and Fractile Hazard

The mean hazard is defined as the expectation of the individual hazard over the epistemic uncertainty variables θ:

$$\lambda(a) = \int_{\theta} \lambda(a|\theta) f_{\Theta}(\theta) d\theta \tag{3}$$

Here, $\lambda(a)$ represents the mean hazard, and $f_\Theta(\theta)$ is the PDF of the epistemic uncertainty variables. Note that equation Eq. (3) corresponds to the discrete form used in logic-tree-based hazard computation, $\lambda(a) = \sum_{i=1}^{N_\theta} \lambda_i(a|\theta_i)\,\omega_i(\theta_i)$, where N_θ is the total number of logic-tree end branches, and $\omega_i(\theta_i)$ denotes the weight associated with the i-th branch.

The distribution of individual hazard with respect to θ, denoted $\lambda(a|\theta)$, represents the range of possible hazard values. Its cumulative distribution function (CDF) is given by:

$$F_\lambda(x) = P[\lambda(a|\theta) \leq x] \tag{4}$$

In practical applications, the probability on the right-hand side of Eq. (4) is computed empirically by cumulatively summing the weights in ascending order of individual hazard values.

The p-th fractile of the hazard, denoted $\lambda_p(a)$, is obtained from the CDF as the smallest value for which the CDF reaches or exceeds $p\%$:

$$\lambda_p(a) = \inf\left\{x : F_\lambda(x) \geq \frac{p}{100}\right\} \tag{5}$$

, where $\lambda_p(a)$ denotes the p-th fractile of the hazard, $F_\lambda(x)$ is the CDF of the individual hazard, and $\inf\{\cdot\}$ identifies the minimum hazard value satisfying the fractile condition. The accuracy of $\lambda_p(a)$ depends on the precision of the estimated $F_\lambda(x)$.

PROPAGATION OF EPISTEMIC UNCERTAINTY THROUGH LOGIC TREES

We examine the issues of accuracy and computational efficiency associated with the traditional approach to propagating epistemic uncertainty in PSHA: the logic tree. We study the case where the continuous distributions of the epistemic uncertainty variables are approximated using discrete distributions through logic trees.

Hazard Bias

We show the potential bias introduced by discretizing a continuous distribution through a simple numerical example (Fig. 1). We consider a point source located 35 km from the site, generating a single characteristic earthquake. For illustration, the simple base ground motion model (GMM) by Sadigh et al. (1997) is used. Three epistemic uncertainty variables are incorporated: the characteristic magnitude, median ground motion, and standard deviation of the GMM. The underlying distribution for the characteristic magnitude is assumed to be normally distributed with a mean of 7.0 and a standard deviation of 0.2, $\mathcal{N}(7.0, 0.2)$. The natural logarithm of the median and the standard deviation (ln) of ground motion are modeled as normally distributed deviations from the base GMM of Sadigh et al. (1997), both with standard deviations of 0.05, i.e., $\mathcal{N}(0, 0.05)$ and $\mathcal{N}(0, 0.05)$, respectively. The distributions are not truncated to upper or lower limits since the ranges yielding unrealistic values, e.g., negative standard deviation,

do not affect the result at the hazard range considered in this example. The annual occurrence rate of the characteristic earthquake is set to 0.01.

To construct the logic tree, we consider three discretization strategies (Table 1 and 2). The Extended Pearson-Tukey three-point approximation of Keefer and Bodily (1983) (denoted KB83), the three-point approximation used in the 2023 USGS seismic hazard model (Petersen et al., 2024) (denoted Pea24), and the five-point approximation proposed by Miller and Rice (1983) (denoted MR83). These approximation schemes have been widely used in many practical applications, and they aim to match the first few statistical moments of the original distribution (Keefer and Bodily, 1983; Miller and Rice, 1983; Coppersmith et al., 2012; Al Atik and Youngs, 2014).

The benchmark solution is obtained using continuous-distribution-based Monte Carlo (MC) simulation with $N = 5 \times 10^8$ samples, which ensures a coefficient of variation (COV) $\leq 1\%$ down to a hazard level of 2×10^{-7}/yr (Houng and Ceferino, 2025).

The mean, 16th, and 84th fractile hazards from each approximation scheme, along with their relative errors compared to the benchmark, are shown in Figure 2. As expected, increasing the number of discrete points improves accuracy. For the five-point MR83 approximation, errors remain within 5.28 % in ground motion intensity up to 0.5 g for the mean, 16th, and 84th fractile hazards. However, a tendency to diverge from the benchmark at higher ground motions is still observed.

The KB83 approximation yields mean hazard estimates comparable in accuracy to the five-point approximation, despite requiring only 22 % of the computational cost ($= 3^3/5^3$). It also produces reasonably accurate 84th fractile estimates, with errors under 15 %, and does not show significant divergence at higher ground motion levels within the range considered. However, for the 16th fractile, KB83 significantly underestimates hazard values beyond 0.1 g (exceedance rate of 6.5×10^{-3}/yr). Although this exceedance level corresponds to a relatively high probability event (probability of exceedance without accounting for earthquake rate is 0.65), the KB83 approximation still deviates substantially, indicating difficulties in accurately capturing the lower tail of the hazard distribution. Note that KB83 underestimated the 16th fractile, suggesting an exaggerated range in the hazard curve, driven by a low bias in the lower fractiles.

Among the three, Pea24 performed the worst. The mean hazard deviates from the benchmark by over 5% at ground motion levels exceeding 0.2 g (hazard level of 1.1×10^{-3}/yr). The 16th and 84th fractile estimates diverge even earlier, at ground motion levels of 0.1 g, corresponding to hazard levels of 1.8×10^{-3}/yr and 3.9×10^{-3}/yr, respectively. Notably, Pea24 overestimates the 16th fractile and underestimates the 84th fractile, suggesting that the scheme underrepresents the full variability of the hazard. Although both KB83 and Pea24 use the same number of branches, their differing weighting schemes lead to substantially different hazard estimates. At 0.5 g, the mean, 16th, and 84th fractile hazard estimates differ by factors of approximately 1.8, 2.7, and 1.8, respectively.

We also evaluated the error in ground motion intensities corresponding to various hazard levels (Figure 2 (g)–(i)). These errors were generally smaller than the errors in hazard given ground motion intensities (Figure 2 (d)–(f)). This is because the largest

discrepancies in hazard occur at high ground motion intensities, where the hazard curve is steep; thus, a perturbation in the horizontal axis (ground motion intensity) results in a smaller shift than the corresponding perturbation in the vertical axis (hazard).

Among the tested methods, the five-point MR83 approximation yielded the best performance, while the Pea24 approximation performed the worst as shown in error of the hazard level. Within the annual exceedance rate range of 10^{-2} to 10^{-5}/year, the maximum errors in mean, 16th, and 84th percentile ground motion intensities for MR83 were 0.27%, 0.43%, and 0.85%, respectively. For the KB83 three-point approximation, the corresponding errors were 0.25%, 3.4%, and 3.0%. In contrast, the Pea24 three-point approximation resulted in significantly larger errors: 6.1%, 5.1%, and 7.6%, respectively.

This example highlights that discretization in logic trees introduces bias into hazard estimates, especially at large ground motions (low exceedance probabilities) and in the lower fractile estimates. Any existing approximation scheme is not sufficient to accurately propagate the full distribution of the underlying continuous distribution, i.e., they do not fully capture the entire range of the hazard. Moreover, the example demonstrates that discretization strategies without statistical justification can result in significantly biased hazard estimates even when they approximate the same underlying distribution.

In summary, this example reveals several limitations of logic-tree-based approximations in PSHA computation: First, the underlying continuous distributions of epistemic uncertainty variables can be significantly distorted when arbitrary weights are assigned, as shown in Figure 2. Second, even the most statistically grounded discretizations cannot fully represent the underlying distributions, introducing errors in hazard estimates. This issue could be further exacerbated when a few epistemic variables dominate hazard sensitivity (Lacour and Abrahamson, 2019). In such cases, discretizing these key variables using only a few points causes an abrupt change in the hazard only when their values change, resulting in a step-like distribution of hazard that deviates significantly from the true fractile.

Computational challenges

In the previous numerical example, we considered only three epistemic uncertainty variables in the hazard calculations, requiring 27 and 125 hazard evaluations for the three- and five-point discretization schemes, respectively. However, practical PSHA applications involve more epistemic variables, resulting in substantial computational demands even after discretization. For instance, the total number of end branches in the 2023 USGS hazard maps for Hawaii is approximately 2.9 trillion (Kwong and Jaiswal, 2023), rendering direct computation infeasible. Also, for less extreme cases, PSHA using UCERF3 results in ∼170,000 logic tree end branches, which is also required to be reduced to calculate the risk to a portfolio of buildings (Porter et al., 2025). These examples clearly illustrate the curse of dimensionality inherent in logic-tree-based propagation of uncertainty, where the computational cost increases exponentially with the number of dimensions ($O(n^d)$).

To demonstrate this effect, we present the estimated computation time for full hazard evaluations with increasing numbers of epistemic uncertainty variables, assuming an individual hazard computation time of five seconds (Figure S1). As seen in the

TABLE 1
Approximation strategies commonly used in practical PSHA, tested in the numerical example for logic-tree

Three points (KB83)			Five points (MR83)		
p	**z**	**weight**	**p**	**z**	**weight**
0.05	-1.65	0.185	0.0349	-1.812	0.1011
0.50	0	0.630	0.2117	-0.800	0.2443
0.95	1.65	0.185	0.5000	0	0.3092
			0.7883	0.800	0.2443
			0.9651	1.812	0.1011

TABLE 2
Approximation strategies used in Petersen et al. (2024)

Three points (Pea24)		
p	**z**	**weight**
0.16	-1	0.185
0.50	0	0.630
0.84	1	0.185

figure, with ten epistemic variables and a five-point discretization, the total number of hazard evaluations becomes 5^{10}, making the computation time without parallelization approximately one year.

Another critical computational consideration—often overlooked in PSHA practice—is the number of marginal hazard evaluations required for accurate computation of each individual hazard curve. To ensure the reliability of the individual hazard estimate, a sufficient number of ground motion slices (in Riemann sum approaches) or samples (in Monte Carlo approaches) must be used, and the required number varies depending on the exceedance probability associated with the target ground motion level.

For example, consider a case incorporating uncertainty in the median ground motion model. Assume that three logic-tree branches are defined: a base case with median ground motion of 0.1 g, and upper and lower cases representing ± 0.3 shifts in the logarithmic ground motion, corresponding to median ground motions of 0.075 g and 0.135 g, respectively (Figure S2). Assuming an aleatory variability of $\sigma = 0.6$ (natural logarithm), the exceedance probabilities at a target ground motion of 0.246 g are 15.9%, 6.7%, and 2.3% for the base, upper, and lower cases, respectively—a variation by a factor of 7. This discrepancy becomes even more pronounced at higher ground motion levels. For instance, at 0.45 g, the exceedance probabilities become 2.23%, 0.62%, and 0.135%, resulting in a variation by a factor of 17. Such differences in exceedance probability directly influence the number of sufficient slices in Riemann sum or sufficient samples required in Monte Carlo simulation to achieve a consistent level of accuracy across branches (Houng and Ceferino, 2025). For example, the number of Monte-Carlo samples for the accurate calculation is inversely proportional to the corresponding exceedance probability ($N \sim 1/\lambda$). Thus, to maintain equivalent accuracy, the lower-case branch would require 17 times more samples than the upper-case.

When Does the Discrete Approximation Fail?

In PSHA, a discrete approximation can break down under two circumstances.

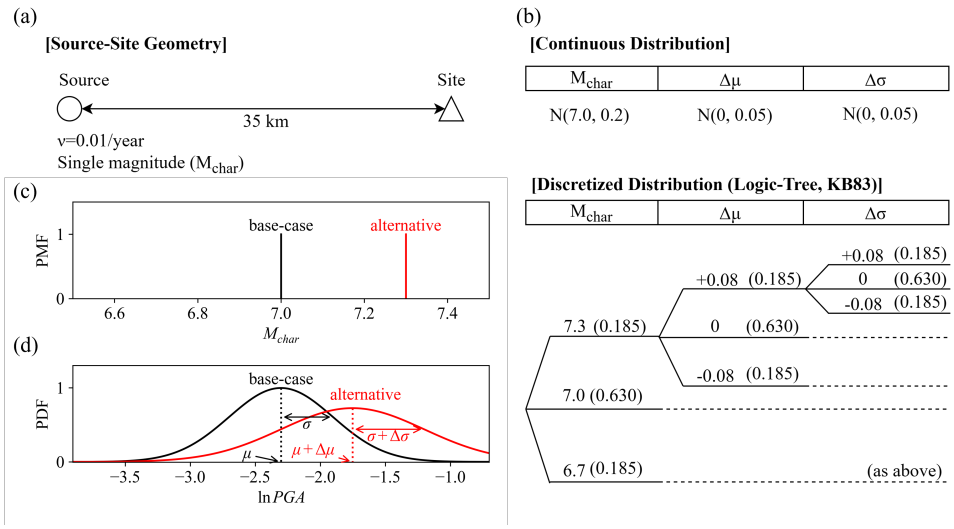

Figure 1. (a) Geometry of the seismic source and site location used in the numerical example to illustrate hazard bias through logic trees; (b) the assumed continuous distributions of the epistemic uncertainty variables, along with an example of the discretized distribution used to construct a logic tree based on the KB83 approximation (Table 1). Three epistemic uncertainty variables are incorporated: the characteristic earthquake magnitude (M_{char}), and the median and standard deviation of the GMM ($\Delta\mu$ and $\Delta\sigma$). (c) Magnitude probability mass function, and (d) peak ground acceleration (PGA) PDF for the base case (black) and one possible alternative case (red). In this example, the base case corresponds to the highest-weighted values in the logic tree structure, while the alternative represents a lower-weighted case. Each curve represents the aleatory uncertainty of the variables, while the different cases reflect epistemic uncertainty.

First, the approximation may not be an adequate proxy for the original probability distribution, failing to capture its statistical characteristics. An optimal discrete representation can be obtained by minimizing an appropriate loss function—e.g., the differences in mean, variance, higher-order moments, or the Kolmogorov-Smirnov D statistic—between the original and the approximated distributions. Most widely used approximation schemes are derived in this way (Keefer and Bodily, 1983; Miller and Rice, 1983). By contrast, an arbitrary discretization can perform poorly unless its fidelity has been thoroughly evaluated.

To illustrate this point, we compared the relative errors in mean and variance produced by three approximation schemes in Table 1 and 2 for several distributions commonly used to represent epistemic uncertainties in PSHA (normal, uniform, and exponential). The results are shown in Table 3. Even with the same number of discrete points, KB83 consistently provided a closer approximation than Pea24. The differences in mean using KB83 were negligible, and the variance errors did not exceed $\sim 10\%$. In contrast, Pea24 underestimated the mean of an exponential distribution by 19.2%, and its variance errors ranged from 48.7% to 72.3%. These results demonstrate that a poorly designed discretization can misrepresent the underlying distribution. Because Pea24 was not tested against multiple statistical moments, it offers no assurance of fidelity; in practice, KB83 or MR83 is therefore preferable.

Note, however, that the analysis above concerns the moments of the input variable itself, not the distribution of seismic hazard. In PSHA mean and fractiles, the relevant quantity is not the input variable's distribution, but the exceedance probability of ground-motion intensity, which leads to the second limitation of discrete approximations.

Secondly, even a discretization that well reproduces the distribution of an epistemic-uncertainty variable can fail if the seismic hazard at certain ground-motion intensities is sensitive to values lying outside the discretized range. Suppose $m_{\max}$ follows

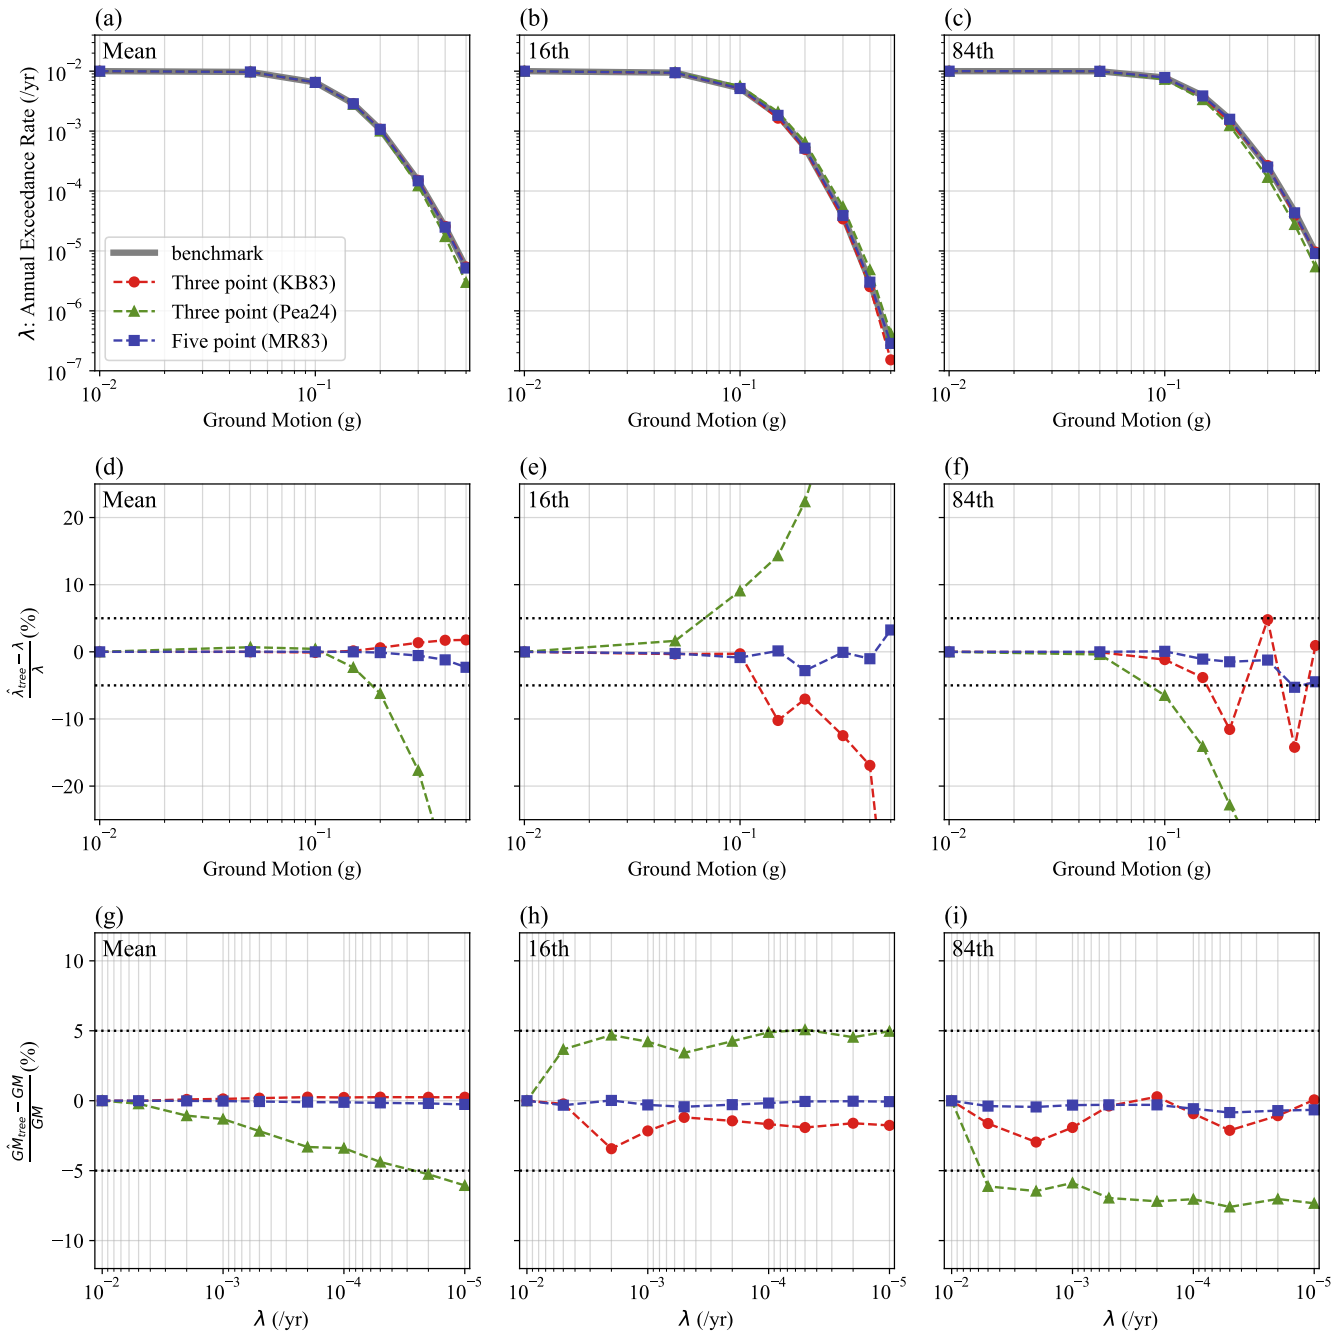

Figure 2. (a)–(c): Mean, 16th, and 84th fractile hazard curves from the numerical example. The benchmark results are shown as thick solid gray lines. Hazard curves based on the three-point approximation of Keefer and Bodily (1983) are shown as black dotted lines; those using the approximation from Petersen et al. (2024) are shown as red dotted lines; and the five-point approximation from Miller and Rice (1983) is presented as blue dotted lines. (d)–(f): Relative errors of the hazard curves obtained using different approximation strategies, evaluated against the benchmark at various target ground motion levels. (g)–(i): Relative errors of the ground motion intensity obtained using different approximation strategies, evaluated against the benchmark ground motion intensities at various hazard levels.

$N(6.5, 0.3^2)$. The KB83 discretization yields points at 6.0, 6.5, and 7.0, which may best represent the original distribution. Nevertheless, events with $m_{\max} = 7.5$ retain non-zero probability but are ignored in the approximation. If such large-magnitude events contribute significantly to the low-probability hazard, the discretized model will result in a biased mean hazard. Figure 2

TABLE 3
Errors in mean and variance of epistemic variables using different approximation strategies

	Normal		Uniform		Exponential	
Error (%)	mean	variance	mean	variance	mean	variance
Three points (KB83)	0.0	-10.5	0.0	-10.1	0.0	-3.73
Five points (MR83)	0.0	0.0	0.0	1.2	0.0	-9.9
Three points (Pea24)	0.0	-63.4	0.0	-48.7	-19.2	-72.3

TABLE 4
Comparison between widely-used and optimal three-point approximations for mean hazard calculation

Three points (KB83)		**Optimal weights for hazard =** 10^{-6}**/yr**				**Optimal weights for hazard =** 10^{-7}**/yr**			
both $\Delta\mu$ and $\Delta\sigma$		$\Delta\mu$		$\Delta\sigma$		$\Delta\mu$		$\Delta\sigma$	
p	**weight**	**p**	**weight**	**p**	**weight**	**p**	**weight**	**p**	**weight**
0.05	0.185	0.394	0.755	0.571	0.887	0.413	0.773	0.645	0.927
0.50	0.630	0.848	0.193	0.932	0.095	0.861	0.182	0.958	0.065
0.95	0.185	0.990	0.052	0.997	0.018	0.992	0.045	0.999	0.008

illustrates this behavior that the approximated logic-tree structure starts to deviate from the true mean hazard as target ground motion increases. This might be critical for facilities such as nuclear-waste repositories, where annual exceedance probabilities below 10^{-6}/year are of interest.

We demonstrate this with the configuration of Figure 1, considering only two epistemic variables—the median and standard-deviation models—while fixing $m_{\max} = 6.5$ and adopting an earthquake rate of one event per year. The errors in mean hazard introduced by the KB83 discretization are -4.2% and -9.7% at exceedance probabilities of 10^{-6} and 10^{-7} /year, respectively. An alternative weighting scheme, summarized in Table 4, reduces these errors to -0.26% and $+0.25\%$. Note that the weighting scheme differs not only depending on the target exceedance probability, but also on the epistemic uncertainty variables, since the sensitivity of each epistemic uncertainty variable to mean hazard is different.

Although this weighting is not unique, it shows that more weight must be assigned to the extreme values of the epistemic variables when estimating low-probability hazard, and that the weights shift further toward the extremes as the target probability decreases.

While the optimized weights improve the accuracy of the mean hazard, they depend on the specific integral being solved—i.e., on the target exceedance probability, the seismic source model, the ground-motion model, and the site location.

The next section presents a systematic approach, based on an importance-sampling framework, that can be applied to any PSHA problem to obtain accurate estimates of both the mean and the fractile hazard efficiently.

ADAPTIVE IMPORTANCE SAMPLING (AIS)

We introduce Adaptive Importance Sampling (AIS) and the concept of Population Monte Carlo (PMC) AIS, which will be employed in the computational framework presented in this paper. Prior to discussing these advanced methods, we provide an

overview of the conventional Monte Carlo and Importance Sampling, as they are prerequisites for the concepts utilized in this work.

Conventional Monte Carlo (MC) Simulation

Monte Carlo simulation is a general numerical method for estimating the expected value of a function. Consider a random variable X that follows a PDF $f_X(x)$. The expected value of a function $g(x)$, denoted by S, is defined as:

$$S = \int g(x) f_X(x) dx \tag{6}$$

The conventional Monte Carlo estimate of S is given by

$$\hat{S}_{MC} = \frac{1}{N} \sum_{i=1}^{N} g(X_i) \tag{7}$$

, where X_i is the i-th random sample drawn from the distribution $f_X(x)$, and N is the total number of samples. As N increases, $\hat{S}_{MC}$ converges to the true value S. The accuracy of the Monte Carlo estimate $\hat{S}_{MC}$ can be quantified by the variance of the estimator, $VAR[\hat{S}_{MC}]$:

$$VAR[\hat{S}_{MC}] = \frac{1}{N} \left[E_f[g^2(X)] - \hat{S}_{MC}^2 \right]$$

The accuracy of the estimator is often represented by the standard deviation of the estimate divided by the mean, the coefficient of variation (COV):

$$COV[\hat{S}_{MC}] = \frac{1}{\sqrt{N}} \sqrt{\frac{E_f[g^2(X)]}{\hat{S}_{MC}^2} - 1}$$

As the COV decreases, the estimate becomes more robust. For example, a COV of 5% indicates that we are 95% confident (Z-score of 1.96) that the true value of S lies within $\pm$ 9.8% ($= 1.96 \times 5\%$) of the estimator $\hat{S}_{MC}$. In contrast, a COV of 1% implies a narrower range of $\pm$ 1.96% ($= 1.96 \times 1\%$) with the same level of confidence.

Importance Sampling (IS)

Importance Sampling (IS) is a generalization of the conventional Monte Carlo method. Eq. (6) can also be rewritten as:

$$S = \int g(x) \frac{f_X(x)}{q_X(x)} q_X(x) dx \tag{8}$$

, where $q_X(x)$ is the proposal density (or distribution), an arbitrary PDF that is non-zero over the domain of integration. The IS estimate of S is given by:

$$\hat{S}_{IS} = \frac{1}{N}\sum_{i=1}^{N} g(X_i)\frac{f_X(X_i)}{q_X(X_i)} \tag{9}$$

, where the random sample X_i is drawn from the proposal distribution $q_X(x)$, and N denotes the total number of samples. $\hat{S}_{IS}$ is also an unbiased estimator (Theodoridis, 2015), and the ratio $f_X(X_i)/q_X(X_i)$ represents the importance weight, which adjusts for the change in the sampling density. The variance of the IS estimate is given by:

$$VAR[\hat{S}_{IS}] = \frac{1}{N}\left[E_q\left[\left(\frac{g(X)f_X(X)}{q_X(X)}\right)^2\right] - \hat{S}_{IS}^2\right] \tag{10}$$

The accuracy metric of $\hat{S}_{IS}$, COV, can also be expressed as:

$$COV[\hat{S}_{IS}] = \frac{1}{\sqrt{N}}\sqrt{\frac{E_q\left[\left(\frac{g(X)f_X(X)}{q_X(X)}\right)^2\right]}{\hat{S}_{IS}^2} - 1}$$

By properly selecting $q_X(x)$, $COV[\hat{S}_{IS}]$ can be reduced. Specifically, we can choose $q_X(x)$ to be

$$q_X^*(x) = \frac{g(x)f_X(x)}{\hat{S}_{IS}} \tag{11}$$

This distribution $q_X^*(x)$ is referred to as the IS optimal density since the variance and COV of the Importance Sampling estimate becomes zero (see Eq. (10)). This indicates that the S can, in theory, be estimated with a single random sample if we know $q_X^*(x)$ at any x. However, identifying the optimal density is nontrivial, as $\hat{S}_{IS}$ is unknown and closed-form expressions for $g(x)$ and $f_X(x)$ are not always available. To address this, we employ an algorithm that iteratively explores the region of importance, a method known as Adaptive Importance Sampling (AIS). Specifically, we introduce Population Monte Carlo (PMC) AIS to apply it to PSHA.

Population Monte Carlo Adaptive Importance Sampling

Numerous AIS algorithms have been developed to iteratively explore regions of importance in numerical integration (Lepage, 1978; Friedman and Wright, 1981; Rubinstein, 1997; Press and Farrar, 1990; Rubinstein and Kroese, 2004; Bugallo et al., 2017). Among them, Population Monte Carlo (PMC) is a straightforward AIS algorithm that utilizes importance weights to update the proposal distribution (Cappé et al., 2004). Similar to other iterative algorithms, each iteration of the PMC AIS consists of three fundamental steps aimed at approximating the optimal IS density: 1) Random sampling from proposal density, 2) estimating the expectation using the generated samples, and (3) updating the proposal density.

First, in the sampling step, N samples ($\mathbf{X}^{(t)} = \{X_i\}_{i=1}^N$) are randomly drawn from the proposal density at t-th iteration step, denoted as $q_X^{(t)}(x)$. The initial proposal density, $q_X^{(0)}(x)$, is typically chosen as either the original distribution $f_X(x)$ or a uniform distribution. If prior knowledge about the region of importance is available, it can be used to construct a more informative initial

proposal. For example, in seismic hazard analysis, if it is known that certain ranges of earthquake magnitude and distance contribute significantly to the hazard, the initial proposal distribution can be tailored to sample more densely within those ranges.

Secondly, in the estimation step, the integrands $g(x)f_X(x)$ in Eq. (6) are evaluated at the sampled points, and the likelihood ratios (w_i) between the target density $g(x)f_X(x)$ and the current proposal density $q^{(t)}X(x)$ are computed. These likelihood ratios serve both to estimate the expectation of $g(x)$ and to guide the update of the proposal density for the next iteration:

$$w_i^{(t)} = \frac{g\left(X_i^{(t)}\right) f_X\left(X_i^{(t)}\right)}{q_X^{(t)}\left(X_i^{(t)}\right)} \propto \frac{q_X^*\left(X_i^{(t)}\right)}{q_X^{(t)}\left(X_i^{(t)}\right)} \tag{12}$$

The estimated value of the integral at the t-th iteration, $\hat{S}^{(t)}$, is obtained by taking the average of w_i:

$$\hat{S}^{(t)} = \frac{1}{N}\sum_{i=1}^{N} g\left(X_i^{(t)}\right) \frac{f_X\left(X_i^{(t)}\right)}{q_X^{(t)}\left(X_i^{(t)}\right)} = \frac{1}{N}\sum_{i=1}^{N} w_i^{(t)} \tag{13}$$

Note that $\hat{S}^{(t)}$ is unbiased as it remains an importance sampling estimator (Eq. (9)).

Third, in the updating step, the weights $w_i^{(t)}$ are normalized such that $\sum w_i^{(t)} = 1$. The normalized $w_i^{(t)}$ represents a probability mass function proportional to the optimal density. Then, we create a new sample set $\mathbf{Y}^{(t)}$ by resampling from the original sample set $\mathbf{X}^{(t)}$, using the weights $\mathbf{w}^{(t)}$ as the resampling probabilities. This resampling process preferentially selects samples where $q_X^*(x)$ is greater than $q_X^{(t)}(x)$ to result in a set of samples that more closely represents the optimal density, $q_X^*(x)$. If the proposal density $q_X^{(t)}(x)$ exactly matches the optimal density $q_X^*(x)$, all weights $w_i^{(t)}$ become equal, and both $\mathbf{X}^{(t)}$ and $\mathbf{Y}^{(t)}$ are effectively drawn from the same distribution, $q_X^*(x)$. If the two densities differ, the proposal distribution is updated by fitting the resampled points $\mathbf{Y}^{(t)}$ to a presumed distribution, denoted by $p_X(x;\eta)$, where η is the parameters for the presumed distribution (e.g., mean and standard deviation in normal distribution), thereby obtaining an updated proposal density $q_X^{(t+1)}(x)$.

Ideally, this iterative process continues until all weights $w_i^{(t)}$ converge to uniformity, i.e., $q_X^{(t)}(x) = q_X^*(x)$. However, in practice, such convergence is rarely achieved because the assumed distribution $p_X(x;\eta)$ may not share the same functional form as the true optimal density. Instead, the iteration is typically terminated once the distributions of normalized weights $w_i^{(t+1)}$ and $w_i^{(t)}$ become sufficiently similar, indicating that further updates to the proposal density are negligible. The Kolmogorov–Smirnov (K-S) D statistics can be used (Kolmogorov, 1933).

After the termination of the loop, we estimate the integral $\hat{S}^{(t)}$ and the interated proposal density $q_X^{(t)}(x)$. In the context of PSHA, $\hat{S}^{(t)}$ represents the mean hazard, and $q_X^{(t)}(x)$ will be key to the computation of fractiles and sensitivity analysis, which will be explained in the following section.

PROPOSED AIS FRAMEWORK FOR PSHA EPISTEMIC UNCERTAINTY PROPAGATION

Before introducing the PSHA computational framework based on PMC-AIS, we first present the mathematical formulation of mean and fractile hazards, and sensitivity analysis in order to establish their connection with the importance sampling framework.

Remind that the individual hazard integrates the aleatory uncertainty variables, and mean and fractile hazards are obtained through the multiple runs of individual hazard calculation across the possible alternative seismic source and ground motion models, the epistemic uncertainty of the model.

Mean hazard

By combining equations Eq. (2) and Eq. (3), the mean hazard can be expressed as:

$$\begin{aligned}\lambda(a) &= \iint_{\phi,\theta} \nu(\theta) I\left(u(\phi;\theta) > a\right) f_{\Phi|\Theta}(\phi|\theta) f_{\Theta}(\theta) d\phi\, d\theta \\ &= \iint_{\phi,\theta} \nu(\theta) I\left(u(\phi;\theta) > a\right) f_{\Phi,\Theta}(\phi,\theta) d\phi\, d\theta\end{aligned}$$

This formulation implies that in order to compute the mean hazard, it is not necessary to separate aleatory and epistemic uncertainty variables. This offers a substantial computational advantage, as it avoids the need for nested integrations—where the inner integral corresponds to individual hazard computation, and the outer integral accounts for the distribution over epistemic uncertainty. The IS formulation of mean hazard is:

$$\lambda(a) = \iint_{\phi,\theta} \nu(\theta) I\left(u(\phi;\theta) > a\right) \frac{f_{\Phi,\Theta}(\phi,\theta)}{q_{\Phi,\Theta}(\phi,\theta)} q_{\Phi,\Theta}(\phi,\theta) d\phi\, d\theta$$

The IS MC estimate of the mean hazard is:

$$\hat{\lambda}(a) = \sum_{i=1}^{N} \nu(\Theta_i) I\left(u(\Phi_i;\Theta_i) > a\right) \frac{f_{\Phi,\Theta}(\Phi_i,\Theta_i)}{q_{\Phi,\Theta}(\Phi_i,\Theta_i)} \tag{14}$$

, where N is the total number of samples, (Φ_i, Θ_i) are sampled from $q_{\Phi,\Theta}(\phi,\theta)$. The optimal density which makes the variance of $\hat{\lambda}(a)$ zero is given by:

$$q^*_{\Phi,\Theta}(\phi,\theta) = \frac{\nu(\theta) I\left(u(\phi;\theta) > a\right) f_{\Phi,\Theta}(\phi,\theta)}{\lambda(a)} \tag{15}$$

.

Fractile Hazard and Sensitivity Analysis

By rearranging Eq. (15), we obtain the following expression (see Appendix A for the derivation):

$$\lambda(a|\theta) = \lambda(a) \frac{q^*_{\Theta}(\theta)}{f_{\Theta}(\theta)} \tag{16}$$

, where $\lambda(a|\theta)$ denotes the hazard (i.e., exceedance rate) conditional on the epistemic uncertainty parameter θ—that is, the individual hazard given the seismic source and ground motion models. Here, $\lambda(a)$ represents the mean hazard, and $q^*_\Theta(\theta)$ is the importance sampling (IS) optimal density of the epistemic uncertainty variables, computed by marginalizing the IS optimal density (Eq. (15)) over the aleatory uncertainty variables ϕ: $\int_\phi q^*_{\Phi,\Theta}(\phi,\theta)d\phi$. Here, $f_\Theta(\theta)$ denotes the original PDF of the epistemic uncertainty variables. By repeatedly calculating Eq. (16) using random Θ samples following $f_\Theta(\theta)$, a list of individual hazards can be obtained.

Recall that the CDF of the hazard, from which fractile hazards can be derived, must be computed empirically by sorting the individual hazards, $\lambda(a|\theta)$, computed from numerical integration of Eq. (1) (see Eq. (4) and Eq. (5)). The Eq. (16) eliminates the need for numerical integration typically required for individual hazard calculation. That is, individual hazard can be efficiently calculated if the mean hazard, $\lambda(a)$, and the optimal density of the epistemic uncertainty variables, $q^*_\Theta(\theta)$, are known.

Eq. (16) also yields an important insight regarding the optimal density and the sensitivity of each epistemic uncertainty variable to the hazard. Suppose we are interested in assessing the sensitivity of the i-th epistemic variable, θ_i, to the hazard. If variations in θ_i have a negligible effect on the hazard, then the individual hazard remains approximately constant across different values of θ_i, i.e., $\lambda(a|\theta_i) \approx \lambda(a)$ for all θ_i. In this case, the marginalized optimal density $q^*_{\Theta_i}(\theta_i)$ should closely resemble the original distribution $f_{\Theta_i}(\theta_i)$. Conversely, if θ_i exerts a substantial influence on the hazard, then $q^*_{\Theta_i}(\theta_i)$ will significantly deviate from $f_{\Theta_i}(\theta_i)$. Therefore, the degree of deviation between the marginalized optimal and original densities can serve as an indicator of the sensitivity of the hazard to the corresponding epistemic uncertainty variable. The hazards with respect to the various values of ith epistemic uncertainty variable can be calculated by reformulating Eq. (16) in terms of the i-th variable θ_i:

$$\lambda_i(a|\theta_i) = \lambda(a) \times \frac{q^*_{\Theta_i}(\theta_i)}{f_{\Theta_i}(\theta_i)} \tag{17}$$

, where $q^*_\Theta(\theta_i) = \int_\phi \int_{\tilde{\theta}_i} q^*_{\Phi,\Theta}(\phi,\theta)d\phi d\tilde{\theta}_i$ and $\tilde{\theta}_i$ denotes the set of epistemic uncertainty variables excluding θ_i.

The variance contribution of the i-th epistemic uncertainty variable to the total hazard variance, denoted C_i, can be computed as:

$$C_i = \frac{\mathrm{VAR}_{\theta_i}[\lambda_i(a|\theta_i)]}{\mathrm{VAR}_\theta[\lambda(a|\theta)]}. \tag{18}$$

, where $\mathrm{VAR}_{\theta_i}[\lambda_i(a|\theta_i)]$ is the variance of the hazard with respect to i-th epistemic uncertainty variable, θ_i, and $\mathrm{VAR}_\theta[\lambda(a|\theta)]$ is the variance of the hazard when all the epistemic uncertainty variables (θ) are varied. Although the computation of $\lambda_i(a|\theta_i)$ also inherently involves numerical integration, the use of Eq. (17) significantly accelerates this process. Interestingly, C_i is equivalent to the first-order Sobol index of θ_i (Sobol', 2001; Saltelli et al., 2004), under the assumption that θ_i is independent of the set of epistemic uncertainty variables excluding θ ($\tilde{\theta}_i$),—i.e., θ_i is propagated through all other variables (see Appendix B for the derivation). Note that this is not always the case when variables are conditionally propagated; for example, different epistemic

uncertainty distributions of the ground motion model may be applied to different seismic source models. The first-order Sobol index, V_i quantifies the direct contribution of each variable to the total variance of the estimate and ranges between 0 and 1. A larger value of V_i indicates that the hazard is more sensitive to θ_i. The sum of C_i, denoted $\sum_i^{d_\theta} C_i$, where d_θ is the number of epistemic uncertainty variables, does not necessarily equal 1. The residual term $1 - \sum_i^{d_\theta} C_i$ represents the interaction effects among the variables.

While Eq. (16)-Eq. (18) are useful in fractile hazard calculation and sensitivity analysis; however, there are three primary challenges in utilizing the equations. First, the mean hazard $\lambda(a)$ must be known. In the traditional PSHA framework, individual hazards are first computed and then averaged to obtain the mean hazard. In contrast, equation Eq. (16) requires the mean hazard to compute the individual hazard, which may seem paradoxical. Second, $q^*_\Theta(\theta)$ (or $q^*_{\Theta_i}(\theta_i)$) can be derived through marginalization of $q^*_{\Phi,\Theta}(\phi,\theta)$, and the accuracy of obtained fractile is directly related to how we can closely approximate the optimal density. However, obtaining $q^*_{\Phi,\Theta}(\phi,\theta)$ is non-trivial problem, requiring well-designed algorithm to approximate this. Third, even if the optimal density, $q^*_{\Phi,\Theta}(\phi,\theta)$, becomes available, obtaining its marginal distribution, $q^*_\Theta(\theta)$ (or $q^*_{\Theta_i}(\theta_i)$), still requires numerical integration, which is computationally expensive in general. In the following section, we propose a framework to address these challenges using Gaussian Population Monte Carlo Adaptive Importance Sampling (G-PMC AIS).

G-PMC AIS PSHA Computational framework

PMC AIS is a general framework for numerically computing the mean and identifying a proposal IS density that closely approximates the optimal density. For a target ground motion, a, applying PMC AIS to PSHA iteratively yields both the mean hazard, $\lambda(a)$, and the iterated proposal IS density, $\hat{q}^*_{\Phi,\Theta}(\phi,\theta)$, which approximates the optimal density $q^*_{\Phi,\Theta}(\phi,\theta)$. This resolves the first two challenges associated with applying Eq. (15) to obtain individual hazard estimates, as noted in the previous paragraph. Furthermore, by employing a joint Gaussian distribution as the assumed distribution in PMC AIS (i.e., $p_X(x;\eta) = \pi_X(x;\mu,\Sigma)$, where μ and Σ are the mean vector and covariance matrix, respectively), the third challenge—computational difficulty in marginalization—is addressed. For a joint Gaussian, marginalization is straightforward: we can simply select the relevant entries from the mean vector and covariance matrix corresponding to the variables of interest, without requiring numerical integration. Additionally, the use of Gaussian distributions provides further advantages: they are easy to sample from, their PDFs are straightforward to evaluate for weight calculations, and they effectively capture dependencies among variables.

We refer to this framework as Gaussian Population Monte Carlo Adaptive Importance Sampling PSHA (G-PMC AIS PSHA). The framework comprises three primary stages: 1) Initialization, 2) Estimation of the mean hazard and the iterated proposal IS density, and 3) Fractile calculation (Algorithm 1, Fig. 3). The following sections describe these steps in detail.

Stage 1. Initialization In Line 1 of Algorithm 1, the initialization stage begins by setting the iteration step $t = 0$ and initializing the Kolmogorov–Smirnov (K-S) D statistic $d = 1$, the maximum possible value, to ensure at least one iteration. The stopping criterion is defined by $\varepsilon = 0.1$, which will be explained later.

Next, the joint probability density of ϕ (aleatory variables such as magnitude, distance, and ground motion residual) and θ (epistemic variables such as the b-value, $m_{\max}$, and alternative ground motion models) is defined (Line 2). We define it as the conditional density of ϕ given θ as $f_{\Phi|\theta}(\phi|\theta)$. This process is equivalent to constructing the logic-tree structure and defining each end branch's PDF of magnitude, distance, and ground motion distributions in conventional PSHA studies. Note that, unlike traditional approaches, the distribution of epistemic uncertainty variables in this framework can be continuous and does not require discretization. A simple two-dimensional example involving one aleatory variable (ϕ: magnitude M) and one epistemic variable (θ: $\Delta\mu$, median ground motion model) is depicted in the "Stage 1. Initialization" box in Fig. 3.

Finally, the initial proposal density $q^{(0)}_{\Phi,\Theta}(\phi,\theta)$ is selected (Line 3). Assuming a joint Gaussian form, the mean vector and covariance matrix must be specified. If prior knowledge exists regarding important regions of each variable for the target ground motion a, the joint Gaussian can be centered on those values. Otherwise, a common strategy is to center the mean at the midpoint of the integration domain and use a covariance matrix with large variances and zero covariances to span the full integration range. These parameters will be updated in subsequent iterations to reflect the important regions and interdependencies.

Stage 2. Estimation of mean hazard and iterated proposal IS density In this stage, we estimate both the mean hazard $\hat{\lambda}(a)$ and the iterated proposal IS density $\hat{q}^*_{\Phi,\Theta}(\phi,\theta)$, which approximates the optimal density. This is an iterative process, and we consider the t-th iteration step.

A random sample set $\vec{X}^{(t)}$ consisting of ϕ and θ is drawn from $q^{(t)}_{\Phi,\Theta}(\phi,\theta)$. The indicator function $I(u(\vec{X}^{(t)})>a)$ is evaluated for each sample, and the results are multiplied by the earthquake occurrence rate $\nu(\theta)$ and the likelihood ratio (importance weight) $f_{\Phi,\Theta}(\vec{X}^{(t)})/q_{\Phi,\Theta}(\vec{X}^{(t)})$ to obtain the weight vector $\vec{w}^{(t)}$ (Lines 5–6). The mean hazard for t-th iteration, $\hat{\lambda}^{(t)}(a)$, is then estimated as the average of $\vec{w}^{(t)}$ (Eq. (14), Line 7).

Next, we update the proposal density to better approximate the optimal density. A new sample set $\vec{Y}^{(t)}$ is formed by resampling from $\vec{X}^{(t)}$ according to weights $\vec{w}^{(t)}$ (Line 8). Note that $\vec{w}^{(t)}$ is proportional to the optimal density, $\nu(\theta)I(u(\vec{X}^{(t)})>a)f_{\Phi,\Theta}(\vec{X}^{(t)})$, divided by the proposal density in the current iteration step, $q^{(t)}_{\Phi,\Theta}(\phi,\theta)$. Thus, the resampling process redistributes the original samples ($\vec{X}^{(t)}$) to follow the optimal density. The updated proposal density $q^{(t+1)}_{\Phi,\Theta}(\phi,\theta)$ is then obtained by fitting a joint Gaussian distribution to $\vec{Y}^{(t)}$ via maximum likelihood estimation (Lines 9–10).

To decide whether to terminate or continue the iteration, the K-S D statistic is evaluated (Line 11). This statistic measures the maximum difference between the CDFs of $q^{(t+1)}\Phi,\Theta$ and $q^{(t)}\Phi,\Theta$: $\sup_{\Phi,\Theta}|Q^{(t+1)}_{\Phi,\Theta}-Q^{(t)}_{\Phi,\Theta}|$. A smaller d indicates convergence, and the iteration terminates when $d<\varepsilon$, where $\varepsilon=0.1$ in this study. The outputs of this stage are the mean hazard estimate $\hat{\lambda}(a)$ and the approximated optimal (iterated proposal) density $\hat{q}^*_{\Phi,\Theta}(\phi,\theta)$ (Line 13). An example is shown in the "Stage 2" box of Fig. 3, where the evolved proposal density captures the integration's important region near $\phi\approx5.7$ and $\theta\approx0.2$.

Stage 3. Fractile calculation In this stage, $\hat{q}_{\Phi,\Theta}(\phi,\theta)$ is marginalized to obtain $\hat{q}^*_{\Theta}(\theta)$ (Line 14). This is done by extracting the θ-related elements of the mean vector and covariance matrix without additional computation. Random samples of θ are drawn

TABLE 5
Parameters and Functions used in Algorithm 1

Parameter/Function	Description
t	Iteration step
N	Number of samples
d	K-S D statistics
ε	Pre-defined iteration stopping criterion (e.g., 0.1)
a	PSHA target ground motion (e.g., 0.5g)
$\vec{p}$	Target fractiles (e.g., 16th, 50th, 84th)
ϕ	Aleatory variability variables (e.g., m, r)
θ	Epistemic uncertainty variables (e.g., b-value, $\Delta\mu$)
$L(X;\cdot)$	Likelihood function given observed data X
$f_{\Phi,\Theta}(\phi,\theta)$	Original joint PDF of ϕ and θ
$f_{\Theta}(\theta)$	Original PDF of θ
$u(\phi,\theta)$	ground motion intensity given ϕ and θ
$I(u(\phi,\theta)>a)$	Indicator function (1 if $u(\cdot)>a$, otherwise 0).
$\pi_{\Phi,\Theta}(\phi,\theta;\mu,\Sigma)$	PDF of joint Gaussian distribution (μ=mean; Σ=covariance)
$q^{(t)}_{\Phi,\Theta}(\phi,\theta)$	t-th iteration step proposal density
$Q^{t}_{\Phi,\Theta}(\phi,\theta)$	CDF of $q^{(t)}_{\Phi,\Theta}(\phi,\theta)$
$\hat{q}^{*}_{\Phi,\Theta}(\phi,\theta)$	Iterated proposal PDF (Approximated optimal density)
$\hat{q}^{*}_{\Theta}(\theta)$	Marginal distribution of $\hat{q}^{*}_{\Phi,\Theta}(\phi,\theta)$ over θ
$\hat{\lambda}(a)$	Estimated mean hazard
$\vec{\lambda}(a\|\theta)$	List of estimated individual hazard
$\hat{\lambda}_{\vec{p}}(a)$	Estimated $\vec{p}$th fractile hazards

from the original epistemic uncertainty distribution $f\Theta(\theta)$ to form the sample set $\vec{\xi}$ (Line 15). These are visualized in the leftmost plot of the "Stage 3" box in Fig. 3, where the blue solid curve represents $f_{\Theta}(\theta)$ and red dots indicate the samples representing key fractiles (1st, 16th, 50th, 84th, and 99th percentiles).

Next, likelihood ratios between $\hat{q}^{*}_{\Theta}(\theta)$ and $f_{\Theta}(\theta)$ are computed (second plot in Fig. 3). These ratios are scaled using the mean hazard $\hat{\lambda}(a)$, producing individual hazard estimates $\vec{\lambda}(a|\theta)$ (Eq. (16), Line 16). The pth fractile hazard, $\hat{\lambda}_{p}(a)$, is then identified such that $p\%$ of the individual hazards are less than or equal to $\lambda_{p}(a)$ (Eq. (5), Line 17). This process is illustrated in the final plots of the "Stage 3" box in Fig. 3.

This overall procedure (Stages 1 to 3) is repeated for each target ground motion intensity of interest.

It is important to note that the overall computational effort is dominated by “Stage 2” due to repeated evaluations of the marginal hazard for numerical integration. In contrast, “Stage 3”, fractile calculation, is computationally lightweight, involving only likelihood ratio computations between known distributions.

Algorithm 1 G-PMC AIS PSHA computational framework pseudocode given target ground motion, a

[1. Initialize]

$t = 0$, $d = 1$, $\varepsilon = 0.1$

Define $f_{\Phi,\Theta}(\phi, \theta)$ (or $f_{\Phi|\Theta}(\phi|\theta) f_{\Theta}(\theta)$)

Initial guess: $q^{(t)}_{\Phi,\Theta}(\phi, \theta) = \pi_{\Phi,\Theta}(\phi, \theta; \mu^{(t)}, \Sigma^{(t)})$

[2. Mean and Iterated Proposal]

while $d > \varepsilon$ **:**

$\vec{X}^{(t)} = \{X_i = (\Phi_i, \Theta_i)\}_{i=1}^{N}$, $\vec{X} \sim q^{(t)}_{\Phi,\Theta}(\phi, \theta)$

$\vec{w}^{(t)} = \frac{\nu(\theta) I(u(\vec{X}^{(t)}) > a) f_{\Phi,\Theta}(\vec{X}^{(t)})}{q^{(t)}_{\Phi,\Theta}(\vec{X}^{(t)})}$

$\hat{\lambda}^{(t)}(a) = \sum \vec{w}^{(t)} / N$

$\vec{Y}^{(t)} \sim \vec{X}^{(t)}$ with probability of selecting $X_i^{(t)}$ is $w_i^{(t)}$

$\mu^{(t+1)}, \Sigma^{(t+1)} = \arg\max_{\mu,\Sigma} L(\pi_X(\vec{Y}^{(t)}; \mu, \Sigma))$

$q^{(t+1)}_{\Phi,\Theta}(\phi, \theta; \mu^{(t+1)}, \Sigma^{(t+1)}) = \pi^{(t+1)}_{\Phi,\Theta}(\phi, \theta; \mu^{(t+1)}, \Sigma^{(t+1)})$

$d = \sup_{\Phi,\Theta} |Q^{(t+1)}_{\phi,\Theta} - Q^{(t)}_{\Phi,\Theta}|$

$t \leftarrow t + 1$

$\hat{\lambda}(a) = \hat{\lambda}^{(t)}(a)$, $\hat{q}^{*}_{\Phi,\Theta}(\phi, \theta) = q^{(t)}_{\Phi,\Theta}(\phi, \theta)$

[3. Fractile]

$\hat{q}^{*}_{\Theta}(\theta) = \pi^{(t)}_{\Theta}(\theta; \mu^{(t)}_{\theta}, \Sigma^{(t)}_{\theta\theta})$

$\vec{\xi} = \{\Theta_i\}_{i=1}^{n}$, $\vec{\xi} \sim f_{\Theta}(\theta)$

$\vec{\lambda}(a|\theta) = \hat{\lambda}(a) \times \hat{q}^{*}_{\Theta}(\vec{\xi}) / f_{\Theta}(\vec{\xi})$

$\hat{\lambda}_{\vec{p}}(a) = \inf \left\{ x : F_{\vec{\lambda}(a|\theta)}(x) \geq \frac{\vec{p}}{100} \right\}$

return $\hat{\lambda}(a), \hat{\lambda}_{\vec{p}}(a)$

NUMERICAL EXAMPLE

Numerical experiment settings

We consider two different seismic source configurations: areal and fault sources (Figure 4). These are commonly used seismic source types in practical PSHA applications (Coppersmith et al., 2012; Petersen et al., 2024). The earthquake occurrences are modeled using a truncated exponential distribution, with a minimum magnitude $m_{\min}$ set to 5.0. The annual occurrence rate for earthquakes with magnitudes exceeding $m_{\min}$, denoted as ν, is set to 0.01. The ground motion model of Abrahamson et al. (2014) is employed, assuming a V_{S30} of 760 m/s.

We incorporate four epistemic uncertainty variables: two related to the seismic source—the b-value and $m_{\max}$—and two associated with the ground motion model—its median and standard deviation. The b-value and $m_{\max}$ are assumed to follow normal distributions: $b \sim \mathcal{N}(1.0, 0.1)$ and $m_{\max} \sim \mathcal{N}(7.0, 0.3)$, with truncation to the ranges [0.7, 1.1] and [5.9, 7.1], respectively, to prevent unrealistic values. Uncertainty in the natural logarithm of the median ground motion is represented as $\Delta\mu \sim \mathcal{N}(0, 0.1)$, corresponding to approximately a 10% variation in peak ground acceleration (PGA) within one standard deviation. Similarly, the uncertainty in the standard deviation of the natural logarithmic ground motion is modeled as $\Delta\sigma \sim \mathcal{N}(0, 0.05)$.

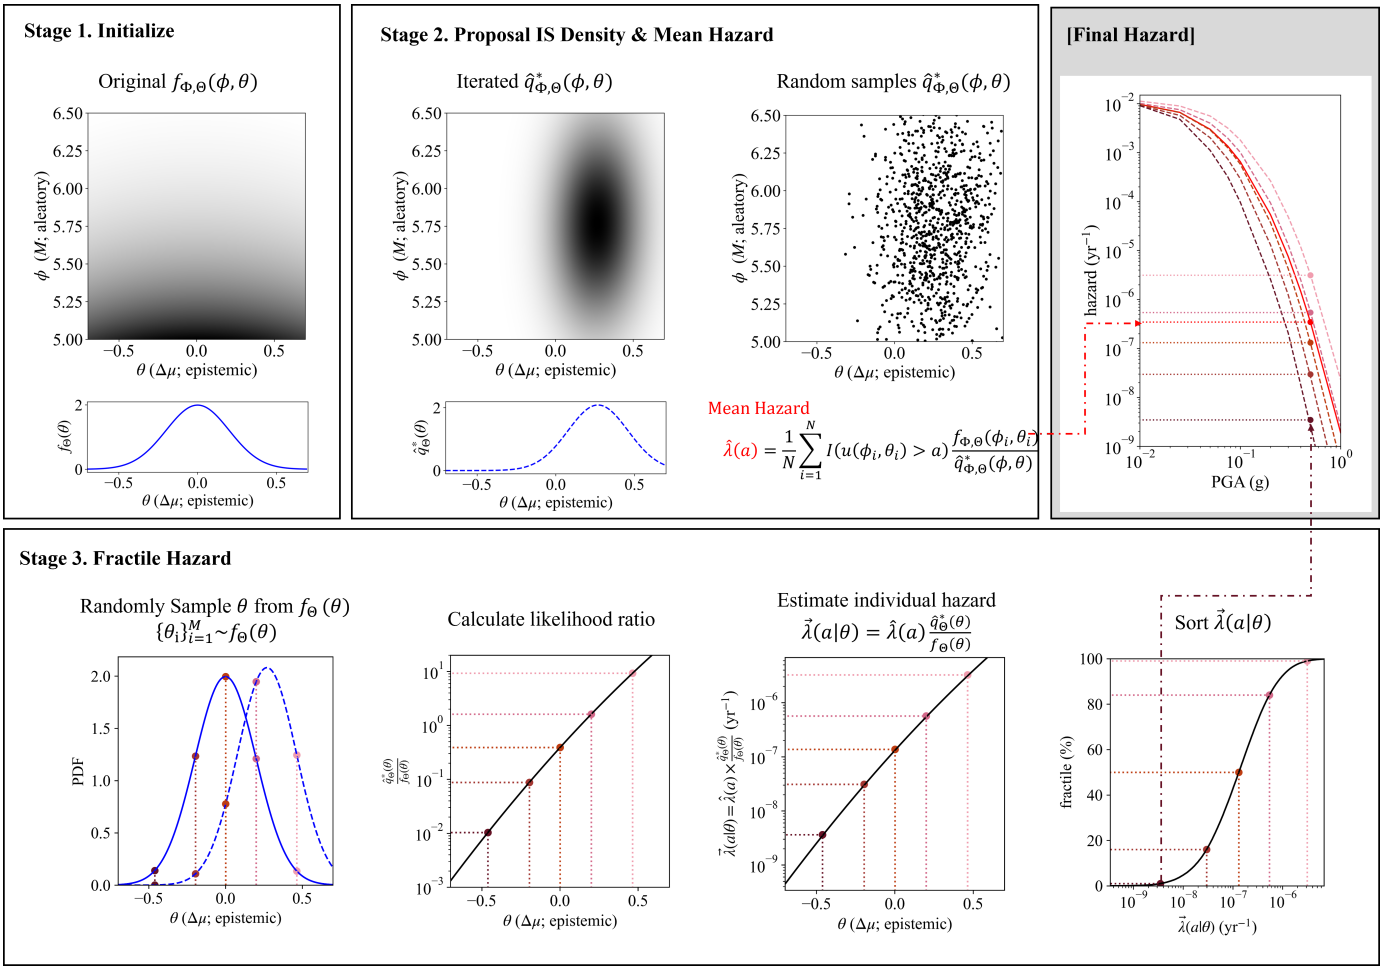

Figure 3. Diagram illustrating the proposed framework using a simplified 2D PSHA problem with one aleatory variable (ϕ: magnitude, m) and one epistemic variable (θ: median ground motion, $\Delta\mu$), for a target ground motion level of 0.5g. The magnitude follows a truncated exponential distribution (M5.0–6.5), and epistemic uncertainty in ln [median ground motion] relative to the base model (Sadigh et al., 1997) is modeled as $\mathcal{N}(0, 0.2)$. In "Stage 1. Initialize", the joint distribution $f_{\Phi,\Theta}(\phi,\theta)$ is shown with darker regions indicating higher density. The marginal epistemic distribution $f_\Theta(\theta)$ is shown as a blue solid curve. "Stage 2. Proposal IS Density and Mean Hazard" presents the iterated proposal density $\hat{q}^*_{\Phi,\Theta}(\phi,\theta)$ and marginalized form $\hat{q}^*_\Theta(\theta)$ (blue-dotted curve), used to compute the estimated mean hazard $\hat{\lambda}(a)$ (red dot in "Final Hazard"). In "Stage 3. Fractile Hazard", fractile sampling from $f_\Theta(\theta)$ is illustrated, with selected samples marked as reddish dots. The blue solid and blue dotted curves represent the original and the iterated proposal densities of epistemic uncertainty, respectively. The likelihood ratio $\hat{q}^*_\Theta(\theta)/f_\Theta(\theta)$ and individual hazard $\vec{\lambda}(a|\theta)$ are computed, and the full hazard distribution is constructed by sorting individual hazards to identify fractile values, also shown in the "Final Hazard" box.

We estimate hazard levels at 0.13g, 0.32g, 0.64g, and 1.1g, which correspond to annual exceedance probabilities of approximately 10^{-3}, 10^{-4}, 10^{-5}, and 10^{-6} for the fault source, and $\sim 1.3 \times 10^{-4}$, 1.2×10^{-5}, 1.2×10^{-6}, and 1.3×10^{-7} for the areal source.

Numerical solutions

We compare the proposed framework, "G-PMC AIS," with logic-tree approaches using widely adopted discretization schemes: three- and five-point approximations (Keefer and Bodily, 1983; Miller and Rice, 1983), denoted as "LT(3)" and "LT(5)", respectively (Table 1). Given four epistemic uncertainty variables, the total number of end-branches (N_θ) is 81 and 625 for the "LT(3)"

and "LT(5)", respectively. The total number of marginal hazard evaluations is $N_t = N_\theta \times N_\phi$, where N_ϕ is the number of samples used per individual hazard computation. For both methods, individual hazard estimation is performed using the framework proposed by Houng and Ceferino (2025), which is significantly faster than traditional techniques. We use $N_\phi = 10^4$, which ensures high accuracy of estimates (COV $\leq 2.5\%$) for most cases considered.

Additionally, we evaluate the performance of a continuous distribution based conventional Monte Carlo ("MC") method. This approach is similar in structure to the logic-tree method but draws samples directly from the continuous distributions without approximation. For consistency, the same $N_\phi = 10^4$ is used. It guarantees the true fractile if sufficient number of individual hazards are estimated. The benchmark CDF is derived from 1,000 MC realizations.

To evaluate the accuracy of mean hazard estimates, we use COV, representing the relative variability of the estimate with respect to mean. A lower COV indicates a more robust estimate. For fractile hazard accuracy, we use the Kolmogorov–Smirnov (K-S) D statistic to compare the CDF of each method's fractile hazard with a benchmark CDF derived from 1,000 MC realizations. Note that a K-S D statistic below 5% generally indicates good agreement. For instance, if hazard is assumed log-normally distributed (Kwong and Jaiswal, 2023), a 5% K-S D corresponds to a mean difference of only 1%–2.3% depending on the assumed standard deviation.

Computational efficiency is measured by the total number of marginal hazard computations, N_t.

Results and Discussion

Figure 5 shows the coefficient of variation (COV) of mean hazard estimates as a function of N_t, and Figure 6 presents the Kolmogorov–Smirnov (K-S) D statistics for the fractile hazards. Across all scenarios—both areal and fault sources and for all ground motion levels—the G-PMC AIS framework demonstrates superior efficiency.

To achieve a 1% COV for PGAs of 0.13g, 0.32g, 0.64g, and 1.1g, G-PMC AIS requires 29, 29, 21, and 16 times fewer samples than the "LT(3)", and 224, 224, 162, and 126 times fewer than "LT(5)", respectively (Figure 5). Compared to "MC", the sample reductions are even more pronounced: 286, 559, 1666, and 3,775 times fewer samples are required by G-PMC AIS, respectively. This underscores the inefficiency of MC method at higher ground motion levels, where rare-event modeling becomes computationally burdensome, unlike G-PMC AIS, which maintains stable performance.

Although G-PMC AIS excels in mean hazard estimation, high accuracy in the mean does not inherently guarantee accurate fractile estimates. To evaluate this, we compute K-S D statistics at N_t's corresponding to 1% COV in mean hazard estimation, achieved using G-PMC AIS. The resulting K-S Ds for 0.13g, 0.32g, 0.64g, and 1.1g are 5.6%, 4.0%, 2.8%, and 3.0%, respectively, indicating good agreement with the benchmark CDFs (Figure 6).

While the five-point logic tree, LT(5), yields more accurate fractile estimates than the three-point, LT(3), as seen in Figure 2, its computational cost increases significantly with the number of epistemic variables. Thus, although preferred for accuracy, it suffers from a computational burden due to the rapid growth in N_t.

An interesting contrast is observed in the performance of the MC method between mean and fractile estimates (Fig. 5 and 6). For the mean, performance deteriorates with increasing ground motion level due to the rarity of higher threshold exceeding events. However, in the case of fractiles, the K-S D statistic is relatively insensitive to ground motion, since all samples contribute equally to the CDF regardless of event rarity.

It is important to note that the efficiency of LT(3), LT(5), and MC methods depends heavily on the number of samples used for individual hazard calculations. While we use $N_\phi = 10^4$, reducing this to 10^3—for example, corresponding to approximately 20 magnitude bins (e.g., increments of 0.1 from 5 to 7) and 50 spatial samples—still results in G-PMC AIS outperforming all other methods. Moreover, optimizing the grid size for such methods demands significant prior knowledge or sensitivity analysis, incurring additional computational cost. In contrast, G-PMC AIS operates effectively without such prerequisites.

Fractile comparisons with the benchmark—derived using MC with $N_t = 10^7$ ($N_\theta = 10^3$), G-PMC AIS ($N_t = 10^5$), and the five-point logic tree ($N_t = 6.25 \times 10^6$, $N_\theta = 625$)—are shown in Figure S4.

For fault sources, the performance gap remains substantial. At a 1% COV, G-PMC AIS is 13–24 times, 98–183 times, and 123–2268 times more efficient than the three-point logic tree, five-point logic tree, and MC methods, respectively. The K-S D values for the 1% COV cases are 3.6%, 2.8%, 3.9%, and 2.8% for 0.13g, 0.32g, 0.64g, and 1.1g, respectively, again demonstrating strong agreement with the benchmark. Interestingly, when using small sample sizes, fault sources outperform areal sources in fractile estimation, likely due to reduced problem dimensionality (Figure 6). In our implementation, one spatial variable (longitude) is eliminated in fault source, due to the vertically lying fault, thereby reducing computational complexity.

These numerical results demonstrate that the proposed G-PMC AIS framework is, in general, faster than any of the alternative methods considered. However, caution is warranted: accurate fractile hazard estimation depends on whether the iterated density adequately approximates the optimal density. In practical PSHA applications, this approximation quality is not known, as verifying each point estimate against the true value would require the same computational effort as traditional numerical integration, negating the benefit of the proposed method. Nevertheless, the numerical examples suggest that when a sufficiently low COV (e.g., $< 3\%$) is achieved for the mean hazard, the corresponding fractiles also exhibit good agreement. Conversely, when the number of samples is insufficient (e.g., $N = 10^3$ in areal source case; Fig. 5 (a)), the K-S D statistic exceed 50% in some cases, indicating a significant deviation from the true fractile and illustrating a failure to approximate the optimal density under limited sampling. This highlights the importance of using an adequate number of samples in G-PMC AIS, though the required sample sizes remain substantially lower than for other methods. Unlike MC or logic-tree approaches, where accuracy improves incrementally with sample size, G-PMC AIS exhibits a "phase transition" behavior—convergence toward the optimal density occurs only after a critical threshold in N_t is surpassed.

For the areal source case, a sensitivity analysis based on the likelihood ratio between the marginalized densities, as defined in Eq. (18), is presented in Figure 7. Note that in this example, since all epistemic uncertainty variables are independently distributed, the first-order Sobol indices can be effectively approximated by C_i, calculated using Eq. (18). The results closely match

the benchmark first-order Sobol indices computed using Eq. (B1), with a maximum difference of 0.013. The results indicate that uncertainties in the ground motion model—specifically, the median ($\Delta\mu$) and standard deviation ($\Delta\sigma$)—are the dominant contributors to overall sensitivity, consistent with findings from previous PSHA studies (Sabetta et al., 2005; Litchfield et al., 2011; Koukouvelas and Papadopoulos, 2024; Bommer and Abrahamson, 2006). In contrast, the influence of seismic source parameters, such as $m_{\max}$ and the b-value, is relatively minor. The notable shift in the distribution of $\Delta\sigma$—highlighting its central role in shaping the hazard—can be seen in the comparison between the marginalized iterated proposal densities (approximated optimal density) and the original prior distributions (Figure S3).

Additionally, the interaction effect among variables, represented by ρ in Figure 7, generally exhibits minimal contribution across most scenarios. However, its influence tends to increase at higher ground motion levels, underscoring the complex nature of epistemic uncertainty interactions in the low exceedance rate hazard.

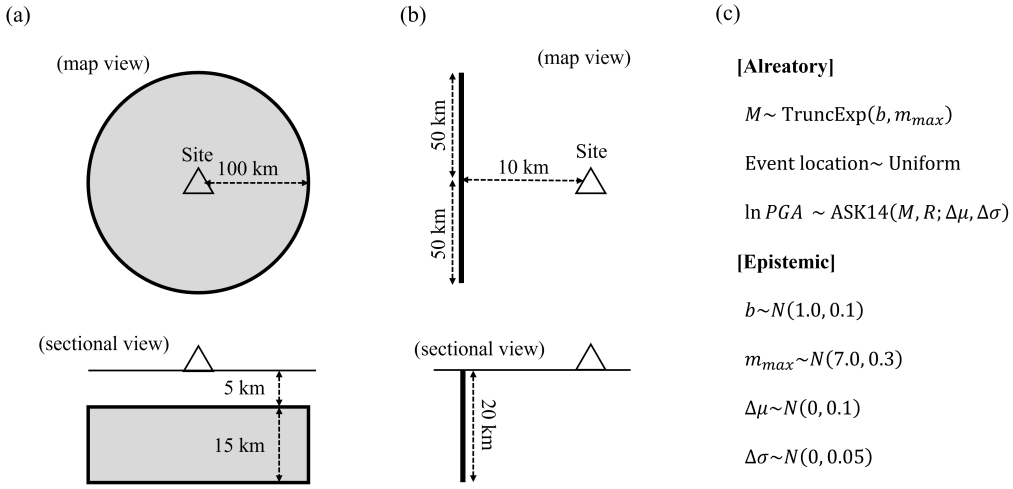

Figure 4. Seismic source geometry for the numerical examples: (a) area source and (b) fault source. (c) Aleatory and epistemic uncertainty variables, along with their associated probability distributions.

CONCLUSION

Probabilistic seismic hazard analysis (PSHA) has traditionally been a computationally intensive task, particularly when complex logic-tree structures with numerous epistemic uncertainty variables are involved.

We demonstrated that approximating the underlying continuous distributions with discrete representations can inherently lead to biased estimates of both mean and fractile hazards, especially when arbitrary weighting schemes are used in the approximation process. From a simple numerical example, it is shown that the mean, 16th, and 84th fractile hazards differ by factors of $\sim$ 1.8, 2.7, and 1.8, respectively, which can be even larger when the lower exceedance probability is considered. Furthermore, the computational burden of hazard calculation increases exponentially with the number of epistemic uncertainty variables.

To address these issues, we suggest the use of continuous distributions to preserve the original characteristics of the epistemic variables. Furthermore, to mitigate the computational challenges associated by incorporating high-dimensional continuous distri-

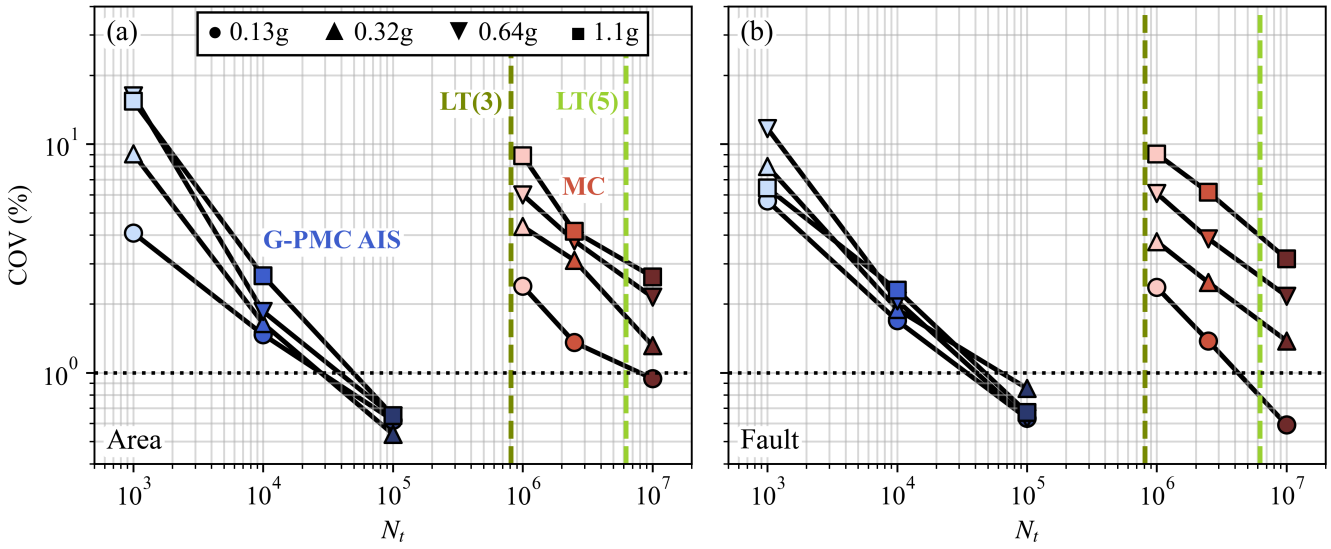

Figure 5. Coefficient of variation (COV) of the mean hazard as a function of the total number of estimated marginal hazards ($N_t = N_\theta \times N_\phi$) for (a) areal and (b) fault source examples. Bluish plots represent G-PMC AIS results, while reddish plots indicate Monte Carlo (MC) results, with color intensity corresponding to the value of N_t. The circle, triangle, inverted triangle, and square markers correspond to ground motion intensities of 0.13g, 0.32g, 0.64g, and 1.1g, respectively. Greenish vertical dashed lines indicate the total number of marginal hazards, N_t, for the three- and five-point logic tree approximations. The black horizontal dotted line marks the 1% COV threshold.

butions, a PSHA computational framework based on Gaussian Population Monte Carlo Adaptive Importance Sampling (G-PMC AIS), is proposed.

The key insights of the G-PMC AIS approach in calculating PSHA are: 1) the mean hazard can be efficiently estimated G-PMC AIS approach based on the formulation treating aleatory and epistemic uncertainty variables jointly, thereby avoiding the need for double-nested integration and improving computational efficiency; 2) individual hazard estimates used for fractile computation can be obtained efficiently, almost without additional cost, by utilizing the likelihood ratio between marginalized optimal density, which is a by-product of mean estimation, and the original density, with the easy marginalization by adopting joint Gaussian distribution; 3) The hazard sensitivity analysis also can be done easily in the same way.

We demonstrated the effectiveness and validity of the proposed methodology through numerical examples. The results indicate that the G-PMC AIS framework outperforms three-point approximation logic tree, five-point approximation logic tree, and continuous distribution based MC approach by a factor of up to 29, 224, and 3775 times while achieving high accuracy (COV < 1%) of mean hazard estimation.

Moreover, K-S D statistic between fractile hazards for those cases with high mean hazard accuracies and the benchmark all fall below $\sim 5\%$, indicating extremely similar hazard fractile distribution, indicating the suggested framework is able to approximate the optimal density well. The likelihood ratio between marginalized each epistemic uncertainty variable's optimal density and marginalized presumed epistemic uncertainty also enables us to quantify the sensitivity of each variable to the hazard.

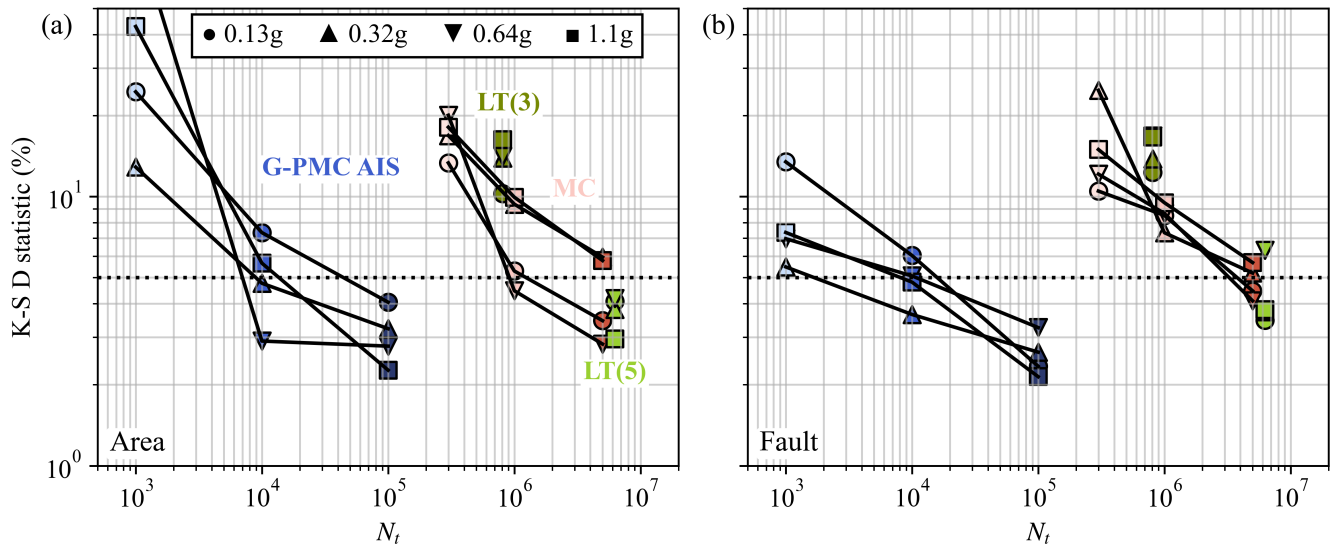

Figure 6. Kolmogorov–Smirnov (K-S) D statistic of the hazard distribution with respect to the epistemic uncertainty variable θ, plotted as a function of the total number of estimated marginal hazards ($N_t = N_\theta \times N_\phi$) for (a) areal and (b) fault source examples. Bluish plots represent G-PMC AIS results, while reddish plots represent Monte Carlo (MC) results, with color intensity indicating the value of N_t. Circles, triangles, inverted triangles, and squares correspond to ground motion intensities of 0.13g, 0.32g, 0.64g, and 1.1g, respectively. Greenish scatter plots indicate N_t and corresponding K-S D values for the three- and five-point logic tree approximations. The 5% threshold for the K-S D statistic is shown as a black horizontal dotted line.

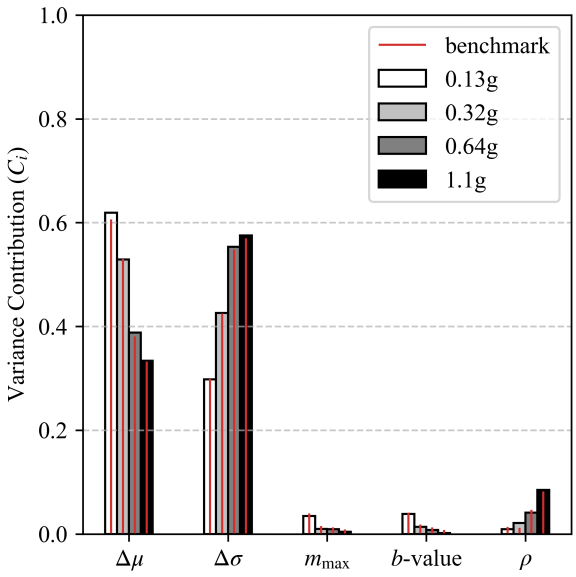

Figure 7. Sensitivity analysis for the areal source case. Grayscale bars indicate the variance contribution (first-order Sobol indices, $\mathcal{V}_i$) of each uncertainty variable at different ground motion levels. The red bars represent benchmark first-order Sobol indices for comparison. The interaction effect, denoted by ρ, is calculated as the remainder term $\rho = 1 - \sum_{i=1}^{4} S_i$.

We anticipate that the proposed framework can be effectively applied to a wide range of PSHA applications involving complex logic-tree structures with extensive epistemic uncertainties.

DATA AND RESOURCES

The source code for computing mean and fractile hazards using the framework explained in this paper is available at https://github.com/sehoung/ais_psha. All the Figures were created using Python Matplotlib (Hunter, 2007) and Microsoft PowerPoint (http://office.microsoft.com). The website was last accessed in May 2025. The supplemental material for this article includes 1 algorithm and 4 figures.

DECLARATION OF COMPETING INTERESTS

The authors have filed a patent application related to the method discussed in this paper. This application is relevant to the research presented and could represent a potential financial interest.

ACKNOWLEDGMENTS

The authors acknowledge the financial support provided by the Department of Civil and Environmental Engineering at the University of California, Berkeley. We are also grateful to Kurt Walter Soncco Sinchi for the valuable discussions and for providing the Python codes used in the numerical examples presented in this study.

REFERENCES

Abrahamson, N. A. and J. J. Bommer (2005). Probability and uncertainty in seismic hazard analysis. *Earthquake spectra* **21**(2), 603–607.

Abrahamson, N. A., W. J. Silva, and R. Kamai (2014). Summary of the ask14 ground motion relation for active crustal regions. *Earthquake Spectra* **30**(3), 1025–1055.

Al Atik, L. and R. R. Youngs (2014). Epistemic uncertainty for nga-west2 models. *Earthquake Spectra* **30**(3), 1301–1318.

ASCE (2022). Minimum Design Loads for Buildings and Other Structures. ASCE/SEI Standard 7-22. American Society of Civil Engineering, Reston, Virginia.

Assatourians, K. and G. M. Atkinson (2013). Eqhaz: An open-source probabilistic seismic-hazard code based on the monte carlo simulation approach. *Seismological Research Letters* **84**(3), 516–524.

Baker, J., B. Bradley, and P. Stafford (2021). *Seismic hazard and risk analysis*. Cambridge University Press.

Blume, J. A. and A. S. Kiremidjian (1979). Probabilistic procedures for peak ground motions. *Journal of the Structural Division* **105**(11), 2293–2311.

Bommer, J. J. (2012). Challenges of building logic trees for probabilistic seismic hazard analysis. *Earthquake Spectra* **28**(4), 1723–1735.

Bommer, J. J. and N. A. Abrahamson (2006). Why do modern probabilistic seismic-hazard analyses often lead to increased hazard estimates?

Bommer, J. J. and F. Scherbaum (2008). The use and misuse of logic trees in probabilistic seismic hazard analysis. *Earthquake Spectra* **24**(4), 997–1009.

Bugallo, M. F., V. Elvira, L. Martino, D. Luengo, J. Miguez, and P. M. Djuric (2017). Adaptive Importance Sampling: The past, the present, and the future. *IEEE Signal Processing Magazine* **34**(4), 60–79.

Cappé, O., A. Guillin, J.-M. Marin, and C. P. Robert (2004). Population monte carlo. *Journal of Computational and Graphical Statistics* **13**(4), 907–929.

Ceferino, L., J. Mitrani-Reiser, A. Kiremidjian, G. Deierlein, and C. Bambarén (2020). Effective plans for hospital system response to earthquake emergencies. *Nature Communications* **11**(1), 4325.

Christou, V., P. Bocchini, M. J. Miranda, and A. Karamlou (2018). Effective sampling of spatially correlated intensity maps using hazard quantization: Application to seismic events. *ASCE-ASME Journal of Risk and Uncertainty in Engineering Systems, Part A: Civil Engineering* **4**(1), 04017035.

Coppersmith, K. J., L. A. Salomone, C. W. Fuller, L. L. Glaser, K. L. Hanson, R. D. Hartleb, W. R. Lettis, S. C. Lindvall, S. M. McDuffie, R. K. McGuire, et al. (2012). Central and eastern united states (ceus) seismic source characterization (ssc) for nuclear facilities project. Technical report, Electric Power Research Institute (EPRI).

Cornell, C. A. (1968). Engineering seismic risk analysis. *Bulletin of the Seismological Society of America* **58**(5), 1583–1606.

Crowley, H. and J. J. Bommer (2006). Modelling seismic hazard in earthquake loss models with spatially distributed exposure. *Bulletin of Earthquake Engineering* **4**, 249–273.

Esteva, L. and M. Ordaz (1963). Regionalizacion sismica de la republica mexicana. *Revista Sociedad Mexicana do Ingenieria Sismica* **1**(1), 31–35.

Friedman, J. H. and M. H. Wright (1981). A nested partitioning procedure for numerical multiple integration. *ACM Transactions on Mathematical Software (TOMS)* **7**(1), 76–92.

Hale, C., N. Abrahamson, and Y. Bozorgnia (2018). Probabilistic Seismic Hazard Analysis Code Verification. Technical report, Pacific Earthquake Engineering Research Center, Berkeley.

Han, Y. and R. A. Davidson (2012). Probabilistic seismic hazard analysis for spatially distributed infrastructure. *Earthquake Engineering & Structural Dynamics* **41**(15), 2141–2158.

Hanks, T. C., N. A. Abrahamson, D. M. Boore, K. J. Coppersmith, and N. E. Knepprath (2009). Implementation of the sshac guidelines for level 3 and 4 pshas—experience gained from actual applications. *US Geological Survey Open-File Report* **1093**, 66.

Houng, S. E. and L. Ceferino (2025). Fast probabilistic seismic hazard analysys through adaptive importance sampling. *Bulletin of the Seismological Society of America* **115**(2), 646–663.

Hunter, J. D. (2007). Matplotlib: A 2d graphics environment. *Computing in science & engineering* **9**(03), 90–95.

Jayaram, N. and J. W. Baker (2010). Efficient sampling and data reduction techniques for probabilistic seismic lifeline risk assessment. *Earthquake Engineering & Structural Dynamics* **39**(10), 1109–1131.

Keefer, D. L. and S. E. Bodily (1983). Three-point approximations for continuous random variables. *Management Science* **29**(5), 595–609.

Kennedy, R. P., C. A. Cornell, R. D. Campbell, S. Kaplan, and H. F. Perla (1980). Probabilistic seismic safety study of an existing nuclear power plant. *Nuclear Engineering and Design* **59**(2), 315–338.

Kiremidjian, A. S., E. Stergiou, and R. Lee (2007). Issues in seismic risk assessment of transportation networks. *Geotechnical, Geological and Earthquake Engineering* **6**, 461–480.

Kolmogorov, A. (1933). Sulla determinazione empirica di una legge didistribuzione. *Giorn Dell'inst Ital Degli Att* **4**, 89–91.

Koukouvelas, I. K. and G. A. Papadopoulos (2024). The main factors controlling probabilistic seismic hazard assessment in greece. *The Seismic Record* **4**(2), 93–115.

Kwong, N. S. and K. S. Jaiswal (2023). Uses of epistemic uncertainties in the usgs national seismic hazard models. *Earthquake Spectra* **39**(2), 1058–1087.

Lacour, M. and N. Abrahamson (2021). Efficient propagation of epistemic uncertainty for probabilistic seismic hazard analyses (pshas) including partial correlation of magnitude–distance scaling. *Bulletin of the Seismological Society of America* **111**(6), 3332–3340.

Lacour, M. and N. Abrahamson (2025). Reducing calculation times for seismic hazard using non-ergodic ground-motion models for areal source zones. *Applied Sciences* **15**(5), 2454.

Lacour, M. and N. A. Abrahamson (2019). Efficient propagation of epistemic uncertainty in the median ground-motion model in probabilistic hazard calculations. *Bulletin of the Seismological Society of America* **109**(5), 2063–2072.

Lepage, G. (1978). A new algorithm for adaptive multidimensional integration. *Journal of Computational Physics* **27**(2), 192–203.

Litchfield, N. J., M. W. Stirling, R. J. Van Dissen, et al. (2011). Consideration and propagation of epistemic uncertainties in new zealand probabilistic seismic hazard analysis. Technical Report GNS Science Consultancy Report 2011/275, GNS Science. Available from GNS Science, New Zealand.

Manzour, H., R. A. Davidson, N. Horspool, and L. K. Nozick (2016). Seismic hazard and loss analysis for spatially distributed infrastructure in Christchurch, New Zealand. *Earthquake Spectra* **32**(2), 697–712.

Marzocchi, W., M. Taroni, and J. Selva (2015). Accounting for epistemic uncertainty in psha: Logic tree and ensemble modeling. *Bulletin of the Seismological Society of America* **105**(4), 2151–2159.

McGuire, R. K. (2008). Probabilistic seismic hazard analysis: Early history. *Earthquake Engineering and Structural Dynamics* **37**(3), 329–338.

Miller, A. C. and T. R. Rice (1983). Discrete approximations of probability distributions. *Management science* **29**(3), 352–362.

Musson, R. M. (2000). The use of monte carlo simulations for seismic hazard assessment in the uk. *Annals of Geophysics* **43**(1), 1–9.

Ordaz, M., M. A. Salgado-Gálvez, and S. Giraldo (2021). R-crisis: 35 years of continuous developments and improvements for probabilistic seismic hazard analysis. *Bulletin of Earthquake Engineering* **19**(7), 2797–2816.

Papadopoulos, A. N. and P. Bazzurro (2021). Exploring probabilistic seismic risk assessment accounting for seismicity clustering and damage accumulation: Part II. Risk analysis. *Earthquake Spectra* **37**(1), 386–408.

Petersen, M. D., A. M. Shumway, P. M. Powers, E. H. Field, M. P. Moschetti, K. S. Jaiswal, K. R. Milner, S. Rezaeian, A. D. Frankel, A. L. Llenos, et al. (2024). The 2023 us 50-state national seismic hazard model: Overview and implications. *Earthquake Spectra* **40**(1), 5–88.

Porter, K., K. Milner, and E. Field (2025). Trimming the ucerf3-td logic tree: Model order reduction for an earthquake rupture forecast considering loss exceedance. *Earthquake Spectra* **41**(1), 636–653.

Press, W. H. and G. R. Farrar (1990). Recursive stratified sampling for multidimensional monte carlo integration. *Computers in Physics* **4**(2), 190–195.

Rahimi, H. and M. Mahsuli (2019). Structural reliability approach to analysis of probabilistic seismic hazard and its sensitivities. *Bulletin of Earthquake Engineering* **17**(3), 1331–1359.

Rubinstein, R. Y. (1997). Optimization of computer simulation models with rare events. *European Journal of Operational Research* **99**(1), 89–112.

Rubinstein, R. Y. and D. P. Kroese (2004). *The cross-entropy method: a unified approach to combinatorial optimization, Monte-Carlo simulation, and machine learning*. Springer.

Sabetta, F., A. Lucantoni, H. Bungum, and J. J. Bommer (2005). Sensitivity of psha results to ground motion prediction relations and logic-tree weights. *Soil dynamics and earthquake engineering* **25**(4), 317–329.

Sadigh, K., C.-Y. Chang, J. Egan, F. Makdisi, and R. R. Youngs (1997). Attenuation relationships for shallow crustal earthquakes based on california strong motion data. *Seismological research letters* **68**(1), 180–189.

Saltelli, A., S. Tarantola, F. Campolongo, and M. Ratto (2004). *Sensitivity Analysis in Practice: A Guide to Assessing Scientific Models*. Chichester, UK: John Wiley & Sons.

Silva, V., D. Amo-Oduro, A. Calderon, C. Costa, J. Dabbeek, V. Despotaki, L. Martins, M. Pagani, A. Rao, M. Simionato, D. Viganò, C. Yepes-Estrada, A. Acevedo, H. Crowley, N. Horspool, K. Jaiswal, M. Journeay, and M. Pittore (2020). Development of a global seismic risk model. *Earthquake Spectra* **36**, 372–394.

Sobol', I. M. (2001). Global sensitivity indices for nonlinear mathematical models and their monte carlo estimates. *Mathematics and Computers in Simulation* **55**(1-3), 271–280.

Theodoridis, S. (2015). *Machine learning: a Bayesian and optimization perspective*. Academic press.

U. S. Nuclear Regulatory Commission (2007). A Performance Based Approach to Define Site Specific Ground Motion. Regulatory Guide 1.208, USNRC.

LIST OF FIGURE CAPTIONS

- **Figure 1.** (a) Geometry of the seismic source and site location used in the numerical example to illustrate hazard bias through logic trees; (b) the assumed continuous distributions of the epistemic uncertainty variables, along with an example of the discretized distribution used to construct a logic tree based on the KB83 approximation (Table 1). Three epistemic uncertainty variables are incorporated: the characteristic earthquake magnitude (M_{char}), and the median and standard deviation of the GMM ($\Delta\mu$ and $\Delta\sigma$). (c) Magnitude probability mass function, and (d) peak ground acceleration (PGA) PDF for the base case (black) and one possible alternative case (red). In this example, the base case corresponds to the highest-weighted values in the logic tree structure, while the alternative represents a lower-weighted case. Each curve represents the aleatory uncertainty of the variables, while the different cases reflect epistemic uncertainty.
- **Figure 2.** (a)–(c): Mean, 16th, and 84th fractile hazard curves from the numerical example. The benchmark results are shown as thick solid gray lines. Hazard curves based on the three-point approximation of Keefer and Bodily (1983) are shown as black dotted lines; those using the approximation from Petersen et al. (2024) are shown as red dotted lines; and the five-point approximation from Miller and Rice (1983) is presented as blue dotted lines. (d)–(f): Relative errors of the hazard curves obtained using different approximation strategies, evaluated against the benchmark at various target ground motion levels. (g)–(i): Relative errors of the ground motion intensity obtained using different approximation strategies, evaluated against the benchmark ground motion intensities at various hazard levels.
- **Figure 3.** Diagram illustrating the proposed framework using a simplified 2D PSHA problem with one aleatory variable (ϕ: magnitude, m) and one epistemic variable (θ: median ground motion, $\Delta\mu$), for a target ground motion level of 0.5g. The magnitude follows a truncated exponential distribution (M5.0–6.5), and epistemic uncertainty in ln [median ground motion] relative to the base model (Sadigh et al., 1997) is modeled as $\mathcal{N}(0, 0.2)$. In "Stage 1. Initialize", the joint distribution $f_{\Phi,\Theta}(\phi, \theta)$ is shown with darker regions indicating higher density. The marginal epistemic distribution $f_{\Theta}(\theta)$ is shown as a blue solid curve. "Stage 2. Proposal IS Density and Mean Hazard" presents the iterated proposal density $\hat{q}^*_{\Phi,\Theta}(\phi, \theta)$ and marginalized form $\hat{q}^*_{\Theta}(\theta)$ (blue-dotted curve), used to compute the estimated mean hazard $\hat{\lambda}(a)$ (red dot in "Final Hazard"). In "Stage 3. Fractile Hazard", fractile sampling from $f_{\Theta}(\theta)$ is illustrated, with selected samples marked as reddish dots. The blue solid and blue dotted curves represent the original and the iterated proposal densities of epistemic uncertainty, respectively. The likelihood ratio $\hat{q}^*_{\Theta}(\theta)/f_{\Theta}(\theta)$ and individual hazard $\vec{\lambda}(a|\theta)$ are computed, and the full hazard distribution is constructed by sorting individual hazards to identify fractile values, also shown in the “Final Hazard” box.
- **Figure 4.** Seismic source geometry for the numerical examples: (a) area source and (b) fault source. (c) Aleatory and epistemic uncertainty variables, along with their associated probability distributions.
- **Figure 5.** Coefficient of variation (COV) of the mean hazard as a function of the total number of estimated marginal hazards ($N_t = N_\theta \times N_\phi$) for (a) areal and (b) fault source examples. Bluish plots represent G-PMC AIS results, while reddish plots indicate Monte Carlo (MC) results, with color intensity corresponding to the value of N_t. The circle, triangle, inverted triangle,

and square markers correspond to ground motion intensities of 0.13g, 0.32g, 0.64g, and 1.1g, respectively. Greenish vertical dashed lines indicate the total number of marginal hazards, N_t, for the three- and five-point logic tree approximations. The black horzontal dotted line marks the 1% COV threshold.

- **Figure 6.** Kolmogorov–Smirnov (K-S) D statistic of the hazard distribution with respect to the epistemic uncertainty variable θ, plotted as a function of the total number of estimated marginal hazards ($N_t = N_\theta \times N_\phi$) for (a) areal and (b) fault source examples. Bluish plots represent G-PMC AIS results, while reddish plots represent Monte Carlo (MC) results, with color intensity indicating the value of N_t. Circles, triangles, inverted triangles, and squares correspond to ground motion intensities of 0.13g, 0.32g, 0.64g, and 1.1g, respectively. Greenish scatter plots indicate N_t and corresponding K-S D values for the three- and five-point logic tree approximations. The 5% threshold for the K-S D statistic is shown as a black horizontal dotted line.
- **Figure 7.** Sensitivity analysis for the areal source case. Grayscale bars indicate the variance contribution (first-order Sobol indices, $\mathcal{V}_i$) of each uncertainty variable at different ground motion levels. The red bars represent benchmark first-order Sobol indices for comparison. The interaction effect, denoted by ρ, is calculated as the remainder term $\rho = 1 - \sum_{i=1}^{4} S_i$.

FULL MAILING ADDRESS FOR EACH AUTHOR

Soung Eil Houng

University of California, Berkeley

Department of Civil and Environmental Engineering

Davis Hall, Berkeley, 94720, CA

shoung@berkeley.edu

Luis Ceferino

University of California, Berkeley

Department of Civil and Environmental Engineering

Davis Hall, Berkeley, 94720, CA

ceferino@berkeley.edu

Norm Abrahamson

University of California, Berkeley

Department of Civil and Environmental Engineering

Davis Hall, Berkeley, 94720, CA

abrahamson@berkeley.edu

APPENDIX A. INDIVIDUAL HAZARD CALCULATION USING MARGINALIZED IS OPTIMAL DENSITY

We can rewrite Eq. (15) as:

$$\nu(\theta) I(u(\phi;\theta) > a) f_{\Phi|\Theta}(\phi|\theta) = \lambda(a) \frac{q^*_{\Phi|\Theta}(\phi,\theta)}{f_\Theta(\theta)}$$

By integrating both sides of the equation with respect to ϕ, we can obtain the following expression:

$$\int_\phi \nu(\theta) I(u(\phi;\theta) > a) f_{\Phi|\Theta}(x|\theta) d\phi = \int_\phi \lambda(a) \frac{q^*_{\Phi|\Theta}(\phi,\theta)}{f_\Theta(\theta)} d\phi$$

Note that the left-hand side of the equation is the same form with the individual hazard (Eq. (1)). Therefore,

$$\begin{aligned}
\lambda(a|\theta) &= \int_\phi \lambda(a) \frac{q^*_{\Phi|\Theta}(\phi,\theta)}{f_\Theta(\theta)} d\phi \\
&= \int_\phi \lambda(a) \frac{q^*_{\Phi|\Theta}(\phi|\theta) q^*_\Theta(\theta)}{f_\Theta(\theta)} d\phi \\
&= \lambda(a) \frac{q^*_\Theta(\theta)}{f_\Theta(\theta)} \int_\phi q^*_{\Phi|\Theta}(\phi|\theta) d\phi
\end{aligned}$$

Since $q^*_{\Phi|\Theta}(\phi|\theta)$ is PDF with respect to ϕ, $\int_\phi q^*_{\Phi|\Theta}(\phi|\theta) d\phi$ reduces to 1, yielding:

$$\lambda(a|\theta) = \lambda(a) \frac{q^*(\theta)}{f_\Theta(\theta)}$$

APPENDIX B. FIRST ORDER SOBOL INDEX AND MARGINALIZED IS OPTIMAL DENSITY

The first-order Sobol index of i-th epistemic uncertainty variable (Sobol', 2001; Saltelli et al., 2004), θ_i, is calculated to be:

$$V_i = \frac{\mathrm{Var}_{\theta_i}[E_{\tilde{\theta}_i}[\lambda(a|\theta)|\theta_i)]]}{\mathrm{Var}_\theta[\lambda(a|\theta)]} \tag{B1}$$

, where $\tilde{\theta}_i$ is the set of epistemic uncertainty variables excluding θ_i, $E_{\tilde{\theta}_i}[\lambda(a|\theta)|\theta_i)]$ is the mean hazard with respect to $\tilde{\theta}_i$ when θ_i is fixed, and the whole numerator, $\mathrm{Var}_{\theta_i}[E_{\tilde{\theta}_i}[\lambda(a|\theta)|\theta_i)]]$, shows how the mean hazard is influenced by different θ_i. The denominator, $\mathrm{Var}_\theta[\lambda(a|\theta)]$, is the total variance of individual hazard with respect to θ. V_i ranges from 0 to 1, and representing how much θ_i directly affects the hazard. The sum of V_i over the all the epistemic uncertainty variables is less than or equal to one, where the remainder, $1 - \sum_i^{N_\theta} V_i$, where N_θ is total number of epistemic uncertainty variables, reflects the interactions between these variables.

Here, we show how Eq. (18) is equivalent to the first-order Sobol index Eq. (B1) if θ_i is independent of $\tilde{\theta}_i$. Eq. (18) and Eq. (B1) share the denominator; thus, we show the numerators of Eq. (18) and Eq. (B1) are equivalent. The integration form of the term in the numerator of Eq. (B1) is:

$$\begin{aligned}\mathrm{VAR}_{\theta_i}[E_{\tilde{\theta}_i}[\lambda(a|\theta)|\theta_i)]] &= \mathrm{VAR}_{\theta_i}\left[\int_{\tilde{\theta}_i}\lambda(a|\tilde{\theta}_i,\theta_i)f_{\tilde{\Theta}_i}(\tilde{\theta}_i)d\tilde{\theta}_i\right] \\ &= \mathrm{VAR}_{\theta_i}\left[\lambda(a)\int_{\tilde{\theta}_i}\frac{q_\Theta(a|\tilde{\theta}_i,\theta_i)}{f_\Theta(a|\tilde{\theta}_i,\theta_i)}f_{\tilde{\Theta}_i}(\tilde{\theta}_i)d\tilde{\theta}_i\right]\end{aligned} \tag{B2}$$

717 If θ_i is independent of all $\tilde{\theta}_i$, the Eq. (B2) reduces to:

$$\begin{aligned}\mathrm{VAR}_{\theta_i}[E_{\tilde{\theta}_i}[\lambda(a|\theta)|\theta_i)]] &= \mathrm{VAR}_{\theta_i}\left[\lambda(a)\int_{\tilde{\theta}_i}\frac{q_\Theta(\tilde{\theta}_i,\theta_i)}{f_{\tilde{\Theta}_i}(\tilde{\theta}_i)f_{\Theta_i}(\theta_i)}f_{\tilde{\Theta}_i}(\tilde{\theta}_i)d\tilde{\theta}_i\right] \\ &= \mathrm{VAR}_{\theta_i}\left[\frac{\lambda(a)}{f_{\Theta_i}(\theta_i)}\int_{\tilde{\theta}_i}q_\Theta(\tilde{\theta}_i,\theta_i)d\tilde{\theta}_i\right] \\ &= \mathrm{VAR}_{\theta_i}\left[\lambda(a)\frac{q_{\Theta_i}(\theta_i)}{f_{\Theta_i}(\theta_i)}\right]\end{aligned} \tag{B3}$$

718 By Eq. (17),

$$\mathrm{VAR}_{\theta_i}\left[\lambda(a)\frac{q_{\Theta_i}(\theta_i)}{f_{\Theta_i}(\theta_i)}\right] = \mathrm{VAR}_{\theta_i}[\lambda_i(a|\theta_i)] \tag{B4}$$

719 , which is equal to the numerator in Eq. (18).

720

End of the Manuscript