A Continuous-Time Markovian Framework for Performance-Based Earthquake Engineering: Integrated Seismic Risk and Resilience Assessment

Chiara Nardin, Stefano Marelli, Bruno Sudret and Marco Broccardo
[1] *Department of Civil, Environmental and Mechanical Engineering, University of Trento, 38132 Trento, Italy - {chiara.nardin, marco.broccardo}@unitn.it*

[2] *Chair of Risk, Safety and Uncertainty Quantification,ETH Zurich, 8093 Zurich, Switzerland - {cnardin,marellis,sudretb}@ethz.ch*

Abstract. Assessing the seismic reliability of structural systems increasingly requires models that jointly represent damage evolution and recovery. Yet, the traditional PEER PBEE formulation does not capture the temporal dynamics of post-event repair, and modern reliability analyses often rely on scarce or uncertain recovery data, motivating the need for tractable methods.

We reformulate the PEER PBEE framework using a continuous time Markov Chain (CTMC) Monte Carlo approach, which naturally models probabilistic transitions among damage and recovery states under time-invariant seismic hazards. Transition rates incorporate metamodel-based state-dependent fragility functions, as well as recovery processes, yielding a unified representation of hazard effects and system evolution. Leveraging the analytical properties of CTMCs, we compute steady-state performance statistics, such as mean time to failure and time spent in degraded states, and define a resilience index adapted from the classical reliability index β.

The approach is tested on a simple yet representative example, showing that the CTMC-enhanced PBEE framework captures the temporal progression from damage to recovery in a computationally efficient and data-flexible manner.

Keywords: PEER PBEE framework, seismic reliability, recovery, CTMC

1. Introduction

The PEER performance-based earthquake engineering (PBEE) framework has revolutionized seismic risk quantification by decomposing performance assessment into hazard, structural response, damage, and loss. Despite its success, the framework's foundational assumptions, formalized in the seminal works of (Cornell and Krawinkler, 2000) and (Moehle and Deierlein, 2004), feature specific limitations regarding temporal dynamics, as noted by (Der Kiureghian, 2005). Two critical gaps persist:

- probability-based vs. rate-based assessment: traditional PBEE evaluates risk through conditional probabilities of damage given an intensity measure. However, reliability is intrinsically rate-based. By focusing on event-based probabilities, the framework often misses the underlying temporal dynamics, specifically the transition rates that govern system evolution over time.

- absence of renewal mechanisms: real-world systems are not static; they degrade and recover. While standard PBEE typically evaluates structures on a per-event basis, it lacks an intrinsic renewal mechanism to account for the stochastic interplay between damage accumulation and restoration.

Significant research has sought to address the first point by modeling seismic damage accumulation. Several researchers (Guida and Pulcini, 2011; Iervolino et al., 2015a; Iervolino et al., 2015b) proposed state-dependent Markov models to capture cumulative path-dependence, while (Andriotis and Papakonstantinou, 2018) developed generalized frameworks for complex damage interactions. While these models represent a major step forward, they remain largely probabilistic in terms of damage exceedance. They do not explicitly represent the infinitesimal transition rates required to seamlessly couple damage accumulation and recovery within a single, time-dependent stochastic engine. Recent work by (Nardin et al., 2025; Nardin et al., 2026) bridged this gap by formulating damage accumulation via transition rates λ_{ij}, directly derived from surrogate model-based state-dependent PBEE fragility curves.

In parallel to these developments, the engineering community has prioritized the treatment of recovery processes, shifting from a risk perspective to a resilience one. Thus, the research on modeling recovery processes, registered a substantial progress over the past decade, not only in conceptual and probabilistic frameworks, but also in developing widely available engineering software that can embed recovery into seismic loss assessment. Indeed, many community- or regional-scale resilience studies rely on recovery functions to be included into classical loss and damage assessment procedure. A prominent example is the FEMA HAZUS model and manual (Federal Emergency Management Agency (FEMA), 2020), which provides a national database of generalized damage, loss, and recovery curves for use in large-scale planning and risk assessment (native to U.S. practice). Similar examples include the P-58 project (Federal Emergency Management Agency (FEMA), 2018a; Federal Emergency Management Agency (FEMA), 2018b), carried out by FEMA, and the ongoing ATC-138 project (Applied Technology Council, 2021), that provides recovery-based design tools and guidelines for modern buildings.

Nevertheless, a gap still remains: most existing formulations treat recovery as a post-event *add-on* rather than a competing mechanism between degradation and repair interventions. Hence, a unified framework capable of representing the continuous-time evolution of system performance under the simultaneous mechanisms of hazard-driven degradation and repair-driven renewal is still missing.

1.1. Contribution

Based on these premises, in this paper, we address these challenges by generalizing the PEER PBEE framework to a continuous-time Markov chain (CTMC) environment. We adopt a twofold strategy: we first propose a unified stochastic dynamics approach, where we represent both damage and recovery as stochastic state transitions, allowing for a rigorous treatment of system performance

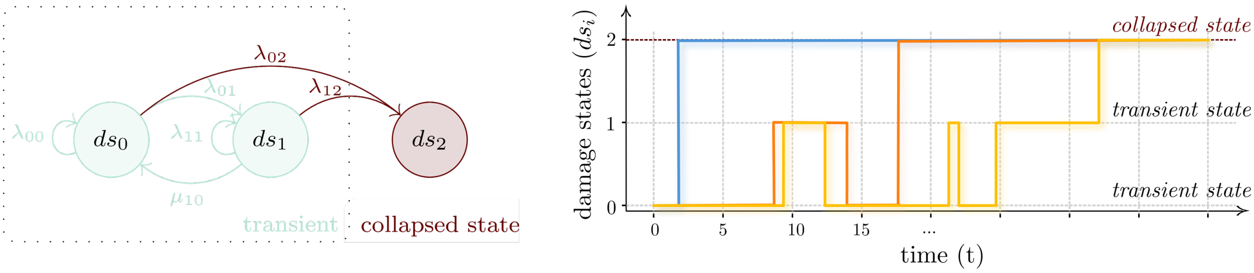

(a) (b)

Figure 1. (a) State transition diagram for the 3-state CTMC model, and (b) example of possible CTMC trajectories simulating transitions from undamaged -0- to damaged -1- and collapse -2- states, along with the recovery path from state -1- to -0-.

over any decision horizon. We then shift to a rate-based approach, by showing that hazard effects can be embedded through state-dependent transition rates rather than static probabilities.

Under this generalized framework, structural collapse emerges naturally as an absorbing state, while repair actions introduce a renewal mechanism that provides a formal basis for resilience metrics. We conclude by validating these analytical CTMC solutions using Monte Carlo simulation, demonstrating the flexibility of the framework for complex engineering applications.

2. Methodological Framework

Classical PEER-PBEE evaluates performance through exceedance probabilities derived under Poissonian assumptions. While computationally efficient, this formulation implicitly enforces renewability and memorylessness: after each event, the structure is assumed to return to its original condition. Such assumptions hinder the consistent representation of cumulative damage, irreversible failure, and recovery dynamics over long service horizons.

To overcome these limitations, we propose a generalization of the PEER PBEE within a continuous-time Markov chain (CTMC) framework. The key conceptual shift is from probabilities of system damage/deterioration to rates of full stochastic dynamic evolution considering both damage/deterioration and recovery processes. Structural performance is represented as a stochastic process evolving among discrete damage states, thereby embedding degradation and recovery directly into the system dynamics.

Let $\mathcal{D} = \{ds_0, ds_1, \ldots, ds_C\}$ denote the finite set of possible structural damage states, ordered by increasing damage severity. Here, ds_0 represents the fully operational state, $ds_1, \ldots, ds_{C-1}$ denote progressively degraded conditions, and ds_C represents collapse. In the CTMC context, collapse is modeled as an absorbing state, representing an irreversible terminal condition.

The structural condition at time t is described by the stochastic variable $DS(t)$, with corresponding state probabilities

$$\pi_i(t) = P(DS(t) = ds_i), \qquad \boldsymbol{\pi}(t) = [\pi_0(t), \ldots, \pi_C(t)],$$

satisfying $\sum_{i=0}^{C} \pi_i(t) = 1$ at each time instant. In this framework, the evolution of the system depends solely on its current state, thus replacing the renewal assumption with a physically consistent memory of accumulated damage.

The CTMC is defined by its infinitesimal generator matrix $\boldsymbol{Q} = \{Q_{ij}\}$. For $i \neq j$, Q_{ij} represents the transition rate from state ds_i to state ds_j, while diagonal entries satisfy

$$Q_{ii} = -\sum_{j \neq i} Q_{ij}. \tag{1}$$

The time evolution of the state probability vector follows the Kolmogorov forward equation (see (Ross, 2014; Trivedi and Bobbio, 2017; Brémaud, 2020)):

$$\dot{\boldsymbol{\pi}}(t) = \boldsymbol{\pi}(t)\boldsymbol{Q}, \qquad \boldsymbol{\pi}(0) = \boldsymbol{\pi}_0, \tag{2}$$

the solution of which is available in closed form through the matrix exponential

$$\boldsymbol{\pi}(t) = \boldsymbol{\pi}_0 e^{\boldsymbol{Q}t}, \tag{3}$$

providing the exact transient occupancy probability of each damage state at any time.

One notable property of the generator matrix is its direct physical interpretation. Indeed, we can see it as a composition of a *damage* (the upper triangular of the matrix) and *recovery* contributions (the lower triangular). To exemplify, consider the 3-state system depicted in Fig. 1. Its generator matrix reads:

$$\boldsymbol{Q} = \boldsymbol{Q}_{\text{dam}} + \boldsymbol{Q}_{\text{rec}} = \begin{bmatrix} -\sum_{j>0}^{j=\text{C}} \lambda_{0j} & \lambda_{01} & \lambda_{0,\text{C}} \\ \mu_{10} & -\left(\mu_{10} + \sum_{j>1}^{j=\text{C}} \lambda_{1j}\right) & \lambda_{1\text{C}} \\ 0 & 0 & 0 \end{bmatrix}. \tag{4}$$

Damage transitions move the system toward more severe damage states; hence $\boldsymbol{Q}_{\text{dam}}$ is upper-triangular and characterized by transition rates λ_{ij} for $j > i$. These rates incorporate state-dependent fragility functions and the seismic hazard, as derived in (Nardin et al., 2025):

$$\lambda_{ij} = \int_{IM} \mathbb{P}(DS = ds_j \mid DS = ds_i, im)\, |d\lambda(im)|, \tag{5}$$

where $\lambda(im)$ denotes the hazard curve and the conditional probability $\mathbb{P}(DS = ds_j \mid DS = ds_i, im)$ represents the fragility conditioned on the current structural state and the hazard. This formulation generalizes PEER-PBEE by replacing exceedance probabilities with transition intensities and by explicitly conditioning on the instantaneous damage state. In the limiting case involving only ds_0 and ds_C, the classical two-state Poisson collapse model is recovered.

Recovery transitions, represented by $\boldsymbol{Q}_{\text{rec}}$, return the structure from more severe to less severe states. These transitions populate the lower-triangular portion of the generator matrix and are governed by the corresponding recovery rates μ_{ij} for $i > j$. Although constant rates are used here for computational tractability, the framework can in principle accommodate more complex time-dependent or even stochastic recovery models if required.

Moreover, leveraging the structure of the generator matrix, we can partition it into the contribution of transient states and absorbing states:

$$\boldsymbol{Q} = \begin{bmatrix} \boldsymbol{Q}_T & \boldsymbol{a} \\ \boldsymbol{0} & 0 \end{bmatrix}, \tag{6}$$

where $\boldsymbol{Q}_T$ corresponds to transitions among the non-collapse states ds_0, ds_1, and $\boldsymbol{a} = (\lambda_{0C}, \lambda_{1C})^\intercal$ to the transitions towards the absorbed state ds_C.

Let $\mathcal{T}$ be the random variable describing the first time instant when the system collapses, and $F_\mathcal{T}$ its cumulative distribution function (CDF). Then, the probability of the system entering the absorbing collapse state by time t is derived from the missing probability mass in the transient sub-generator $\boldsymbol{Q}_T$:

$$F_\mathcal{T}(t) = \mathbb{P}(\mathcal{T} \leq t) = 1 - S_\mathcal{T}(t) = 1 - \sum_{i \in \mathcal{D}_T} \pi_i(t) = 1 - \boldsymbol{\pi}_T(0) \cdot e^{\boldsymbol{Q}_T\, t} \cdot \boldsymbol{1}_T, \tag{7}$$

where $\mathcal{D}_T = [ds_0, ds_1]$ and $S_\mathcal{T}$ is the survival function. The product $\boldsymbol{\pi}_T(0) \cdot e^{\boldsymbol{Q}_T\, t}$ gives the probability distribution over all "safe" (i.e., non-collapsed) states at time t, and multiplication by the column vector of ones $\boldsymbol{1}_T$ sums these probabilities. This formulation highlights an important structural property of the transient sub-generator $\boldsymbol{Q}_T$: its rows do not sum to zero. In a standard generator matrix $\boldsymbol{Q}$, each row sums to zero because total probability is conserved. In contrast, the rows of $\boldsymbol{Q}_T$ sum to strictly negative values. The "missing" probability mass corresponds exactly to the rate at which probability flows out of the transient state space and into the absorbing collapse state (through the coupling vector $\boldsymbol{a}$).

Following standard structural reliability jargon (Ditlevsen, 1979), we consider transition to the absorbing state as failure, and therefore $F_\mathcal{T}(t)$ in Eq. (7) as a time-dependent probability of failure. We can therefore map this failure probability to a time-dependent reliability index β_t:

$$\beta_t = -\Phi^{-1}(F_\mathcal{T}(t)), \tag{8}$$

where Φ is the standard normal CDF.

Thus, the reliability index β_t, derived from the transient subgenerator $\boldsymbol{Q}_T$, measures the probability of no collapse by time t, but ignores how the system degrades beforehand. This limitation motivates a shift to structural resilience, which emphasizes sustained performance over time. Accordingly, in a structural resilience perspective, the focus moves from the time-to-collapse $\mathcal{T}$ to the occupation time of the undamaged state ds_0.

First consider the occupation time of the undamaged state ds_0 over the interval $[0, t]$, defined as

$$A_0(t) := \int_0^t \boldsymbol{1}_{\{DS(s)=ds_0\}}\, ds, \tag{9}$$

where $\boldsymbol{1}_{\{\cdot\}}$ is the indicator function. The random variable $A_0(t)$ represents the total time spent in the undamaged state during $[0, t]$.

The expected occupation time admits a closed-form expression in terms of the transient sub-generator $\boldsymbol{Q}_T$. Indeed, recalling that the mean of an indicator function equals the corresponding probability, one obtains

$$\mathbb{E}[A_0(t)] = \int_0^t \mathbb{E}[\mathbf{1}_{\{DS(s)=ds_0\}}]\, ds, = \int_0^t \mathbb{P}(DS(s) = ds_0)\, ds = \int_0^t \left[e^{\boldsymbol{Q}_T s}\right]_{00} ds, \tag{10}$$

where $\mathbb{P}(DS(s) = ds_0)$ is the unconditional probability of being in the undamaged state ds_0 at time s. One recognizes that this quantity coincides with the $(0,0)$ entry of the truncated fundamental matrix

$$\boldsymbol{N}(t) = \int_0^t e^{\boldsymbol{Q}_T s}\, ds = \boldsymbol{Q}_T^{-1}\left(e^{\boldsymbol{Q}_T t} - \boldsymbol{I}\right), \tag{11}$$

so that $\mathbb{E}[A_0(t)] = n_{00}(t)$, where $n_{ij}(t)$ denotes the (i,j) entry of $\boldsymbol{N}(t)$ and represents the expected total time spent in state ds_j, starting from state ds_i, over the interval $[0,t]$. As $t \to \infty$, the truncated fundamental matrix converges to the well-known fundamental matrix (see (Ross, 2014; Trivedi and Bobbio, 2017)) $\boldsymbol{N} = -\boldsymbol{Q}_T^{-1}$, whose (i,j) entry n_{ij} represents the expected total time spent in state ds_j before absorption, starting from state ds_i.

We now define the unconditional metric $\mathcal{R}_0^{\text{unc}}$ as the expected fraction of the horizon T_{hor} during which the system does not remain in the undamaged state ds_0, starting from ds_0:

$$\mathcal{R}_0^{\text{unc}} = 1 - \frac{n_{00}(T_{\text{hor}})}{T_{\text{hor}}}, \tag{12}$$

where $n_{00}(T_{\text{hor}}) = [\boldsymbol{N}(T_{\text{hor}})]_{00}$ is the $(0,0)$ entry of the truncated fundamental matrix evaluated at T_{hor}. This metric is unconditional in the sense that the expectation is taken over all possible trajectories, including those that reach the collapse state before T_{hor}. Equivalently, $\mathcal{R}_0^{\text{unc}}$ admits the probabilistic interpretation

$$\mathcal{R}_0^{\text{unc}} = \mathbb{P}(DS(U) \neq ds_0), \qquad U \sim \mathcal{U}(0, T_{\text{hor}}), \tag{13}$$

that is, the probability that a time instant U drawn uniformly at random from $[0, T_{\text{hor}}]$ falls within a period during which the system is not in the undamaged state, without conditioning on survival. Consequently, a high value of $\mathcal{R}_0^{\text{unc}}$ indicates a higher likelihood of sustained or progressive damage or early collapse, while a lower value reflects the system spends most of its time undamaged.

To maintain consistency with the reliability index β, we introduce the resilience index ρ^{unc} as

$$\rho^{\text{unc}} = -\Phi^{-1}(\mathcal{R}_0^{\text{unc}}), \tag{14}$$

where $\mathcal{R}_0 \in [0,1]$ ensures that ρ^{unc} is well defined. Since $\mathcal{R}_0^{\text{unc}}$ is the expected fraction of time spent in damaged states, the index ρ^{unc} provides a one-to-one transformation of the resilience metric into the standard normal space, directly comparable to β_t.

Interestingly, both the reliability and the resilience indices are fully consistent with the underlying damage accumulation and recovery mechanisms encoded in the generator matrix $\mathbf{Q}$. The latter provides therefore a unified basis for comparing alternative structural configurations, retrofitting strategies, or recovery policies within a single probabilistic framework.

Moreover, since both indices are directly from $\mathbf{Q}_T$, their estimation does not require costly Monte Carlo simulations, but they can be obtained directly by integrating Eq. (11). This makes the approach particularly suitable for long-time horizon assessments, where repeated simulation of event sequences would otherwise be computationally demanding.

3. A Simple Example: time evolution of a 3-state System

3.1. System Configuration and Parameters

To demonstrate the capabilities of the CTMC-generalized PBEE framework, we analyze a 3-state example that captures the fundamental competition between hazard-driven degradation and repair-driven recovery. This system allows for a clear visualization of how infinitesimal transition rates govern long-term structural performance.

As shown in Fig. 1, we consider a 3-state space $\mathcal{D} = \{0, 1, 2\}$, where state $c = 2$ represents structural collapse. The transition rates within the generator matrix $\boldsymbol{Q}$ are defined as:

$$\lambda_{01} = 0.02, \quad \lambda_{02} = 0.0001, \quad \lambda_{12} = 0.005. \tag{15}$$

and a time horizon T_{hor} of 50 years. These values are selected to reflect a physically consistent engineering scenario: the rate of direct collapse from an undamaged state (λ_{02}) is significantly lower than the rate of recoverable-damage from a pristine initial condition state (λ_{01}), capturing the general desired structural behavior.

Nevertheless, we let μ_{10}, i.e., the only recovery rate of the system, to range in the interval 10^{-3} to 10^{1} [years^{-1}], corresponding to a mean recovery time that varies from months to over a century. Fig. 2 illustrates the time-dependent occupancy probabilities over a 50-year horizon (T_{hor}) for two extreme scenarios: a low-recovery case ($\mu_{10} = 10^{-3}$) and a high-recovery case ($\mu_{10} = 10^{1}$). In panel (a), the near absence of recovery means that the system evolves in only one direction: probability mass gradually shifts from the undamaged state ds_0 to the damaged state ds_1, and eventually to collapse ds_2. This behavior closely reflects the implicit assumption of classical PEER-PBEE, where each seismic event affects a structure that is not repaired between events. As a result, long-term performance is governed almost entirely by the damage rates λ_{ij}. In this limit, the CTMC reproduces the traditional formulation, providing a useful consistency check.

Panel (b) shows a fundamentally different regime which can not be captured by the classical framework. When the recovery rate μ_{10} becomes comparable to the damage rate λ_{01}, damage and repair compete with each other and reach a balance. Instead of continuously degrading, the probability of being in the undamaged state ds_0 remains relatively high over time, reflecting ongoing repair. The probability of being in the damaged state ds_1 stays limited, and the rate of collapse is significantly reduced. This type of behavior cannot be described using event-based exceedance probabilities alone; it requires accounting for the transition rates between states in both directions, which is what the generator matrix $\mathbf{Q}$ represents.

Taken together, these panels highlight the main advantage of the CTMC reformulation. In the PEER-PBEE framework, hazard, fragility, and loss are evaluated separately for each event. In contrast, the CTMC approach models damage and recovery together as a continuous process over

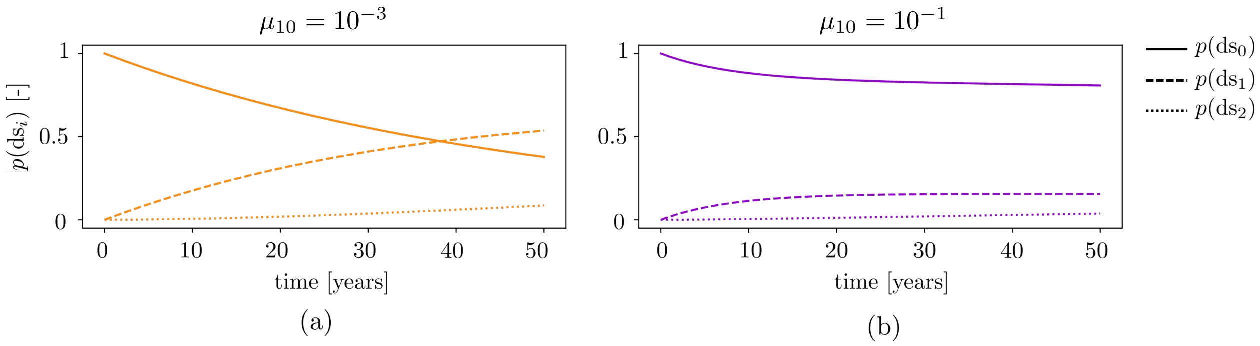

Figure 2. Time-dependent occupancy probabilities for transient states (0, 1) and the absorbing collapse state (2) over a 50-year horizon, comparing (a) negligible recovery ($\mu_{10} = 10^{-3}$), and (b) normal recovery ($\mu_{10} = 10^{-1}$).

time, so that the proportion of time spent in each damage state reflects the combined effect of these competing mechanisms. This shifts the focus from the probability of exceeding a damage threshold in a single event to the expected fraction of time the structure spends in each performance state over its lifetime, which is a more informative and practical measure for lifecycle management, retrofit prioritization, and recovery planning.

3.2. Verification and Performance Metrics

Next, we evaluate the reliability and resilience of the system using the metrics derived in Section 2: the first-passage probability to collapse $F_{\mathcal{T}}(T_{\mathrm{hor}})$ (Eq. (7)) and the expected damaged fraction of lifetime $\mathcal{R}_0^{\mathrm{unc}}$ (Eq. (12)). Analytically, both quantities are computed from the sub-generator $\boldsymbol{Q}_T$ of the transient states via matrix exponentiation and the truncated fundamental matrix $\boldsymbol{N}(T_{\mathrm{hor}})$, as described in Section 2. To verify these predictions, we generate $N_{\mathrm{sim}} = 10^5$ independent sample trajectories of the CTMC using the full generator $\boldsymbol{Q}$: holding times are drawn as $\Delta t \sim \mathrm{Exp}(-Q_{ii})$ and the next state j is sampled with probability $Q_{ij}/(-Q_{ii})$, censoring each path at T_{hor}. A trajectory contributes to the MCS estimate $\hat{F}(T_{\mathrm{hor}})$ if it reaches the absorbing state within the horizon, while $\mathcal{R}_0$ is estimated as the average fraction of $[0, T_{\mathrm{hor}}]$ over which damage accumulates across all trajectories. The agreement between analytical and MCS results confirms the correctness of both the theoretical expressions and their numerical implementation.

Figure 3 shows near-perfect agreement between the analytical predictions and the Monte Carlo estimates across the full range of recovery rates μ_{10}. The results confirm that $F_{\mathcal{T}}$ and $\mathcal{R}_0^{\mathrm{unc}}$ are smooth, monotone functions of μ_{10}, and that even a modest increase in recovery rate produces a substantial reduction in collapse probability and a marked increase in the fraction of time spent in the operational state. This sensitivity is not a numerical artifact: it reflects the physical competition between damage and recovery encoded in $\boldsymbol{Q}$, and it is precisely what the CTMC paradigm makes quantitatively accessible. In the classical PEER-PBEE setting, such a sensitivity curve cannot be so easily constructed, because recovery has no direct representation in the framework.

The convergence of the MC estimates is examined in Fig. 4. In the case of $\mu_{10} = 10^{-2}$/years, the running mean stabilizes well before the 10^5 trajectory limit and confidence interval consistently narrows, confirming the reliability of the simulation baseline. The analytical solution via $\boldsymbol{Q}_T$ reaches

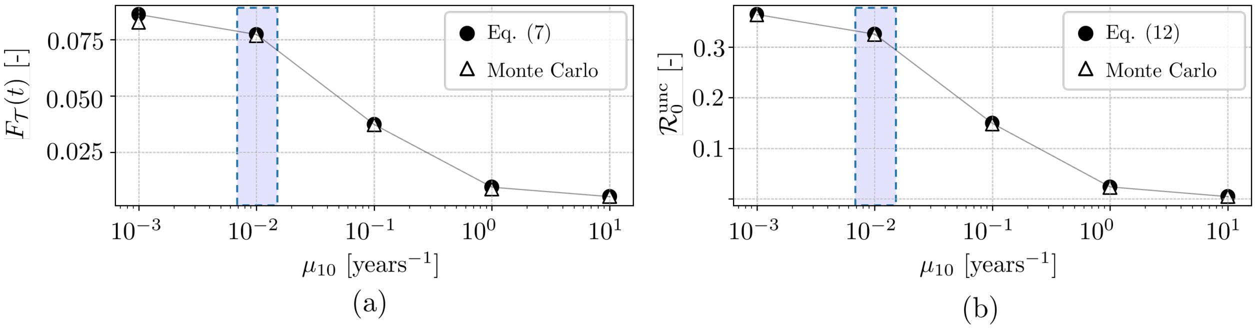

Figure 3. Comparison of (a) first-passage probability $F_{\mathcal{T}}(\cdot)$ and (b) resilience metric $\mathcal{R}_0^{\mathrm{unc}}$ calculated via analytical CTMC estimation and Monte Carlo simulation across a range of recovery rates. The colored box indicates the region zoomed in Fig. 4 to illustrate Monte Carlo convergence.

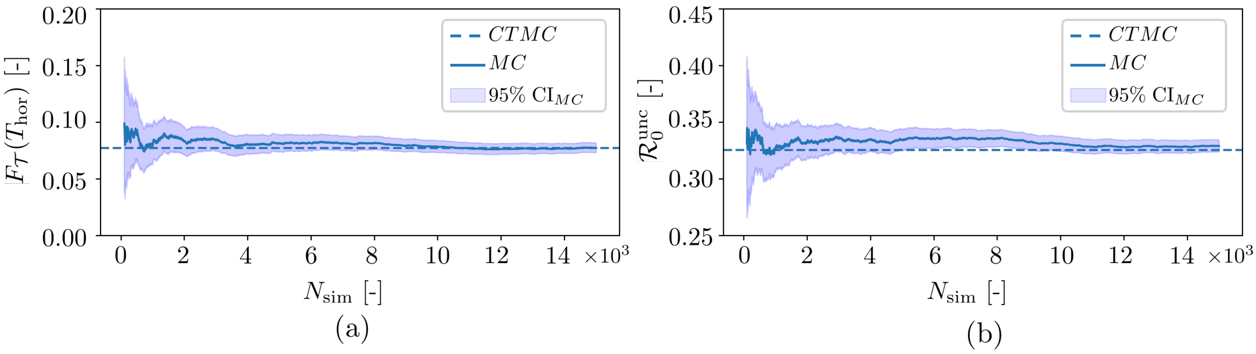

Figure 4. Monte Carlo convergence analysis showing the running mean and standard error for (a) the first-passage probability $F_{\mathcal{T}}(T_{\mathrm{hor}})$ and (b) resilience metric $\mathcal{R}_0^{\mathrm{unc}}$ as a function of the number of simulations N_{sim}.

the same result instantaneously, which is a practical bonus for long-horizon assessments. More importantly, both routes agree, lending credibility to the CTMC model itself as a representation of the underlying engineering problem.

Finally, Fig. 5 maps the system performance onto the β–ρ^{unc} plane as μ_{10} varies across several orders of magnitude. Grey-filled markers correspond to the specific values of μ_{10} adopted in Fig. 3, providing a direct link between the two figures. The green-shaded region identifies the high-performance quadrant, defined by $\beta \geq 2$ and $\rho \geq 2$, which represents configurations where the system is simultaneously highly reliable and highly resilient; this threshold can be used as a practical design target. The various points concretely illustrates how recovery capacity shapes overall performance: systems with negligible repair rates cluster in the low-reliability, low-resilience corner, while faster recovery pushes the system into the desirable quadrant.

One can also observe a positive correlation between β and ρ^{unc} along the trajectory. This is not incidental: by construction, ρ^{unc} is defined through $\mathcal{R}_0^{\mathrm{unc}}$, which accounts for the fraction of lifetime spent in damaged states across all trajectories, including those that reach collapse. For low values of μ_{10}, collapse events are more frequent, meaning a larger share of trajectories contributes damaged

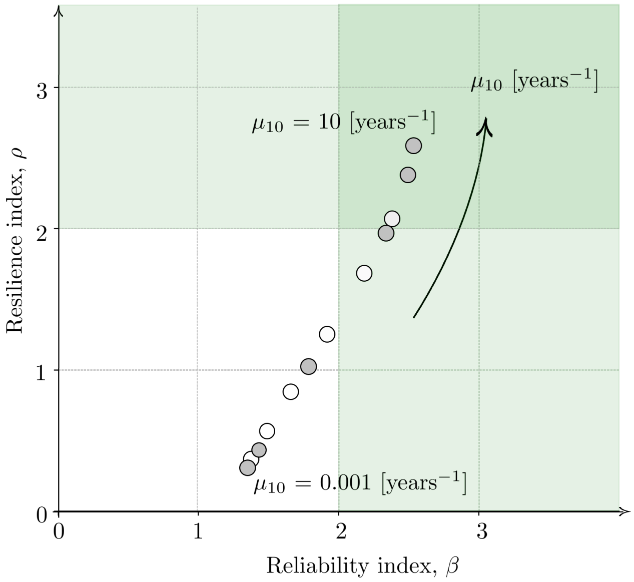

Figure 5. System performance trajectory in the resilience (ρ^{unc}) vs. reliability (β) domain, illustrating how increasing recovery rates move the system into the high-performance quadrant.

time up to absorption, which simultaneously depresses both β and ρ^{unc}. As recovery rates increase, fewer trajectories reach collapse, damage accumulation is curtailed, and both indices rise together.

Crucially, this entire mapmap is a direct output of the CTMC formulation and does not require re-running a full probabilistic analysis for each configuration: it naturally follows from varying a single entry of the generator matrix $\boldsymbol{Q}$. This makes the β–ρ^{unc} map a practical design and decision tool, allowing an engineer to read off, for a target reliability or resilience level, the minimum recovery rate that must be achieved through repair policies or structural intervention.

4. Conclusions

This paper presents a generalization of the PEER-PBEE framework that addresses two fundamental limitations in contemporary seismic risk assessment: the implicit Poissonian renewal hypothesis and the exclusion of recovery dynamics from the structural evolution. By adopting a continuous-time Markov chain (CTMC) approach, we have established a framework that naturally integrates hazard-driven damage accumulation and restoration processes under a unified stochastic engine.

Several key findings emerge from this work. First, we demonstrate that structural reliability and resilience are not merely secondary outputs, but are directly encoded within the system's infinitesimal generator matrix $\mathbf{Q}$. This mathematical coupling allows for the simultaneous evaluation of state-dependent deterioration and recovery mechanisms. Such a formulation provides a rigorous basis for comparing alternative structural configurations, retrofitting strategies, and post-disaster recovery policies within a single, consistent probabilistic environment.

Second, the analytical solutions derived for first-passage times and state occupancy were validated through extensive Monte Carlo simulations. The convergence between these methods highlights the computational efficiency of the CTMC approach; by utilizing analytical matrix exponentials, the framework enables long-horizon performance assessments that would otherwise be computationally prohibitive or statistically unstable using simulation alone.

The elegant, modular nature of this framework paves the way for several future research directions. The current time-invariant hazard assumption can be extended to incorporate time-variant phenomena, such as seismic aftershock sequences or climate-induced degradation. Furthermore, the rate-based formulation is flexible enough to accommodate diverse recovery models, ranging from linear restoration to complex, resource-constrained logistical functions.

Acknowledgements

This research is supported by the Marie-Sklodowska Curie program and the REACTIS project, GA no. 101147351. Views and opinions expressed are however those of the authors only and do not necessarily reflect those of the European Union, or REA or any sponsor. Neither the European Union nor the granting authority can be held responsible for them.

References

Andriotis, C. P. and K. G. Papakonstantinou: 2018, 'Extended and generalized fragility functions'. *Journal of Engineering Mechanics* **144**(9).

Applied Technology Council: 2021, 'ATC-138: Seismic Performance Assessment of Buildings. Volume 8: Methodology for Assessment of Functional Recovery Time'. Federal Emergency Management Agency (FEMA). `https://femap58.atcouncil.org/documents/fema-p-58/34-atc-138-3-volume-8-methodology-for-assessment-of-functional-recovery-time/file`.

Brémaud, P.: 2020, *Probability Theory and Stochastic Processes*, Universitext. Cham: Springer Nature Switzerland.

Cornell, C. A. and H. Krawinkler: 2000, 'Progress and Challenges in Seismic Performance Assessment'. PEER Center News. Spring issue.

Der Kiureghian, A.: 2005, 'Non-ergodicity and PEER's framework formula'. *Earthquake Engineering & Structural Dynamics* **34**(13), 1643–1652.

Ditlevsen, O.: 1979, 'Generalized Second Moment Reliability Index'. *Journal of Structural Mechanics* **7**, 435–451.

Federal Emergency Management Agency (FEMA): 2018a, 'FEMA P-58-1: Seismic Performance Assessment of Buildings, Volume 1 – Methodology'. `https://femap58.atcouncil.org/documents/fema-p-58/24-fema-p-58-volume-1-methodology-second-edition/file`.

Federal Emergency Management Agency (FEMA): 2018b, 'FEMA P-58-2: Seismic Performance Assessment of Buildings, Volume 2 – Implementation Guide, Appendix C: PACT Software'. `https://femap58.atcouncil.org/documents/fema-p-58/25-fema-p-58-volume-2-implementation-second-edition/file`.

Federal Emergency Management Agency (FEMA): 2020, 'HAZUS Earthquake Model 4.2.'. `https://www.fema.gov/sites/default/files/2020-10/fema_hazus_earthquake_technical_manual_4-2.pdf`.

Guida, M. and G. Pulcini: 2011, '"A Continuous-State Markov Model for Age- and State-Dependent Degradation Processes"'. *Structural Safety* **33**(6), 354–366.

Iervolino, I., M. Giorgio, and E. Chioccarelli: 2015a, '"Age- and State-Dependent Seismic Reliability of Structures"'. In: *12th International Conference on Applications of Statistics and Probability in Civil Engineering (ICASP12)*. Vancouver, Canada.

Iervolino, I., M. Giorgio, and E. Chioccarelli: 2015b, '"Markov Chain Modeling of Seismic Damage Accumulation"'. *Earthquake Engineering & Structural Dynamics* **44**(7), 1015–1032.

Moehle, J. and G. G. Deierlein: 2004, 'A Framework Methodology for Performance-Based Earthquake Engineering'. In: *Proceedings of the 13th World Conference on Earthquake Engineering*. Vancouver, Canada.

Nardin, C., S. Marelli, O. S. Bursi, B. Sudret, and M. Broccardo: 2025, '"UQ state-dependent framework for seismic fragility assessment of industrial components"'. *Reliability Engineering & System Safety* p. 111067.

Nardin, C., S. Marelli, B. Sudret, and M. Broccardo: 2026, '"A Generalized Framework for Performance-Based Earthquake Engineering: Integrated Seismic Risk and Resilience Assessment"'. *"In preparation"* pp. –.

Ross, S. M.: 2014, *"Introduction to Probability Models"*. Boston: Academic Press, 11th edition.

Trivedi, K. S. and A. Bobbio: 2017, *"Reliability and Availability Engineering: Modeling, Analysis, and Applications"*. Cambridge: Cambridge University Press.