Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Exposure Uncertainty and its Impact on Regional Seismic Risk Assessment

AnnaElisa Huynh
Stanford University https://orcid.org/0009-0000-5392-5849
Meredith Lochhead
mlochhea@stanford.edu
Stanford University https://orcid.org/0009-0006-7460-7456
Adam Zsarnóczay
Stanford University https://orcid.org/0000-0001-6895-1417
Gregory Deierlein
Stanford University https://orcid.org/0000-0002-0067-3817

Research Article

Keywords: Regional Risk Assessment, Exposure Model, Building Inventory, Earthquake Engineering

Posted Date: November 25th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11865927/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: The authors declare no competing interests.

Exposure Uncertainty and its Impact on Regional Seismic Risk Assessment

AnnaElisa Huynh*, Meredith Lochhead*, Adam Zsarnóczay, and Gregory Deierlein

Department of Civil and Environmental Engineering, Stanford University, California
mlochhea@stanford.edu

ABSTRACT

High-resolution regional risk assessments can forecast the impact of natural hazards on the built environment, providing essential information for community mitigation and resilience initiatives. These assessments are underpinned by building inventories (i.e., exposure models), which are required for robust risk quantification. Developing detailed footprint-level inventories is becoming increasingly feasible with greater data availability and computational power. However, key challenges remain, as existing data sources often have partially or completely missing attributes, incomplete spatial coverage, outdated records, or disagreements between sources. Previous research has shown that the choice of data sources and synthesis methods can lead to large, systematic differences in the resulting inventory and the corresponding risk metrics under a scenario earthquake, but the uncertainty that arises between reasonable candidate exposure models has not been comprehensively quantified. To address this gap, this study applies a logic tree to quantify exposure model uncertainty. Using a case study of seismic risk in building inventories derived from nationally available data, we (1) systematically identify sources of uncertainty in asset-level exposure modeling, (2) discuss robust, repeatable methods to remove known biases in national data, and (3) quantify the remaining epistemic uncertainty resulting from reasonable exposure modeling decisions and compare it against other uncertainties in the risk framework. To do so, a suite of possible inventories is generated and evaluated under a scenario earthquake to capture how bias and uncertainty are distributed across the inventory. We find that, for regional building inventories using nationally available data, significant biases can be addressed through robust exposure modeling methods. Furthermore, the remaining exposure model uncertainty is non-negligible relative to other components of the risk framework, such as ground motion model selection, and is not distributed uniformly across inventory characteristics, downstream risk metrics, occupancy classes, or space. This paper aims to establish methods for quantifying uncertainty between candidate exposure models and motivate greater attention to footprint-level building inventory generation, given its significance in accurately and meaningfully understanding regional risk.

1 INTRODUCTION

Regional natural hazard risk assessments, composed of hazard, exposure, and vulnerability, can be used to anticipate community performance and prioritize mitigation efforts. Exposure data, which characterizes the buildings, infrastructure, or other assets at risk from natural hazards, plays a critical role in quantifying this risk. In state-of-practice regional risk models for buildings in a city or wider region, exposure data (i.e., building inventory data) is typically available at an aggregated scale, such as the census block level [1]. However, recent studies have shifted toward higher *spatial resolution*, often using building-level exposure, which characterizes individual asset locations and corresponding building features (e.g., occupancy class, year of construction, building material, number of residential units). Building features serve two main functions in regional risk assessments: (1) informing vulnerability model selection (e.g., structural system is used to select an appropriate fragility function), and (2) informing downstream risk metrics (e.g., the number of units and/or population are used to assess displacement and other human impacts). The building features required in a given study are therefore governed by the *typology resolution*, which is the level of granularity required for vulnerability model selection and downstream impact assessment [2]. As studies push towards higher spatial and/or typology resolution, exposure data requirements increase significantly.

While exposure data is a critical input to regional risk assessments, no single data source comprehensively describes exposure with sufficient detail and accuracy. Available datasets often have partially or completely missing attributes, incomplete spatial coverage, or outdated records, limiting their use. Additionally, exposure data is often only available at a spatial or typology resolution coarser than a given study requires. When available *exposure data* is insufficient for a given study, an *exposure model* is required. Exposure models use one or more exposure data sources and data processing techniques to meet the needs of a given study. By moving past available data, exposure models introduce

*These authors contributed equally to this work.

epistemic uncertainty into the risk framework. If a study shifts toward higher spatial and/or typology resolution without a corresponding increase in available exposure data, epistemic uncertainty in the exposure model grows.

Different exposure models result from different input data sources or data synthesis and processing methods. Indeed, previous studies have found that the building inventories resulting from different exposure modeling decisions can lead to meaningful differences in the corresponding quantification of risk [2, 3, 4]. Lochhead et al. [2] provide a framework that seeks to standardize steps required for exposure modeling, alongside repeatable methods for data synthesis. However, even within this framework, significantly different exposure models can still arise due to input data source selection, modeling choices, gaps in the data, and other factors.

In the context of regional seismic risk, the three primary components of risk (hazard, exposure, vulnerability) are all characterized by large uncertainties. Here, the term *vulnerability model* is used as an umbrella term for fragility functions, loss functions, consequence functions, and other models that map hazard intensity to a downstream outcome. There are standard practices in modern seismic risk assessment models that address many of the uncertainties in the final risk estimates; however, other sources of uncertainty are often neglected, leading to an underestimation of uncertainty and potentially biased damage and loss estimates [5]. In this study, we consider two main types of uncertainty: (1) uncertainty within a chosen candidate model, which can include both aleatory and epistemic uncertainty, and (2) uncertainty between multiple candidate models, which is primarily epistemic uncertainty due to modeling decisions. There are different levels of maturity for the treatment of these two types of uncertainty across components of risk, as shown in Figure 1.

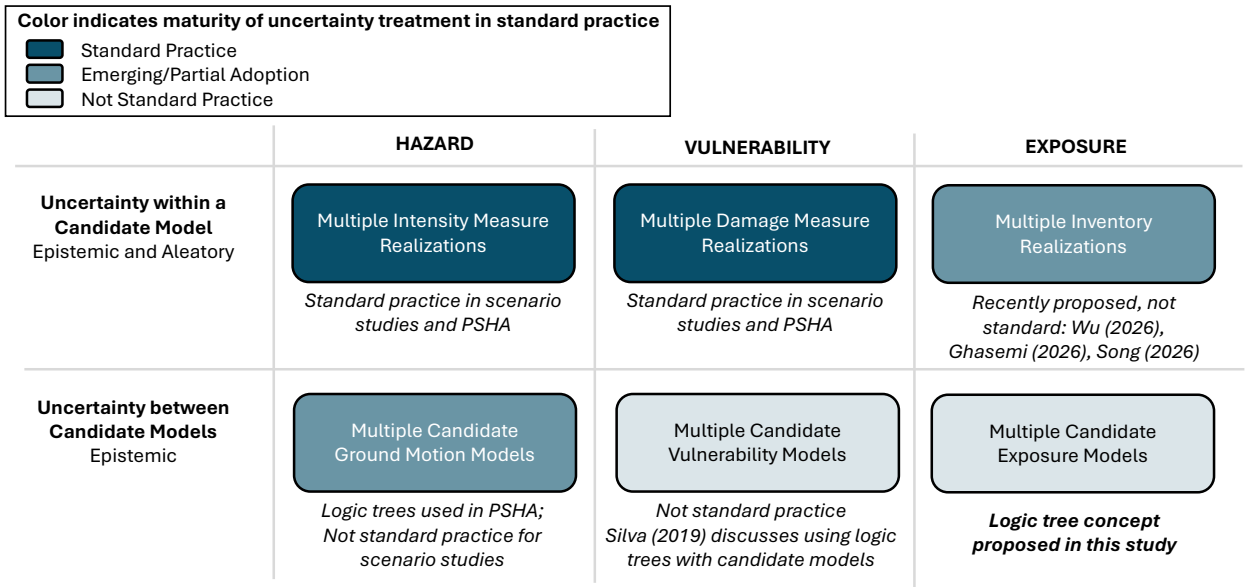


Fig. 1. Maturity of uncertainty quantification practice for the hazard, exposure, and vulnerability components of risk for both uncertainty within a candidate model and uncertainty between candidate models.

In regional assessments, there are significant uncertainties in the hazard component of risk, both in single-rupture earthquake scenario studies and in probabilistic seismic hazard analysis (PSHA). For a selected ground motion model, aleatory uncertainty in the intensity measure (IM) at a given site under a specified rupture can be high, and it is typically propagated via correlated IM realizations at each site [6]. There is additional epistemic uncertainty between multiple candidate ground motion models [7]. Epistemic uncertainty between ground motion models is commonly addressed and propagated in PSHA, including in the development of the United States National Seismic Hazard Model [8], through logic trees that incorporate appropriate candidate models. While widely used in PSHA, logic trees with candidate ground motion models are not standard practice for earthquake scenario studies.

For vulnerability, fragility functions encode uncertainty explicitly, capturing uncertainty in response under given intensity measures [9]. In regional risk assessments, fragility curves are sampled multiple times to account for this uncertainty [10]. Furthermore, loss and consequence calculations are also treated probabilistically to account for uncertainty, using methodologies such as FEMA P-58 [11]. Thus, it is generally standard practice in regional risk

assessments to sample vulnerability probabilistically to propagate this uncertainty. However, these sampling methods do not capture the second form of uncertainty, uncertainty between candidate models. Using multiple candidate vulnerability models is not standard practice. Silva et al. [12] state that including multiple fragility and consequence models within logic trees requires further attention and should become standard practice in the future, both to provide appropriate uncertainty bounds and to track improved data collection and modeling over time. This has not been widely adopted due to the limited number of available vulnerability models.

In this study, an *exposure model* refers to the selection of input data sources, data-processing techniques, and imputation and inference methods used to produce a final inventory. A single exposure model can generate multiple inventory realizations if any probabilistic step, such as filling in data gaps, is sampled multiple times using the same method. When underlying methods or data sources change, this study treats that as a different exposure model, separating the top and bottom rows shown in Figure 1. In current practice, neither type of uncertainty is accounted for in exposure models. Most regional studies treat exposure as deterministic, with minimal discussion of its source, expected accuracy, or the impact of assumptions [13]. This can limit stakeholders' ability to use regional risk assessment results for decision-making.

Several recent studies have explored exposure uncertainty in aggregated inventories for regional seismic risk assessment. Kalakonas et al. [5] quantify exposure uncertainty using alternate distributions of construction material and a range of spatial aggregation levels. Additionally, Gómez Zapata et al. [14] explore uncertain features in aggregated inventories (e.g., the proportion of building classes in an aggregated area), showing that when exposure composition is poorly characterized due to assumptions over aggregated data, its contribution to loss uncertainty may exceed that imposed by spatially distributed ground motion fields.

Several studies have also considered exposure uncertainty in higher resolution inventories (i.e., asset-level, rather than aggregated). However, these studies primarily focus on the uncertainty of multiple inventory realizations generated using a single exposure model (i.e., filling gaps probabilistically rather than deterministically), rather than uncertainty between candidate models (i.e., different input data sources or selected methods). In other words, recent studies have been primarily focused on the upper right-hand corner of Figure 1, not the lower right-hand corner. For example, Song et al. [15] identify the issue of partially missing features and the corresponding need for imputation. This study quantifies epistemic uncertainty due to features missing at random by imputing features multiple times. Other studies have moved beyond quantifying the impact of partially missing building features to address features that are completely missing (i.e., not available for any asset). This typically arises when a feature that is needed to inform vulnerability model selection is not available in a given data source. Ghasemi et al. [16] focus specifically on missing attributes in the National Structure Inventory (NSI), namely structural system; in this study, structural system is sampled from Hazus based on available building features. This study uses Monte Carlo sampling to generate complete building inventory realizations to propagate uncertainty in structural system sampling through a risk assessment, ultimately recommending a threshold number of samples that is appropriate for propagating structural system uncertainty. Wu and Burton [13] use the term *exposure information uncertainty* to describe the epistemic uncertainty that arises when imperfect imputations are used to infer missing attributes in exposure modeling. Using a case study of bridges, they explore two related concepts: (1) additional epistemic uncertainty due to imputing gaps with multiple realizations, relative to baseline uncertainty of epistemic and aleatory non-exposure uncertainties, and (2) bias due to misrepresentation in exposure. Wu and Burton [13] also discuss uncertainty-informed inspections of individual assets to reduce the impact of exposure information uncertainty in assessing risk. Together, these studies take a significant step forward in quantifying uncertainty via multiple inventory realizations, which captures the impact of filling gaps probabilistically rather than deterministically. Findings suggest that exposure characterization should not be treated deterministically, particularly when specific attributes are unavailable, and that incorporating exposure uncertainty is essential for more reliable risk assessments.

In the context of asset-level inventories, gaps in the data are the primary source of exposure uncertainty explored in the aforementioned literature; however, they are not the only source. As shown in Figure 1, there are also broader epistemic uncertainties between candidate exposure models that remain unaddressed. Unlike vulnerability, where the use of multiple candidate models is limited by the availability of robust models and the high effort required to develop new ones, exposure offers an opportunity to quantify uncertainty that can arise between candidate models.

Previous literature on asset-level inventories has generally not explored uncertainty due to multiple candidate exposure models. Several studies have shown that the choice of input data sources dramatically affects the resulting seismic risk [2, 4]; however, there have been few efforts to systematically quantify or propagate this uncertainty. There also may be uncertainty in the number and/or location of assets, which is also not typically addressed in the literature. For example, this topic is not included in Ghasemi et al. [16] because they focus on census blocks with equivalent numbers of buildings across data sources, thus filtering out blocks with uncertainty in asset count and location. More

often, studies select a single authoritative source for asset location to avoid this uncertainty, as Wu and Burton [13] does with the National Bridge Inventory. In addition, the previous literature does not go in depth on the uncertainty and bias due to inaccuracies in existing data. Most assume existing exposure data is fully accurate and focus only on data gaps; however, systematic inaccuracies can exist depending on data provenance, which can significantly bias results [2]. Finally, the previous studies on asset-level uncertainty quantification focus on building features that inform vulnerability model selection. While this is critical, there are other building features that can become important for downstream impact assessment (e.g., replacement value, number of residential units), and it is important to consider uncertainty in those features as well.

This study applies the concept of logic trees to quantify uncertainty between candidate exposure models for regional seismic risk assessment. Using a building stock characterized from nationally available data, this study makes the following contributions. First, this study systematically identifies sources of uncertainty in asset-level exposure modeling and quantifies the epistemic uncertainty resulting from reasonable exposure modeling decisions in terms of both inventory characteristics and seismic risk estimates. To do so, a suite of possible inventories is generated and evaluated under a scenario earthquake to better understand how bias and uncertainty are distributed across the inventory. Second, this study discusses repeatable methods to remove known biases in national data and demonstrates that these methods are robust across the residual (remaining) uncertainty in the inventory. Without resorting to uncertainty propagation or on-the-ground inspections, these are methods modelers can use to reduce the bias in their inventory generation process. Third, the total exposure model uncertainty, as well as the residual uncertainty once known biases have been resolved, are quantified and compared against the uncertainty due to other risk modeling choices (i.e., ground motion model selection) to assess whether exposure uncertainty is meaningful to quantify.

2 EXPOSURE UNCERTAINTY

In this study, an *exposure model* is defined as a suite of exposure modeling decisions, including selecting input data sources, defining assets, and selecting methods for data attribution, feature prioritization, imputation, and inference. A single exposure model (defined as a suite of decisions) can produce multiple inventories if any step involves probabilistic assignment (e.g., assigning structure type probabilistically using a pre-defined methodology). Multiple candidate exposure models can produce different inventories due to different input data or methodological decisions.

To be useful, an exposure model must produce an inventory with sufficient information to support all downstream modeling for a given study, including both fragility/vulnerability model selection and downstream impact assessment, at the required spatial resolution. While the explicit data requirements of each exposure model vary, the methods, challenges, and steps required to develop such models are similar, as discussed in Lochhead et al. [2]. There are also similarities in the sources of exposure model uncertainty and the way it can be defined and classified.

2.1 Sources of Exposure Model Uncertainty

As shown in Figure 1, exposure models have uncertainty within a candidate model as well as between candidate models. This study focuses primarily on the latter. Figure 2 outlines several common exposure modeling decisions that can become sources of uncertainty between candidate exposure models, identified based on existing literature and the authors' own experiences with exposure modeling. These decisions fall into two categories: *spatial modeling decisions*, which concern the spatial characteristics of the inventory (e.g., data sources, asset definition, and data attribution), and *feature modeling decisions*, which determine how each building feature is assigned.

Uncertainty due to *data source selection* (Figure 2a) refers to how the use of different input data sources can contribute to large differences in the resulting inventory and risk. For example, NSI data [17] will produce different results than tax parcel data [2]. Thus, selecting one or more input data sources is a decision that contributes to exposure model uncertainty.

In most exposure models, there is a single baseline geometry that is used to define the asset location and count. Thus, *asset definition*, or the process of selecting and defining that baseline geometry, can also create exposure model uncertainty, as the geometry varies between different candidate baseline geometry sources (Figure 2b). For example, in developing a footprint-level building inventory, asset definition will depend upon the selection of the building footprint source (e.g., OpenStreetMap versus Overture Building Footprints). Additional variation can arise given data processing (e.g., filtering out small building footprints) or data merging across multiple candidate baseline geometry sources.

Uncertainty can also arise due to *data attribution*, where all selected input data sources are attributed to the selected baseline geometry (Figure 2c). The need for methods here is due to inexact geometries and lack of one-to-one mapping across different input data sources. For example, NSI points do not map cleanly to footprint geometries [2]. Multiple candidate methods for data attribution produce different results, thus contributing to exposure model uncertainty.

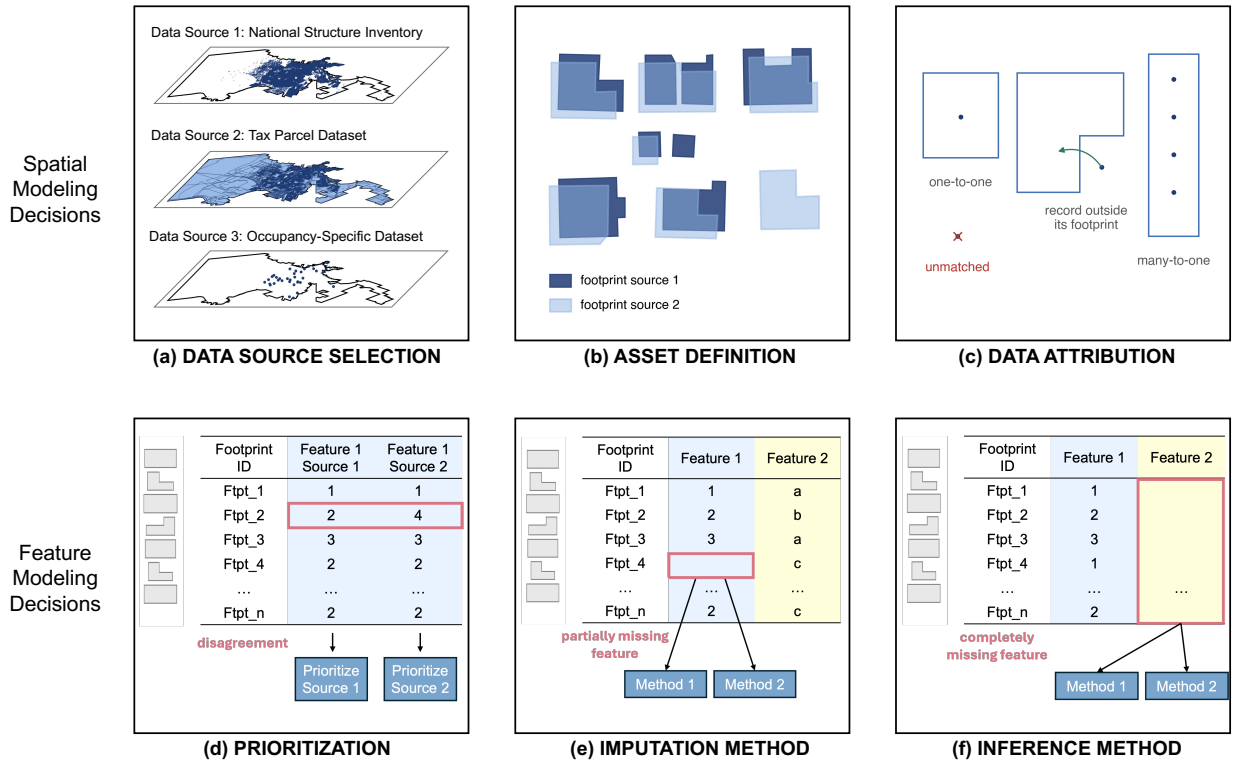


Fig. 2. Identified sources of uncertainty between candidate exposure models (lower right-hand corner of Figure 1), including both spatial modeling decisions and feature modeling decisions.

There are also a series of feature modeling decisions, including methods for prioritization, imputation, and inference, which can be applied independently to different building features in the inventory, as appropriate. *Prioritization* (Figure 2d) refers to decisions where there are multiple data sources or methods that are used to describe a single building feature (e.g., replacement value), and the disagreement between them is resolved by prioritizing one over the other(s). *Imputation* (Figure 2e) refers to the process of filling in partially missing building features using statistical methods based on patterns observed in existing data for that and other features [2]; here, the uncertainty does not come from sampling multiple times using a single imputation method, but instead from multiple possible imputation methods. Finally, *inference* (Figure 2f) refers to the process of filling in gaps in building features through the use of external models or reasonable assumptions based on available building features. There may be multiple possible ways to infer a given feature (e.g., selecting a structural system based on available data). These feature-level decisions are important both for building features that are used to select a vulnerability model, as well as those used to inform other downstream modeling efforts.

2.2 Logic Tree Framework for Exposure Uncertainty and Bias

Multiple candidate methods can be selected for each of the six categories shown in Figure 2, and different combinations of these methods produce distinct downstream inventories. Thus, similarly to the ground motion models, exposure model uncertainty can be quantified using a logic tree, where each exposure modeling decision represents a branch split. Thus, if all reasonable candidate data sources and methods are included as branches in the logic tree for inventory generation in a given region, the *total exposure model uncertainty* could be quantified via the uncertainty in resulting inventories. Likewise, we can quantify the resulting impact on seismic risk metrics by propagating each resulting inventory through a scenario earthquake assessment.

What constitutes *reasonable* candidate data sources and methods, however, is not simple. In this study, reasonable refers to exposure modeling decisions that would not raise significant red flags in the peer review process if used in a standalone regional risk assessment study. However, given that definition, not all reasonable decisions are necessarily recommended. Lochhead et al. [2] and others have identified systematic inaccuracies in existing datasets, and there

are commonly used methods that can lead to significant bias in the downstream inventory. While decisions that do not account for these biases may be considered reasonable under the specified definition (i.e., use the raw NSI data), they are not necessarily recommended given currently available evidence. Recognizing these data biases is important, and several examples of how to do so are listed here.

- **Compare against an external data source.** A data source with significant differences relative to an external data source may indicate bias. For example, Lochhead et al. [2] demonstrate that the NSI lists significantly fewer mobile homes than the tax parcel data for a case study city, indicating bias in the NSI data.
- **Leverage previously-identified biases from appropriate reference locations.** Not all locations have extensive external data by which to evaluate bias; however, it is reasonable to assume that biases identified in a given data source may exist across similar case study locations. For example, if there is significant bias in the mobile home count in the NSI for the case study in Lochhead et al. [2], it is likely that bias may exist in other locations as well.
- **Spot check with widely-available data.** Bias can also be identified by spot checking against widely-available data, such as Google Maps. Satellite and street-level imagery can both be powerful tools in identifying potential biases in existing inventory data.
- **Use engineering judgment.** In the absence of external data, engineering judgment can also be used to identify possible biases. For example, if a city has a mandatory retrofit ordinance for unreinforced masonry buildings that has been completed, there should not be unretrofit unreinforced masonry buildings in the inventory.

Identifying bias in exposure data is the first step in resolving it; the second is to assess possible ways to resolve this bias and ensure the correction is robust across other exposure modeling decisions. If a certain data source or method consistently reduces the bias regardless of other exposure modeling decisions, that branch of the logic tree is better-supported by currently available evidence than its alternative branches. By collapsing the logic tree to the better-supported branch and removing the other(s), the number of exposure models in the logic tree is reduced. This, in effect, reduces exposure uncertainty by removing the component of total exposure model uncertainty resulting from known biases. This study refers to the component of total exposure model uncertainty resulting from such decisions as *resolvable bias*.

Ideally, all exposure modeling decisions would have sufficient evidence to determine which branch is better supported (i.e., which branch reduces bias); if so, all decisions in the logic tree would collapse to a single exposure model. However, not all decisions have sufficient evidence or sufficiently robust methods to categorize them as resolvable bias, so a smaller logic tree remains after removing all branches carrying resolvable biases. The exposure model uncertainty captured by the remaining logic tree (i.e., after collapsing all resolvable biases to their preferred alternative) is referred to in this study as *residual uncertainty*. Several reasons may keep a decision in the category of residual uncertainty rather than resolvable bias, meaning that both branches stemming from that decision remain in the logic tree. First, there may be multiple data sources and/or methods that appear equally defensible given currently available evidence, meaning that there is no clearly preferred alternative. Second, it is possible that a bias is identified using the above methods, but that no robust branch exists to resolve it given the currently available data.

Total exposure model uncertainty thus comprises two components: resolvable bias, which can be collapsed given current evidence, and residual uncertainty, which remains. Unlike total exposure model uncertainty, which should be treated as a diagnostic quantity, this study recommends propagating residual uncertainty through regional risk assessments, as it represents remaining exposure modeling decisions with multiple defensible alternatives. This allows us to quantify the impacts of exposure model uncertainty on downstream risk metrics. Table 1 provides definitions for the terms introduced in this section.

3 CASE STUDY

To explore exposure uncertainty, this study adopts the inventory synthesis framework from Lochhead et al. [2]. This framework provides a general set of steps for synthesizing multiple input data sources into an asset-level inventory for a regional risk assessment. Lochhead et al. [2] presents three examples of inventory generation using the same general framework, including examples using only locally collected data, using only nationally available data, and a 'best estimate' using all available data. The National Synthesis Workflow (i.e., the example using only nationally available data) is shown to produce a footprint-level inventory that more closely approximates the seismic risk of an inventory generated using high-quality local data than reliance on any one nationally available data source alone. It can be used with minimal changes in any region in the United States.

TABLE 1. Definitions of Terms Used to Describe Exposure Model Uncertainty

	Definition
Total Exposure Model Uncertainty	The spread in exposure and downstream risk metrics arising from the range of reasonable exposure modeling decisions, which can be represented using a logic tree. This is diagnostic (not probabilistic); it indicates the extent to which reasonable decisions in inventory generation can change the results of a regional seismic risk assessment, and is not itself recommended for propagation as a probabilistic quantity.
Resolvable Bias	The component of the total exposure model uncertainty contributed by decisions for which currently available evidence implies one branch to be systematically more accurate than its alternatives. Resolving such a decision means collapsing the logic tree to the better-supported branch, thus reducing the uncertainty captured by the logic tree.
Residual Uncertainty	The component of the total exposure model uncertainty that remains once all decisions associated with resolvable bias have been collapsed to their preferred alternative. This is the epistemic uncertainty carried into subsequent risk assessments and treated probabilistically. The resulting distribution is not guaranteed to be centered on the true exposure, since bias likely remains from decisions that current evidence cannot resolve.

Since the synthesis framework from Lochhead et al. [2] was originally developed, an updated version of the National Structure Inventory, a key input data source, has been released (NSI 2022 versus NSI 2026). Changes observed in the data include an improvement in the characterization of masonry residential housing in Northern California (which was dramatically over-represented in NSI 2022), a notable increase in the replacement value, and a reduction in stray points. These improvements improve data quality, but the inventory produced by the National Synthesis Workflow using the updated NSI 2026 data still offers clear benefits over using the raw NSI 2026 data.

In this study, we adopt a modified version of the National Synthesis Workflow to evaluate exposure uncertainty. This workflow is selected to make results applicable to other areas by reducing reliance on local data sources and/or on-the-ground inspections to improve accuracy, which are not always available. We identified 20 exposure modeling decisions throughout the National Synthesis Workflow where user choices can lead to different outcomes. Each decision has between two and five ways to proceed, introducing uncertainty in the resulting inventory. The impact of the 20 decisions is used as a proxy for total exposure model uncertainty. The individual details of all decisions are provided in the supplemental materials, separated into spatial and feature modeling decisions, per Figure 2. The naming convention for these decisions, introduced in Table S1 in the supplemental materials, is used throughout this study. Figure 3 shows an overview of the National Synthesis Workflow from Lochhead et al. [2], and decisions included in the exposure uncertainty calculation are marked appropriately.

We apply the workflow to the case study of Hayward, California, which lies along the Hayward Fault in the San Francisco Bay Area and consists primarily of low-rise residential buildings. While this study focuses on exposure models using national data, Hayward's robust local data is adopted as a reference point against which to identify potential biases and evaluate robustness. Other than its use as an external reference, the local data is not used in any of the inventories in this study. More information on all relevant data sources can be found in Lochhead et al. [2].

To evaluate the impact of exposure uncertainty on downstream risk metrics, this study adopts a scenario earthquake on the Hayward Fault. Specifically, a M6.95 earthquake on the Hayward-Rodgers Creek rupture is obtained from OpenSHA UCERF2 and used in the SimCenter R2D Earthquake Event Generation Tool to generate 50 intensity measure (IM) realizations. The realizations use the ground motion model of Boore et al. [18], the V_{S30} map of Thompson [19], and the spatial correlation model of Jayaram and Baker [6]. Ground failures are not considered. Vulnerability is modeled with PGA-based Hazus fragility functions from the SimCenter Damage and Loss Model Library [20]. The uncertainty in each building's damage state and repair cost is characterized by running 1000 damage and loss simulations using the SimCenter's Pelicun tool [21].

4 QUANTIFYING EXPOSURE MODEL UNCERTAINTY

The goal of this study is to quantify the total exposure model uncertainty, resolvable bias, and residual uncertainty for a case study application. To do so, this study follows an inventory generation and evaluation approach similar to the studies discussed previously: generate multiple inventories, propagate each inventory through a scenario earthquake,

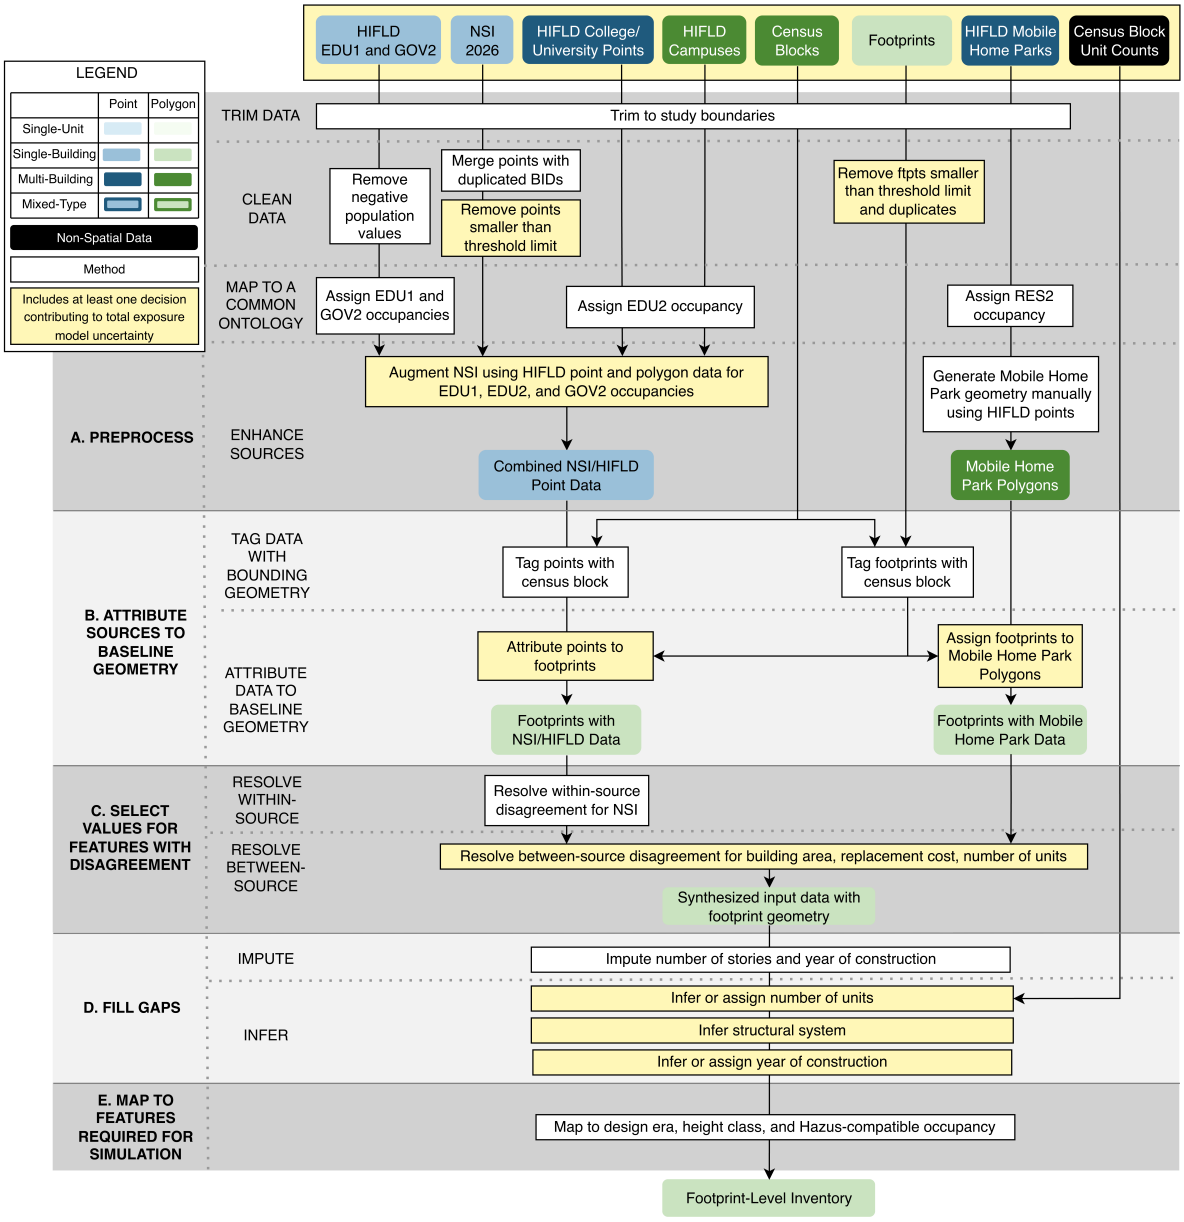


Fig. 3. Modified National Synthesis Workflow from Lochhead et al. (2026) with decision points highlighted in yellow.

and evaluate statistics across the inventories and resulting seismic risk metrics. Within this overall setup, total exposure model uncertainty, resolvable bias, and residual uncertainty are each evaluated using related but distinct methods, as shown in Figure 4.

4.1 Total Exposure Model Uncertainty

Total exposure model uncertainty is the spread resulting from all exposure modeling decisions included in the logic tree. The method used to quantify it is shown in Figure 4. For the Hayward case study, the logic tree contains 20 decisions, which are fully described in the supplemental materials. Because each decision has 2 or more possible outcomes, exhaustively generating every combination of decisions would produce nearly 9 million inventories. Generating this number of inventories and propagating each through a scenario earthquake is computationally prohibitive. Instead, 10,000 exposure models are developed using Monte Carlo sampling of the logic tree, with equal weights assigned to all branches. Each exposure model is then used to generate one inventory. Because a finite sample is

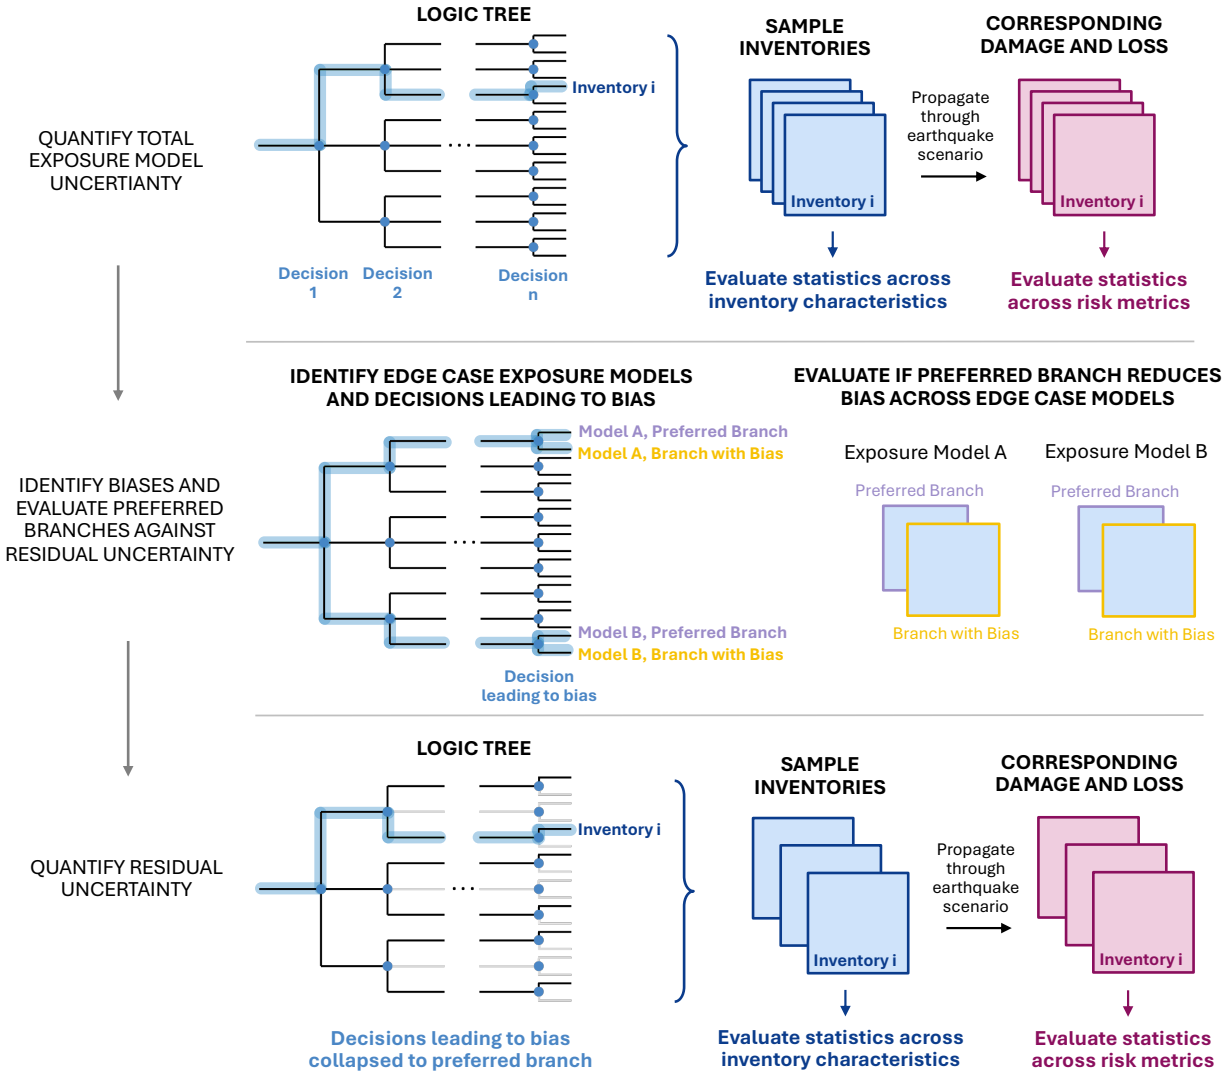


Fig. 4. Overview of uncertainty quantification methods for total exposure model uncertainty, resolvable bias, and residual uncertainty.

unlikely to capture the most extreme branch combinations, the resulting spread may slightly underestimate the total exposure model uncertainty; however, it is considered an appropriate approximation for this study.

All decisions in the logic tree contribute to uncertainty *between* candidate exposure models, and each exposure model is realized only once. Thus, any probabilistic step within an exposure model (e.g., sampling structural system or imputing gaps) is also sampled only once per branch. Because these steps are sampled independently across branches, uncertainty from repeated sampling *within* a single exposure model is assumed to be captured implicitly in the spread across the 10,000 inventories.

4.2 Resolvable Bias

A related but distinct analysis demonstrates robust methods for resolving bias (Figure 4). Bias is resolvable when currently available evidence implies one branch is systematically more accurate than its alternatives. Thus, to show that a bias is resolvable, the analysis must demonstrate that the preferred choice produces more accurate results regardless of the other exposure modeling decisions. This could be done by exhaustively generating all inventories with the bias-resolving choice and all inventories without it, and comparing both sets against a reference inventory. However, as with the total exposure model uncertainty, this is computationally demanding. Furthermore, the goal here is not

to capture the resulting spread in inventories, but to demonstrate that one branch is preferred across the remaining decisions that constitute the residual uncertainty.

Instead of exhaustive generation, this study evaluates two exposure models, each defined as a suite of decisions, intentionally developed to represent significant differences in the remaining decisions (i.e., edge-case exposure models). It is assumed that if one branch resolves the bias in both exposure models, it is robust across all residual uncertainty. For the Hayward case study, the two exposure models are referred to as Exposure Model A, which selects options informed by the original NSI data when possible, and Exposure Model B, which selects options that are not informed by the original NSI data. Full descriptions of both models are provided in Section S1.2 in the supplemental materials.

4.3 Residual Uncertainty

Once resolvable biases have been identified and the corresponding bias reduction methods are shown to be robust, those decisions are collapsed to their preferred alternatives, and all remaining decisions in the logic tree contribute to residual uncertainty. The method for quantifying residual uncertainty closely follows that used for total exposure model uncertainty, meaning inventories are generated, propagated through a scenario earthquake, and uncertainty is evaluated using statistics on the outputs. However, unlike total exposure model uncertainty, which is a diagnostic quantity, residual uncertainty is the between-model uncertainty that this study recommends propagating through a regional earthquake scenario assessment. Collapsing the resolvable-bias decisions reduces the logic tree to only 6,144 exposure models, which is feasible to sample comprehensively. Therefore, all combinations are used to generate exposure models, each of which is realized once and weighted equally.

4.4 Risk Metrics for Measuring Uncertainty

Several downstream risk metrics are computed for each inventory to quantify the impact of exposure uncertainty on seismic risk. These include the total residential repair cost (see Equation 1) and the number of residential buildings (see Equation 2) and residential units (see Equation 3) with a specified percentage of damage realizations meeting or exceeding a threshold damage state. Repair cost is adopted because it is a common quantity in regional risk assessments and has direct financial implications. The selected damage metrics are adopted because separating damage from dollar value allows us to separate uncertainty in the value of the buildings from uncertainty that contributes to their seismic performance. Damage can also impact downstream models other than financial loss (i.e., recovery, displacement). We focus specifically on residential buildings and units as they have clear implications for human displacement and other community-based risk assessment metrics.

$$Repair\ Cost = \sum_{i=1}^{N_{bldgs}} RC_i \cdot \left[\frac{1}{N_{DL}} \sum_{j=1}^{N_{DL}} LR_{ij} \right] \tag{1}$$

where:

- RC_i = full structural replacement value of building i (in dollars). Throughout this study, replacement value represents the full (undepreciated) structure value, excluding content and vehicle value.
- LR_{ij} = loss ratio for building i in realization j (as a fraction of full structural replacement value, from 0 to 1)
- N_{bldgs} = total number of buildings in the inventory
- N_{DL} = total number of damage and loss realizations

$$RB_{DS,p} = \sum_{i=1}^{N_{bldgs}^{\text{RES}}} \mathbf{1}\left(D_{1-p/100,i} \geq DS\right) \tag{2}$$

$$RU_{DS,p} = \sum_{i=1}^{N_{bldgs}^{\text{RES}}} U_i \cdot \mathbf{1}\left(D_{1-p/100,i} \geq DS\right) \tag{3}$$

where:

- $RB_{DS,p}$ = number of residential buildings meeting or exceeding the damage state threshold (DS) in p% of damage realizations

- $RU_{DS,p}$ = number of residential units meeting or exceeding the damage state threshold (DS) in p% of damage realizations
- $D_{1-p/100,i}$ = highest damage state that building i meets or exceeds in at least p% of damage realizations. For example, when $p = 10$, this is the highest damage state reached or exceeded in at least 10% of realizations.
- N_{bldgs}^{RES} = total number of residential buildings in the inventory
- U_i = number of residential units in building i

More specifically, this study uses three combinations of exceedance probability thresholds and damage state thresholds as the primary evaluation metrics to assess exposure model uncertainty. Table 2 lists each metric and explains the rationale for its selection. A more comprehensive list of potential damage metrics considered for this study can be found in Section S1.4 in the supplemental materials.

TABLE 2. Selected Damage Metrics Used to Evaluate Exposure Uncertainty

Metric	Selection Rationale
$\mathbf{RB_{3,10}}$, $\mathbf{RU_{3,10}}$	**Residential buildings or units with a possibility of extensive damage.** Buildings or units have extensive damage (damage state 3 or higher) in at least 10% of damage realizations (at least 1 in 10 probability of extensive damage).
$\mathbf{RB_{3,20}}$, $\mathbf{RU_{3,20}}$	**Residential buildings or units with an elevated possibility of extensive damage.** Buildings or units have extensive damage (damage state 3 or higher) in at least 20% of damage realizations (at least 1 in 5 probability of extensive damage).
$\mathbf{RB_{2,80}}$, $\mathbf{RU_{2,80}}$	**Residential buildings or units that are likely damaged under the scenario earthquake.** Buildings or units have moderate damage (damage state 2 or higher) in at least 80% of damage realizations (at least 4 in 5 probability of damage).

5 RESULTS

This study defines 20 exposure modeling decisions within the National Synthesis Workflow (see Section S1.1 in the supplemental materials). Together, these decisions define the logic tree representing total exposure model uncertainty. Following the uncertainty quantification methods for total exposure model uncertainty shown in Figure 4, results show that the total number of buildings in Hayward varies between approximately 25,400 and 35,500. In other words, before introducing any uncertainty from the hazard or vulnerability components of risk, the number of buildings in the city can change by 30-40% due to reasonable alternative exposure modeling decisions, illustrating that regional study results depend heavily on these choices. This pattern holds true across other inventory metrics, as well as downstream risk metrics resulting from each inventory being passed through a scenario M7 earthquake. For example, the total number of residential units in Hayward ranges from approximately 32,400 to 59,600, a change of roughly 80% (see Figure 5a). Figure 5b shows the range of various inventory metrics, including the total number of buildings and number of units in the city, as well as the number of buildings and/or units for a given damage metric, as specified in Table 2. These results show that reasonable exposure modeling decisions can have large implications for the resulting inventories and seismic risk estimates.

Total exposure model uncertainty is intended as a diagnostic quantity of the spread of inventories resulting from reasonable exposure modeling decisions, not a metric recommended for propagation into regional seismic risk estimates. It is not recommended for propagation because although all decisions in the logic tree are considered reasonable, not all are necessarily recommended, as some decisions lead to bias. Once decisions causing bias are resolved in the logic tree, residual uncertainty should be propagated. Thus, the remainder of this section focuses primarily on the evaluation of methods to resolve bias, as well as the residual uncertainty and its drivers. Total exposure model uncertainty is further assessed in the Discussion section.

5.1 Resolvable Bias

This section evaluates which exposure modeling decisions can be resolved to better-supported branches given currently available data. Building on Lochhead et al. [2], comparisons with external data and engineering judgment are used to identify systematic biases associated with selected branches of the logic tree. When a bias is identified, Exposure Models A and B, which represent contrasting sets of other exposure modeling decisions (see Figure 4), are

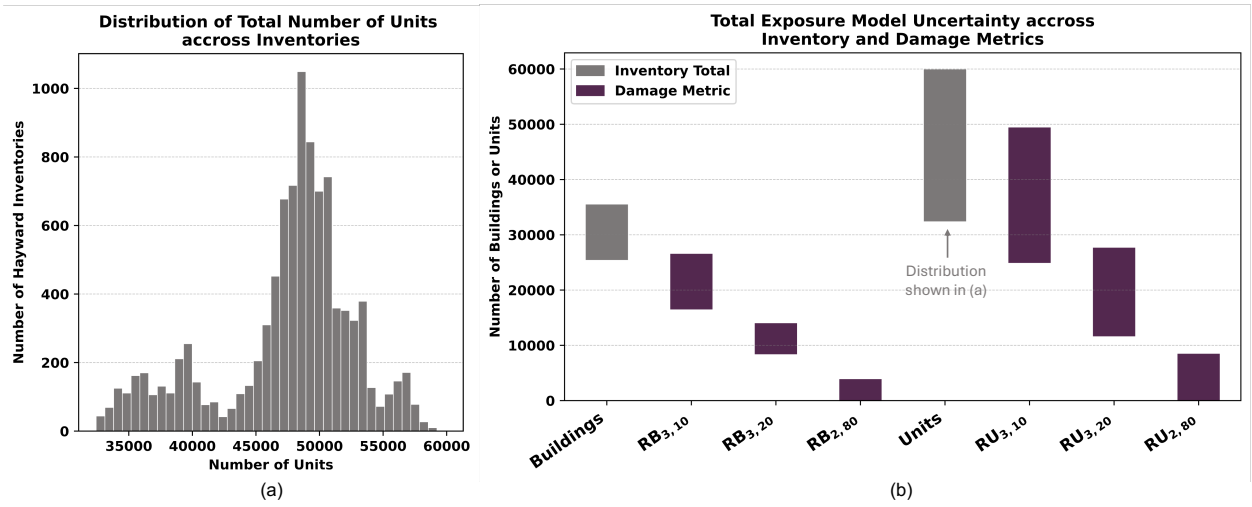


Fig. 5. Total exposure model uncertainty across metrics. (a) Distribution of total number of residential units across inventories. (b) Range of the number of residential buildings (RB) and residential units (RU) across various inventory and damage metrics. Damage metrics are derived according to Equations 2 and 3 and described in Table 2.

used to evaluate whether one branch systematically reduces the bias relative to its alternatives. In this case, we evaluate whether a bias is being reduced by comparing to a reference inventory that incorporates detailed local data. If one branch better approximates local data in both Exposure Models A and B, we assume that the alternative branches lead to bias. A consistent preference across both models provides evidence that the branch selection is robust, indicating that the associated decision in the logic tree should be collapsed to the better-supported branch in subsequent residual uncertainty analysis. The following results present the evidence of bias, the comparison of alternative branches to assess which reduces bias, the evaluation of whether that selection is robust across residual uncertainty, and the recommended change to the logic tree based on this evidence.

Asset Definition and Data Attribution (Decisions S1 and S6 in Table S1)

Two of the primary spatial modeling decisions, per Figure 2, are asset definition (e.g., selecting a building footprint source) and specifying a data attribution method to link input data sources to the footprints.

Candidate branches: We evaluate combinations of five commonly-used building footprint sources: Microsoft Footprints [22], OpenStreetMap [23], USA Structures [24], Overture Footprints [25], and a local Hayward database [26]. In addition, we also consider several data attribution methods, including a simple spatial join, where points that overlap with footprints are attributed to those footprints, a commonly used nearest-neighbor approach where points are assigned to their nearest footprint, and the enhanced spatial join method developed in Lochhead et al. [2], which aims to retain points during attribution if possible. Additional decisions relate to asset definition and data attribution, such as a lower-bound footprint area threshold (Decision S2) or upper-bound distance limits for data attribution (Decision S11). However, these decisions could not be resolved to a single better-supported branch (they remain in residual uncertainty), so they are not discussed here.

Evidence of bias: As shown in Figure 6a, the total building count varies by up to 36% (nearly 9,500 buildings missing in one inventory relative to another). The variation in the total number of buildings is largely driven by variable footprint coverage across individual sources, with a smaller contribution from the attribution method. Figure 7 shows that OpenStreetMaps is missing large quantities of footprints across the city of Hayward, with the same being true of other footprint sources to a lesser extent. Figure 7 represents the most recent update across all sources as of October 2026. Hayward, which is in the Bay Area, is generally a data-rich place; considering how much the inventory changes just based on footprint source selection, the same is likely true in other case studies, particularly less data-rich areas. Data attribution can also create bias, as shown in Figure 6b. In this study, population is taken from NSI data, with points being dropped if not attributed to footprints. The implemented simple spatial join and nearest neighbor method results in population being dropped relative to the enhanced spatial join method, which allows for multiple points to be absorbed into a single footprint. The enhanced spatial join retains population in the final inventory, resulting in a

better match with the 2020 Decennial Census population of Hayward. Furthermore, Figure 6 displays the results for Exposure Model B, but key patterns hold in both Exposure Models A and B, shown in the supplemental materials.

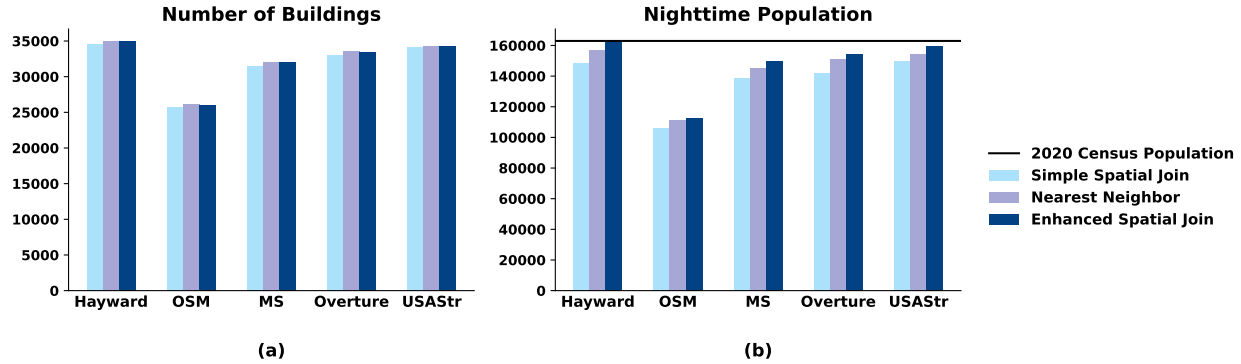


Fig. 6. Comparison of (a) the total number of buildings and (b) the total nighttime population for inventories generated using different combinations of footprint source and data attribution method. Population comparison against the 2020 US Census is also shown in (b).

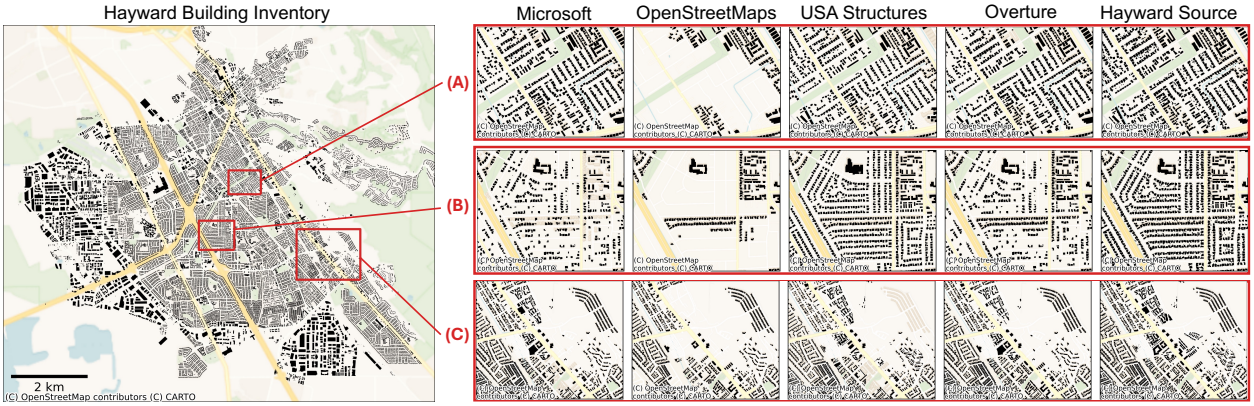


Fig. 7. Examples of differences in footprint coverage that contribute to differences in the total number of buildings in Figure 6a. Figure adapted from Lochhead et al. [2] with updated footprint data.

Branch selection and implications: Based on these results, the five footprint source branches and the three data attribution branches can each be resolved into a single branch that reduces bias. In this case, selection includes the Hayward building footprints and the enhanced spatial join method. In future inventory development applications in other cities, these findings point to two recommendations. First, prior to developing an inventory, it is beneficial to visually compare available footprint sources and select the source with the most comprehensive coverage to avoid large downstream biases; common data sources may still have significant gaps in footprint coverage. Second, the enhanced spatial join method proposed by Lochhead et al. [2] reduces bias relative to the other two data attribution methods considered in this study and is recommended for future use.

Mobile and Manufactured Homes (Decision F8 in Table S1)

Candidate branches: We compare exposure models with and without the manual mobile home correction proposed by Lochhead et al. [2], which uses manually drawn mobile-home-park polygons to assign appropriate occupancy classes, subsequently influencing the structural system designation. This correction is feasible to implement at the city or county scale. Both branches are included in the total exposure model uncertainty because use of uncorrected NSI data meets this study's definition of a reasonable modeling choice.

Evidence of bias: Lochhead et al. [2] identify systematic undercounting of mobile and manufactured homes in the 2022 NSI data; additional analysis shows this systematic bias remains in the 2026 NSI data. The manual correction

method resolves this bias at the city scale: inventories without the correction have only 8-9% of the mobile homes in the local reference inventory, whereas inventories with the correction have between 102-105%. An example Hayward neighborhood demonstrates this pattern in Figure 8, showing how the occupancy class for the corrected inventories more closely matches the reference inventory (Figure 8a-c). Again, results are shown for Exposure Model B, but the correction is consistent for both A and B, as shown in the supplemental materials.

Branch selection and implications: The better-supported branch for this decision is incorporating the manual correction. We then resolve this decision in the logic tree and collapse to that branch in the residual uncertainty analysis. This is particularly important because this misidentification has implications for seismic risk assessments, as mobile homes are significantly more likely to be damaged than wood-frame housing (see Figure 8d). Thus, undercounting mobile homes significantly underestimates earthquake damage. Future studies using NSI should evaluate mobile home representation and apply this correction.

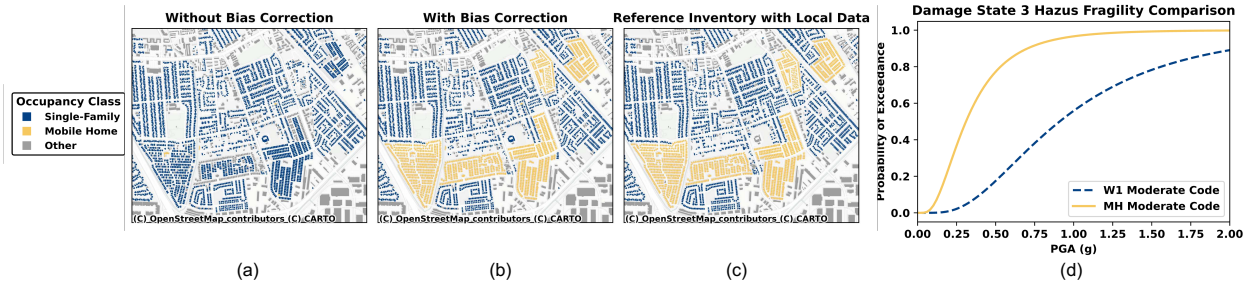


Fig. 8. Resolving bias using manual mobile home park polygons, illustrated using Exposure Model B. (a) Inventory without bias correction, (b) inventory with bias correction, which better aligns with (c) reference inventory with local data. (d) Fragility curves showing resulting difference in estimated seismic risk.

Structural System Inference Modifications (Decision F6 in Table S1)

Structural system is frequently missing across exposure data sources, but is a key feature in selecting an appropriate vulnerability model. Thus, it typically is inferred as part of exposure modeling.

Candidate branches: A common method for inferring structural systems uses the Hazus building stock mapping scheme tables [27], which provide region-specific percentages of floor area for different structural systems, given other building features such as occupancy class. In this study, we evaluate two branches related to the decision of how to infer the structural system. The first directly samples from the Hazus tables, and the second uses the regional adjustments proposed by Lochhead et al. [2] (see Figure 7, Possible Modification 2 in [2]). The adjusted branch combines three rules: (1) sample two-to-four-unit residential buildings using single-family mappings, based on their assumed construction similarity; (2) restrict the mobile-home structural system to mobile-home occupancy; and (3) exclude unreinforced masonry based on Hayward's retrofit efforts. A separate decision (Decision F5) constrains structural-system sampling using NSI building material (see Figure 7, Possible Modification 1 in [2]). This was not recommended with NSI 2022 because of large biases in the building material; improvements in NSI 2026 warrant retaining both alternatives in residual uncertainty. The regional adjustments are evaluated under contrasting assumptions about NSI material use, alongside other exposure modeling decisions via Exposure Models A and B.

Evidence of bias: Resulting inventories with and without the proposed adjustments are compared to the reference inventory with local data to assess which better approximates local data. For 2-4 unit residential buildings, the branch with the adjusted sampling (9b) more closely matches the building material distribution in the reference inventory (9c) than the branch that directly samples using Hazus (9a). Figure 9 shows the results for Exposure Model B, but the same holds true for Exposure Model A.

Branch selection and implications: We retain the branch with the structural system sampling modifications for the residual uncertainty analysis. Like the mobile home bias, this decision affects fragility curve assignment and modeled seismic risk. Overcounting masonry and concrete buildings in 2-4 unit residential buildings inflates damage and loss compared with wood-frame buildings, per the fragility curves shown in Figure 9d. Future studies using Hazus mapping scheme tables to sample structural systems should consider similar adjustments; modifying the sampling to treat 2-4 unit residential buildings as single-family homes and assigning mobile home structural systems only to mobile home occupancies seems logical based on engineering judgment and are recommended for adoption in future studies.

The modification to exclude unreinforced masonry is locally relevant in Hayward, but may not transfer elsewhere, depending on local retrofit ordinances.

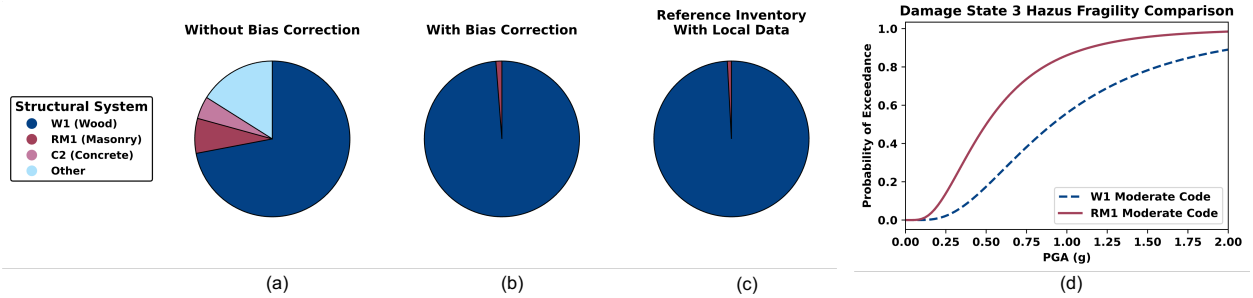


Fig. 9. Resolving bias using adjusted structural system inference, illustrated using Exposure Model B. (a) Inventory without bias correction, (b) inventory with bias correction, which better aligns with (c) reference inventory with local data. (d) Fragility curves showing resulting difference in estimated seismic risk.

Year of Construction (Decision F3 in Table S1)

The year of construction is a key building feature for assessing seismic risk, as it informs the design era and fragility selection. NSI data, in both the 2022 and 2026 versions, do not report the year of construction for individual buildings; instead, they report only the median year across the census tract, informed by the American Community Survey (ACS). Building-specific construction years are available in detailed local data (used as an external reference here) but are unavailable in the national data sources used in this study, motivating alternative assignment strategies.

Candidate branches: Three candidate branches are considered for modeling the year of construction. The first adopts the NSI median value directly, and the second uses a lognormal distribution to sample building age around the median with a pre-defined standard deviation of 0.2 (see supplemental materials for details). The third uses the *acs_year_built_sampler* module developed within the SimCenter, which produces a probability distribution over the ten decade bins of ACS table B25034 for each building, then draws from that distribution to assign a year of construction [28].

Evidence of bias: Comparisons with the external reference inventory reveal two limitations of assigning NSI median construction years to individual buildings. First, this assignment eliminates variation within each tract. In doing so, median assignment omits older construction years that can affect design-era classification and modeled seismic risk. In the Hayward inventory, all NSI-assigned years are 1955 or later, whereas 21% of buildings in the reference inventory were constructed before 1955 (Figure 10a). Second, NSI medians differ by up to 30 years from tract medians calculated using the reference inventory (Figure 10b). These discrepancies are distinct from the loss of within-tract variation, and preliminary analysis suggests that bias comes from NSI adopting the ACS median years, which are based on incomplete samples of census tracts. The lognormal and ACS-based sampling branches assess whether representing a distribution of construction years improves agreement with the reference inventory.

Figure 11 shows the results of the comparison across the three candidate branches relative to the reference inventory. The majority of this comparison centers on a neighborhood in the northern part of the city, as shown in Figure 11a, which shows the predicted mean loss ratio for each building using the reference inventory with local data. As shown in Figure 11b, the reference inventory produces the highest mean loss ratios primarily among older masonry buildings, illustrating the importance of the relationship between construction year and building material. Assigning NSI medians omits these older construction years and the corresponding high-loss buildings. Both sampling branches restore some older construction years and produce higher mean loss ratios for some buildings, to different extents. Among the three candidates, ACS-based sampling most closely reproduces the reference distributions of construction year and mean loss ratio in the evaluated neighborhood. Across all three branches, part of the discrepancy in mean loss ratio arises from differences in building material distributions, as shown in Figure 11b. Thus, there are biases in the nationally available data to characterize building material. In addition, none of the three branches reproduce the observed association between construction year and building material, because none explicitly conditions year assignment on material. Improved agreement in the individual distributions therefore does not establish accurate identification of the buildings with the highest modeled losses. Correlations between building features such as this are a compelling area for future work, but are not assessed further in this study.

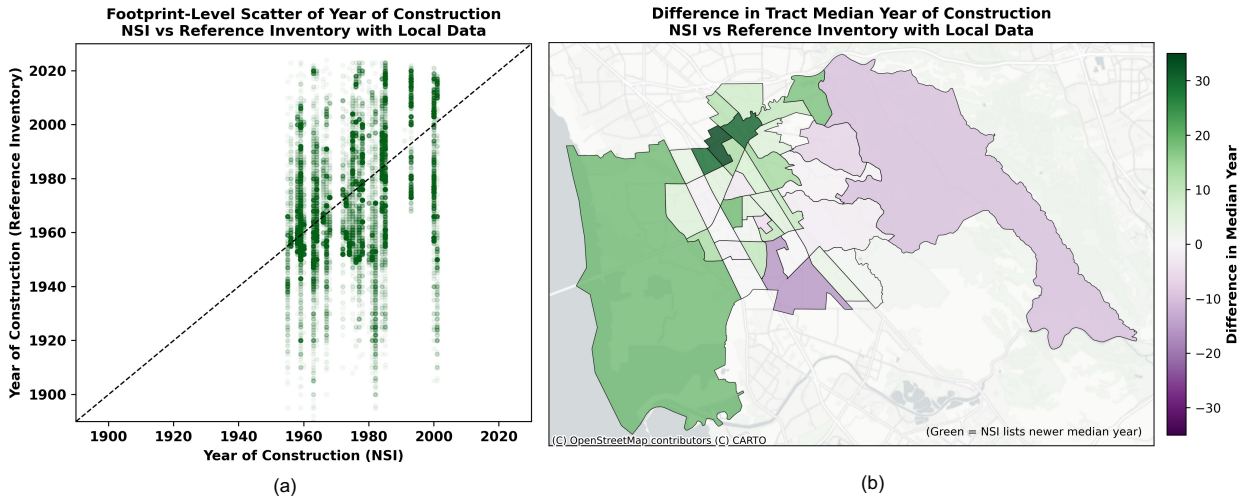


Fig. 10. Comparison between NSI year of construction versus footprint-specific year of construction from local reference inventory. (a) Footprint-level scatter for footprints where both data sources are available. (b) Difference in tract-level medians for tracts with at least ten footprints.

Branch selection and implications: Among the three candidate branches, the ACS-based sampling most closely reproduces the reference distributions of construction year and mean loss ratio in the evaluated area. This preference holds under both Exposure Models A and B (Figure 11 shows Exposure Model A), supporting selection of the ACS-based branch in the residual uncertainty analysis. However, none of the three considered branches reproduce the observed relationship between construction year and building material. Collapsing the logic tree to only the ACS-based method therefore reduces the identified bias without eliminating it, and the remaining dependence-related bias is not necessarily captured in the residual uncertainty. Despite these ongoing challenges, the results support using ACS-based sampling for future NSI-based inventories.

5.2 Residual Uncertainty

Seven of the twenty identified exposure modeling decisions can each be resolved to a single, better-supported branch of the logic tree, per Section 5.1 (see Table S1 for details). The remaining 13 decisions were also evaluated as possible cases of resolvable biases, but could not be resolved due to (1) a lack of reliable external data to serve as a reference when comparing logic tree branches, or (2) underdeveloped methodology to improve the existing estimation of detailed local data using national data. This section explores the impact of these unresolved decisions by generating inventories that represent all unique branch combinations remaining after biases are resolved in the logic tree, resulting in a spread of values for inventory characteristics and risk metrics (see Figure 4), which we define as residual uncertainty. In this section, we quantify this spread in inventory characteristics and downstream risk results, identify nonuniform patterns in the distribution of uncertainty across the building inventory, and investigate which decisions drive meaningful uncertainty in risk metrics.

Uncertainty in Inventory Characteristics

Based on the suite of inventories used to measure residual uncertainty, it is clear that some inventory characteristics have more uncertainty than others. This is best illustrated by comparing uncertainty in the number of buildings and the number of units in Hayward, two metrics that both carry implications for quantifying earthquake damage and its impact on the population. For example, the number of buildings per inventory, shown in Figure 12a, has a standard deviation of 160 buildings and a range of 677 buildings. This uncertainty is small relative to the total number of buildings. The uncertainty in the number of units per inventory, shown in Figure 12b, is greater, with a standard deviation of 2,467 units and a range of 8,246 units based solely on exposure modeling decisions. Comparing standard deviation or range directly is limiting when different metrics have different totals (e.g., total number of buildings and units), so to compare uncertainty, we take the coefficient of variation, which is the standard deviation divided by the mean. Comparing the coefficients of variation for these metrics, 0.0046 and 0.046, respectively, reveals that the number of units is more uncertain by a factor of ten. This can be explained by the decision to employ the most comprehensive data footprint

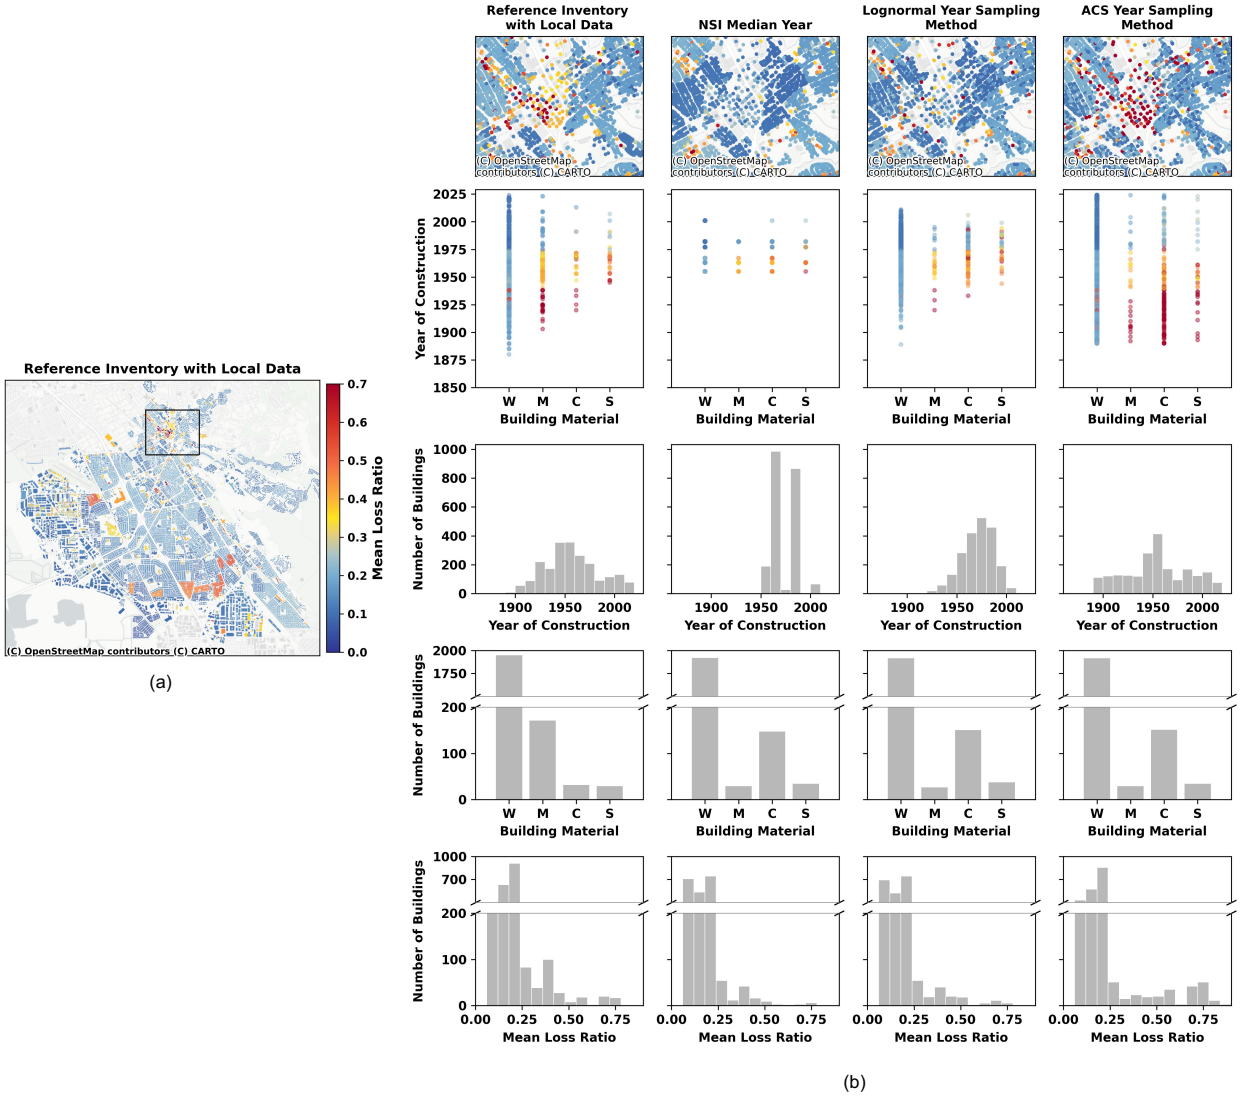


Fig. 11. Year of construction branch comparison for Exposure Model A. (a) Mean loss ratio across Hayward for the reference inventory, with the selected neighborhood boxed. (b) Neighborhood results for the reference inventory and three candidate branches, with rows (in descending order) showing mean loss ratio by footprint, construction year by building material colored by mean loss ratio, histograms of year of construction, bar charts of building material, and histograms of mean loss ratio. Mean loss ratio (map and scatter) share the color scale in (a). W, M, C, and S denote wood, masonry, concrete, and steel.

source and robust data attribution method available across all inventories used to measure residual uncertainty. In other words, the bias is reduced by not dropping data. The higher levels of uncertainty in other inventory characteristics that are not impacted by choosing a preferred branch for this decision underline the efficacy of removing this bias, outlined in the previous section.

Another instance of significant uncertainty is observed in the full structural replacement value of the inventory, whose distribution of values is shown in Figure 12c. Full structural replacement value is calculated for each generated inventory by summing up the starting value of all the buildings within that inventory. Relative to a mean of $34.6 billion, the full structural replacement value of Hayward can range from $26.9 billion to $39.7 billion with a standard deviation of $4.7 billion. For comparative purposes, this uncertainty can be captured by a coefficient of variation of 0.14. Its distribution shows three distinct clusters, consistent with many other inventory characteristics that exhibit

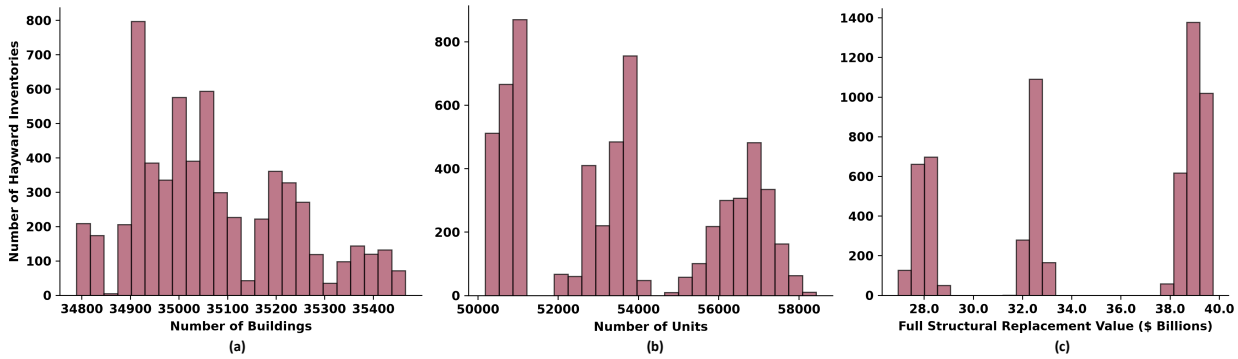


Fig. 12. Distributions of selected inventory characteristics across inventories: (a) number of buildings, (b) number of units, and (c) full structural replacement value.

multimodal, nonnormal distributions separated into groupings. By tracing the methodological decisions that produce distinct inventory clusters and applying visual filters to isolate the influence of each decision on the resulting inventories, we identified which decisions drive uncertainty for a given inventory characteristic. While a one-way ANOVA was conducted to validate these graphical patterns, the resulting groupings in inventory characteristics are sufficiently distinct such that two-way ANOVA or other analyses designed to capture higher-order interactions between decisions were not required to support the conclusions drawn.

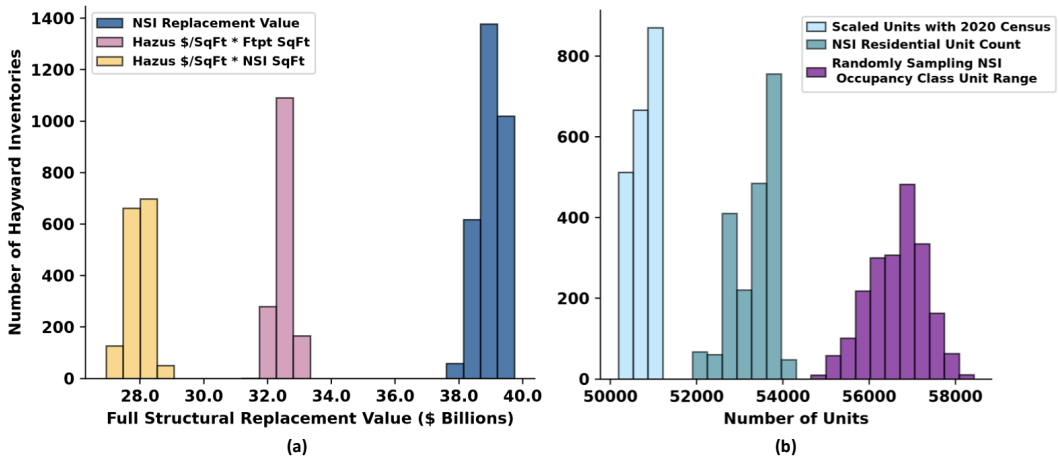


Fig. 13. (a) Distribution of full structural replacement value mapped by decisions F1 and F4 and (b) distribution of number of units in Hayward mapped by decision F2.

Upon examining residual uncertainty in the number of units and the full structural replacement value of the Hayward inventory, it is visually apparent that single decisions govern these patterns. The three distinct groupings seen in the distribution of Hayward's full structural replacement value, seen in Figure 13a, are driven by the decision of whether to employ NSI-assigned full structural replacement value or calculate full structural replacement value using Hazus occupancy-class specific value per square foot ratios multiplied by the total area of each building (Decision F4); if we choose to use Hazus ratios, the area can be pulled from NSI reported values or from the building footprint source used to define underlying geometries (Decision F1). These two decisions result in the three clusters that constitute the distribution in full structural replacement value of Hayward. The distribution of number of units, in Figure 13b, can be traced to the decision of using the residential unit count that NSI reports within its dataset, by randomly sampling the number of units per occupancy class defined by NSI's HAZUS taxonomy (i.e., a RES3C buildings is defined as having 5-9 units), or estimating the number of units per building by scaling NSI nighttime population against 2020 Census block unit totals, as proposed in Lochhead et al. [2]. The distinctive separation of these groupings in these two

metrics underlines the potential for key decisions in the inventory synthesis workflow to develop significantly different inventories and drive uncertainty.

Uncertainty in Damage and Loss Metrics

After all inventories are propagated through the same M7 scenario earthquake, downstream risk metrics also exhibit spreads of possible values from the different inventories due to effects of decisions in the inventory synthesis workflow. The shapes of these distributions reveal that decisions which drive uncertainty in inventory metrics also impact the outcomes of a risk assessment. This can be seen in Figure 14a, the distribution of values for the total repair cost of Hayward. Repair cost per inventory is calculated according to Equation 1. From the three distinctive groupings in the distribution mirroring the distribution of full structural replacement value, it is clear that this uncertainty is again caused by the decision of source and calculation method for full structural replacement value. Such uncertainty is on a magnitude that would significantly impact stakeholder decision-making: relative to a mean of $7.3 billion, Hayward's total repair cost has a standard deviation of $1.1 billion and a total swing of $3.3 billion, which represents 46% of the mean.

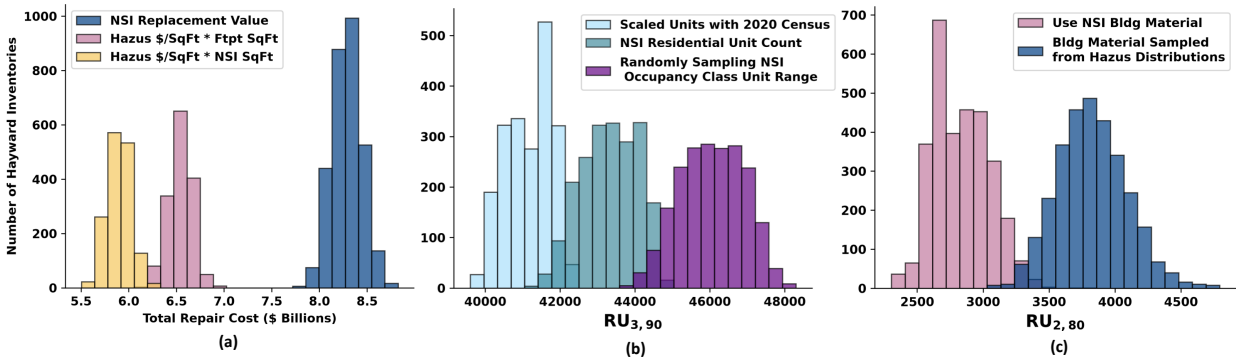


Fig. 14. (a) Distribution of repair cost mapped by decisions F1 and F4, (b) distribution of number of residential units (RU) with a 1 in 10 probability of extensive damage in Hayward mapped by decision F2, and (c) distribution of number of residential units (RU) with a 4 in 5 probability of damage mapped by decision F5. Metrics are derived according to Equations 1 through 3.

A similar pattern is observed in the distribution of values for $RU_{3,10}$ (the number of units in Hayward with a 1 in 10 probability of extensive damage). This spread consists of three distinctive groupings which are governed by the same decision which impacted uncertainty in the number of units in Hayward overall. As is evident in Figure 14b, the uncertainty in this damage metric stems from the decision of estimating the number of units per building by scaling NSI nighttime population against 2020 Census block unit totals, using the residential unit count that NSI reports within its dataset, or by randomly sampling the number of units per occupancy class defined by NSI's HAZUS taxonomy. This results in a swing of 8,729 units captured by the damage metric relative to a mean of 43,524 units and a coefficient of variation of 0.05.

We also observed a decision whose impacts appeared after the propagation of the inventories through the risk assessment: whether or not to employ NSI-reported building material or sample the building material from the provided Hazus distribution. This decision drove uncertainty in $RU_{2,80}$, the damage metric capturing the number of units with a 4 in 5 probability of suffering damage following the simulated M7 earthquake, and resulted in a swing of 2,489 units captured by the damage metric relative to a mean of 3,328 units, and a coefficient of variation of 0.17. The appearance of this decision in driving risk metric uncertainty demonstrates the potential for different exposure modeling decisions to impact risk assessments downstream, even if immediate effects on inventory characteristics are not obvious. As such, the metrics investigated in this study do not represent an exhaustive view of the residual uncertainty's impact on possible risk metrics. Nonetheless, they illustrate that the uncertainty is of a magnitude such that decision-making at a city scale would be hindered and that key decisions in the inventory synthesis workflow are important drivers of uncertainty in downstream outcomes.

Uneven Distribution of Uncertainty Across the Inventory

Furthermore, this uncertainty in inventory characteristics and results of a risk assessment were shown to be not impact all parts of the Hayward inventory to the same extent. In this paper, we limit our examination of this nonuniform distribution to occupancy classes of buildings and geospatial distribution; however, this preliminary investigation demonstrates that cross-sections of an inventory should not be assumed to carry proportional amounts of uncertainty. Furthermore, the analysis and the plots created as a result can prove useful diagnostic tools for understanding the distribution of uncertainty across an inventory, in their revealing of trends or of "hotspots" of uncertainty.

In Figure 16a, by comparing the coefficient of variation across residential building occupancy classes, it is clear that uncertainty does not impact all occupancy classes to the same extent. Furthermore, we observed that the uncertainty in the full structural replacement value for mobile homes in Hayward was unusually high; this can be attributed to NSI having labeled the majority of mobile homes as single-family homes or duplexes, thus inflating their value. This impacts the uncertainty created by decisions in the inventory workflow of how to determine building value. We also observed that multi-family residential occupancy classes with higher unit counts tended to have slightly more uncertainty in the number of units. Although patterns and outliers may vary depending on the composition of a given case study, nonuniformity in the impact of uncertainty is likely to persist based on differences in the quality and extent of data available across buildings in an inventory.

As was observed for inventory statistics, uncertainty in damage characteristics does not impact buildings of all occupancy classes to the same extent, as is seen in 16b. Continuing to use the coefficient of variation as the basis of comparison, we find a clear pattern that uncertainty in the number of buildings and units with a given probability of damage impacts multi-family residential buildings to a greater extent. This pattern holds when considering $RU_{3,10}$ (the number of units with a 1 in 10 chance of sustaining extensive, displacement-inducing damage) and with $RU_{3,20}$ (2 in 10 chance of sustaining such damage) and, although not within the scope of this study, likely across other damage metrics.

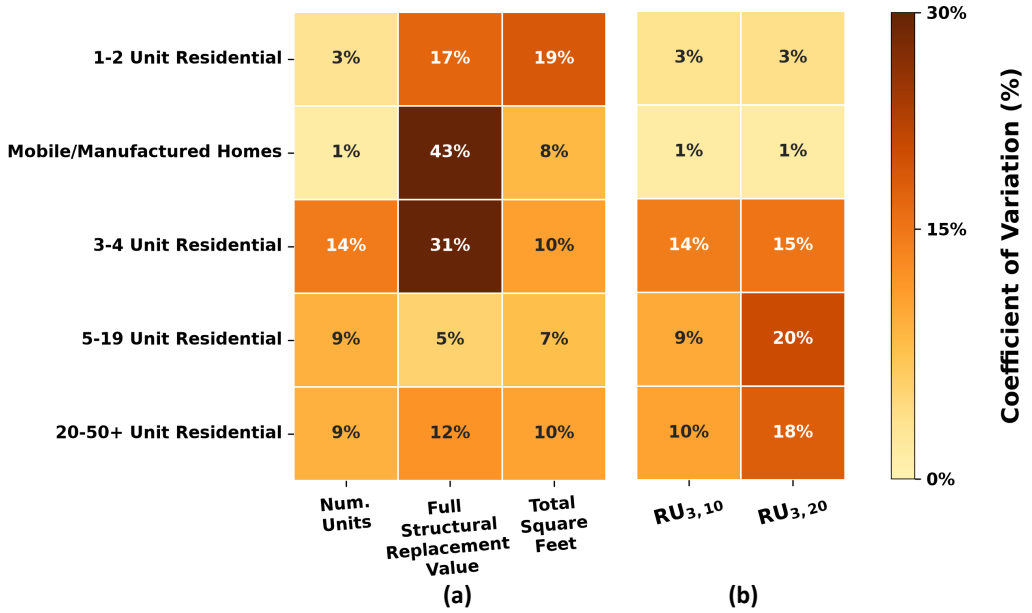


Fig. 15. Coefficient of variation by residential occupancy class across (a) inventory characteristics and (b) selected damage metrics.

Due to the nonuniformity in the distribution of uncertainty across building occupancy classes and the clustering of buildings constructed in the same design era, among other factors, we observed nonuniformity in the spatial distribution of uncertainty for inventory statistics and risk metrics. Furthermore, the impact of propagating inventory uncertainty through a risk assessment on downstream risk metrics is apparent when visually comparing the spatial distribution of full structural replacement value for residential buildings (16a), loss ratio (16b) for residential buildings, and repair cost (16c) of these buildings; repair cost is the most variable because it is the product of the first two, but it is also

the most critical to city stakeholders. Such spatial visualizations highlight local hotspots of uncertainty, thus better directing the efforts of inventory developers to mitigate its magnitude and improve understanding of city-level risk.

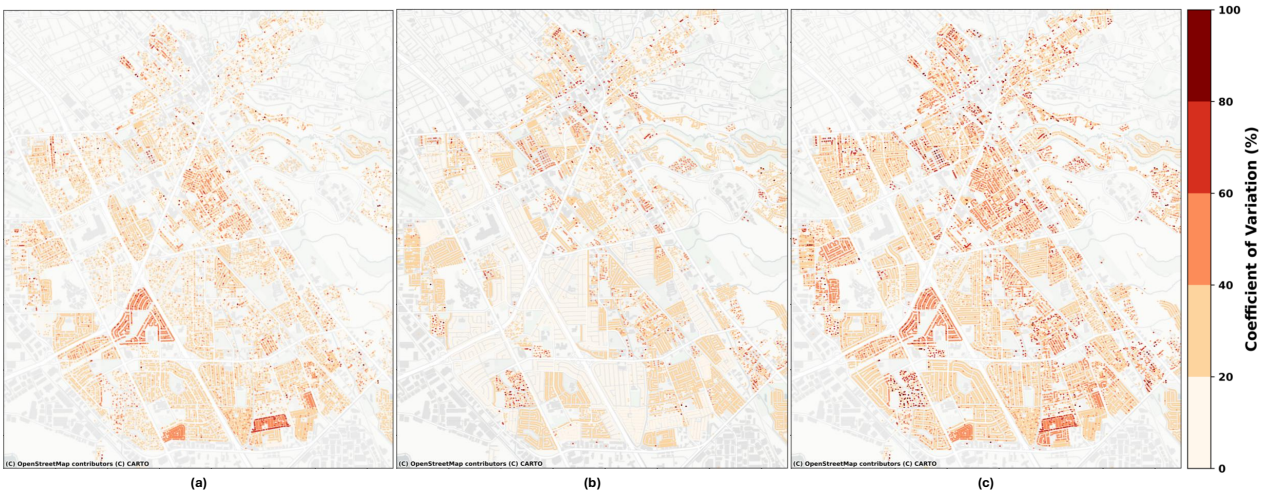


Fig. 16. Coefficient of variation of (a) full structural replacement value, (b) loss ratio, and (c) repair cost across inventories for each residential building footprint.

Drivers of Uncertainty in Risk Metrics

Alongside the generation of histogram distributions to visually observe the effect of decisions on uncertainty in downstream risk metrics, a one-way ANOVA test was conducted as a complementary numerical analysis. This test considered all exposure modeling decisions included in the residual uncertainty analysis, outlined in the supplemental materials, and returned the relative proportion of uncertainty in each metric attributable to each decision. In this way, the more influential decisions were numerically identified; if efforts were to be targeted at resolving only a few decisions, this analysis pinpoints where the biggest impact can be made with regards to reducing uncertainty in metrics that impact stakeholder decision-making.

In the following plots, we broadly categorized the decisions included in this study as spatial modeling decisions and feature modeling decisions. The reasoning for the distinction between these groupings was outlined above in Figure 2. We separated the exposure modeling decisions into these two categories, distinguished by color family. The inner ring of the circular sunburst plots displays the proportion of uncertainty in a given metric attributed to the broad spatial modeling decision and feature modeling decision families, while the outer ring further displays the proportion of uncertainty from each of the individual, contributing decisions within these broad categories. For example, 17(a) displays the relative influence of decisions on the total repair cost of Hayward. From the inner ring, it is apparent that feature modeling decisions constituted much of the uncertainty in that metric, whereas spatial modeling decisions did not contribute. Furthermore, the outer ring reveals that the decision of whether to use NSI as the source for building value or Hazus occupancy-class specific ratios was the most influential individual decision. We can also see that the decision of pulling area from NSI reported values or from the building footprint source also drove some of the uncertainty. A similar analysis can be done for the repair cost of residential buildings in Hayward, shown in 17(b). It is also important to consider the absolute value of these metrics; thus, the ranges of each repair cost metric are provided in 17(c). A decision that appears quite influential in a given risk metric may matter little to a stakeholder if such a metric has a small cost or a small number of buildings potentially impacted by the simulated hazard.

This same analysis is applied to the damage metric used in this study, whose notation mapping can be found in 2. Once again, the absolute number of buildings or units (18(a)) in each of these metrics likely factors into their importance for stakeholders. For all damage metrics in this study, we calculated the average influence of each individual decision on the resulting uncertainty, plotted in 18(b). This revealed that the decision of whether or not to employ NSI-reported building material or sample the building material from the provided Hazus distribution contributed more than half of all uncertainty, highlighting this decision as an important target of future efforts to further standardize the exposure modeling process. Another important decision was revealed to be the method of determining the number of units. In

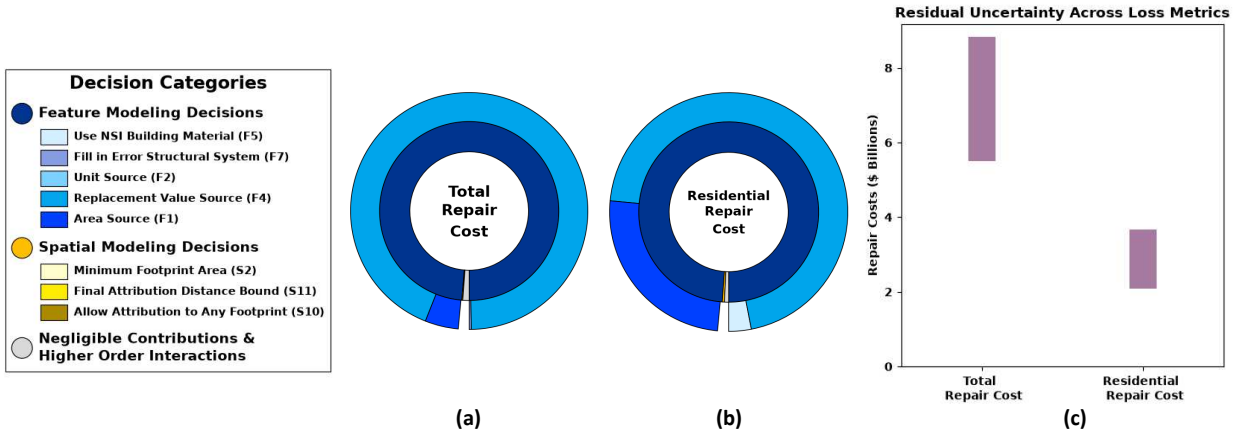


Fig. 17. Influence of exposure modeling decisions on (a) total repair cost of Hayward and (b) repair cost of residential buildings in Hayward. (c) Residual uncertainty in selected loss metrics. The legend references labels given to the decisions that correspond to their in-depth explanations in the supplemental materials.

the individual plots for number of buildings and units in the damage metrics (18(c-h), it is evident that feature modeling choices are more influential overall, and spatial modeling choices have more influence on the uncertainty in damage metrics pertaining to buildings, rather than units.

The findings from such an analysis will vary as the case study inventory varies; however, by identifying which decisions are most influential in creating uncertainty, these methods can be helpful in determining where research should be directed to improve confidence in the metrics outputted from regional risk assessments.

6 DISCUSSION AND IMPLICATIONS

Per the results, there is substantial uncertainty in regional risk metrics due to exposure modeling decisions, both in terms of total exposure model uncertainty, which represents the spread due to reasonable exposure modeling decisions, and the residual uncertainty, which represents the uncertainty still present in the workflow once all known biases have been addressed. Two questions remain about the implications of this uncertainty, including whether it is meaningful relative to other uncertainties in the risk framework and whether capturing only within-model uncertainty is sufficient to approximate exposure model uncertainty.

Regional seismic risk assessments have many uncertainties, and the question remains whether the exposure model uncertainty is meaningful relative to other epistemic uncertainties in the risk modeling framework. One such uncertainty in the regional seismic risk modeling framework is the epistemic uncertainty due to ground motion model selection. Both the exposure model uncertainty considered in this study and the uncertainty between candidate ground motion models are the result of modeling choices and represent epistemic uncertainty (see Figure 1). To compare these two uncertainties, we conduct a risk assessment using a single inventory across four commonly used ground motion models [29, 18, 30, 31]. Intensity measures are defined using the same method as previously discussed; for each of the four models, the same UCERF2 rupture is selected, the selected ground motion model is used to generate 50 realizations of IMs, and the spatial correlation model from Jayaram and Baker [6] is used across all four ground motion models. Each suite of IMs is then paired with a single inventory (Exposure Model A and B, previously discussed) and used to sample 1000 damage realizations. The results can then be compared between the case of 1) a single inventory across four ground motion models, and 2) a single ground motion model used across all inventories, as was done throughout all analyses.

The second question to consider is whether capturing within-model exposure uncertainty is sufficient. In this case, within-model exposure uncertainty is defined as using a single exposure model (i.e., one set of input data, modeling decisions, methods, etc.), but sampling multiple times. Thus, any probabilistic step, such as sampling a structural system, is done repeatedly. This constitutes the within-model versus between-model distinction in Figure 1. Several recent papers have evaluated within-model uncertainty [13, 16], so this comparison aims to assess whether most exposure modeling uncertainty comes from within-model uncertainty or whether meaningful between-model uncertainty exists.

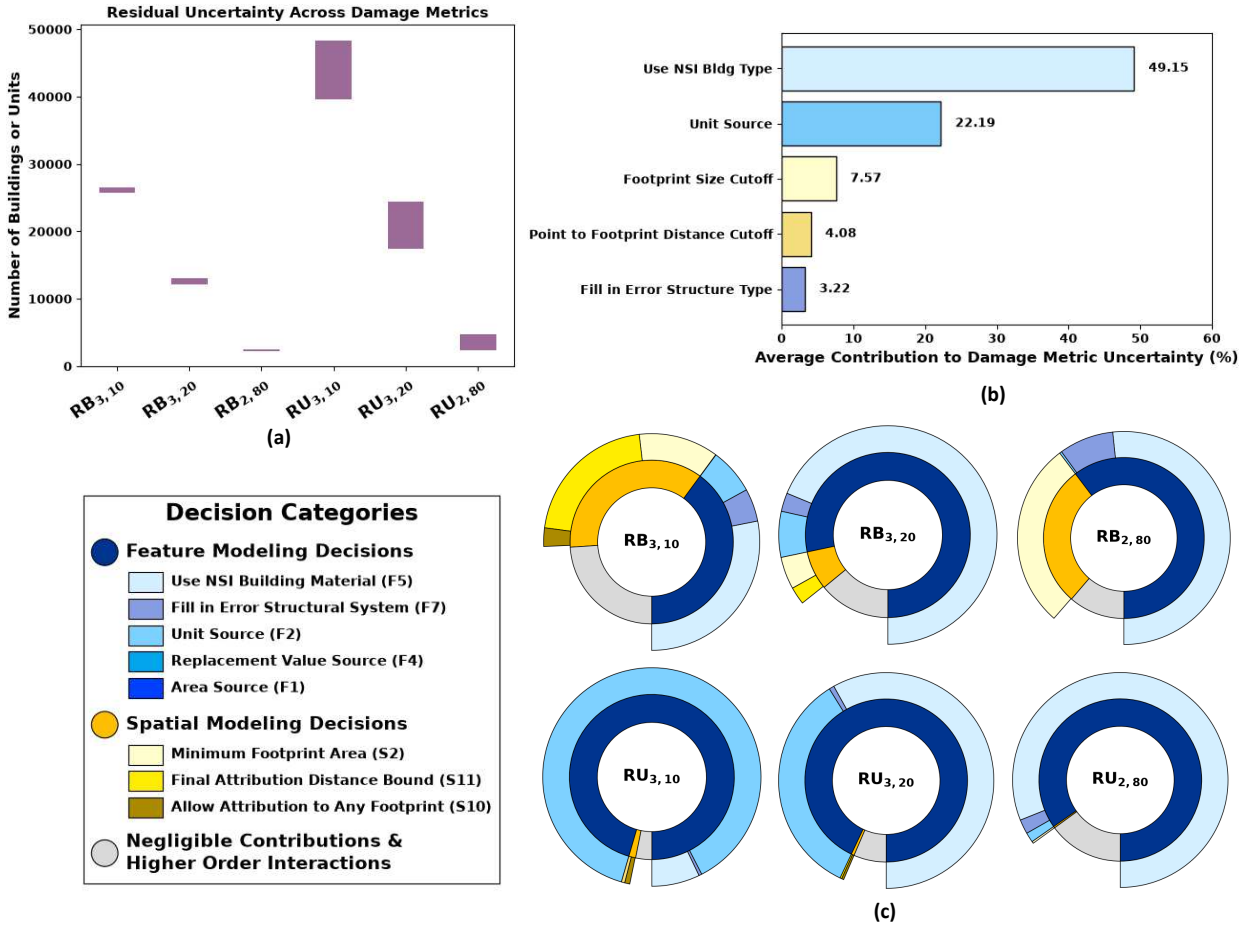


Fig. 18. (a) Residual uncertainty across damage metrics. (b) Average influence of exposure modeling decisions on damage metrics. (c) Analysis of variance results showing the effects of individual decisions on damage metrics.

Figure 19 shows the results of both comparisons across various inventory, damage, and loss metrics. Several primary observations emerge from this figure. First, there is substantially more total exposure model uncertainty than there is residual uncertainty, demonstrating the clear benefits of robust exposure modeling methods and bias reduction efforts. In other words, very large ranges in metrics can result from seemingly reasonable exposure modeling decisions, and simple strategies can meaningfully reduce this uncertainty. The second main conclusion is that exposure model uncertainty (including both total and residual) is non-negligible relative to the epistemic uncertainty due to ground motion model selection. For the plots of inventory metrics prior to the scenario earthquake (buildings, units, and full structural replacement value), the ground motion model does not create any uncertainty, as the ground motions have not yet been applied. Across all damage and loss metrics, total exposure model uncertainty exceeds that of the ground motion models. Similarly, residual uncertainty meets or exceeds ground motion model uncertainty for most damage and loss metrics. This demonstrates that exposure model uncertainty is meaningful in the broader risk assessment context, further demonstrating the need for exposure model uncertainty quantification and reduction.

Figure 19 also shows that capturing only within-model uncertainty greatly underestimates the residual uncertainty across all metrics. This implies that while sampling structural system probabilistically is meaningful in terms of uncertainty for a single building, these effects are minimal in determining city-level risk as compared to exposure modeling decisions such as selecting input data sources or methods. Thus, it is insufficient to focus only on within-model uncertainty and probabilistic gap filling if the goal is exposure model uncertainty quantification.

Finally, Figure 19 demonstrates that the amount of uncertainty due to various sources changes based on the metric being assessed, which is consistent with the ANOVA findings. For example, the uncertainty in the number of buildings

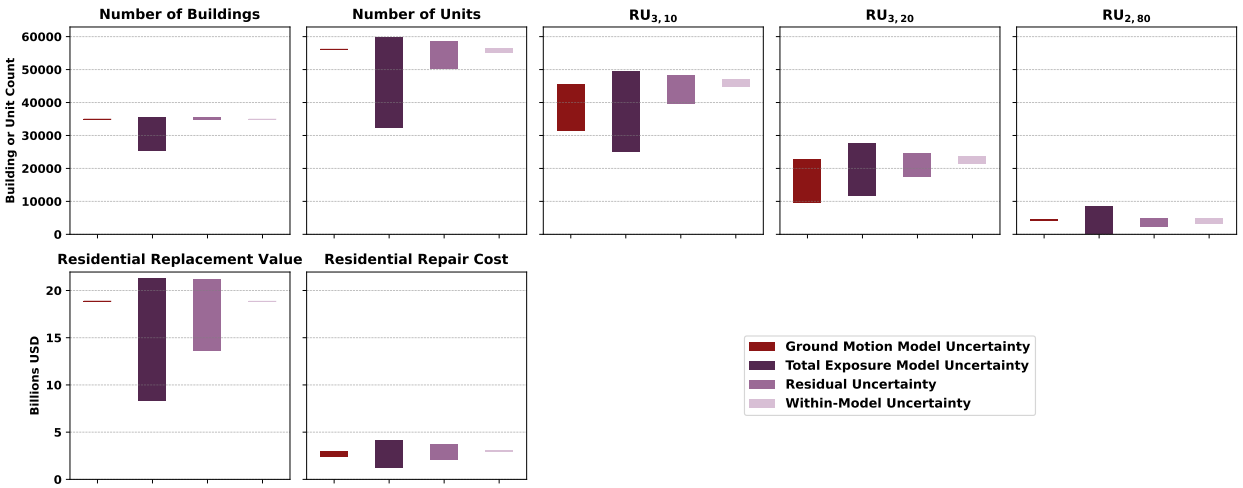


Fig. 19. Comparison of total exposure model uncertainty, residual uncertainty, and within-model uncertainty across inventory, damage, and loss metrics, as compared to the epistemic uncertainty due to ground motion model selection.

is almost exclusively due to total exposure model uncertainty; once biases have been resolved by using the enhanced spatial join method and a robust footprint source, residual uncertainty is small. Furthermore, while there is substantial residual uncertainty and uncertainty due to ground motion models for $RU_{3,10}$ and $RU_{3,20}$, the uncertainty in $RU_{2,80}$ is much lower for these sources. In this case, this pattern can be attributed to mobile homes, as $RU_{2,80}$ describes buildings that are very likely damaged and thus serves as a proxy for mobile homes. Since mobile homes appear consistently in the inventory once bias is resolved, this metric has little ground motion model, residual, or within-model uncertainty. This broadly demonstrates that uncertainty is nonuniform and that forms of uncertainty (i.e., ground motion model versus total exposure versus residual) are more impactful in some inventory, damage, and loss metrics than in others.

7 CONCLUSIONS

This study proposes the use of a logic tree to evaluate exposure modeling uncertainty in the context of regional seismic risk assessment. Using the case study city of Hayward, California, and the National Synthesis Workflow defined in Lochhead et al. [2], this study quantifies total exposure model uncertainty, defined as the spread in inventory and downstream risk metrics resulting from different reasonable exposure modeling decisions. In addition, this study explores several methods to reduce bias in exposure models by comparing inventories, generated by traversing different combinations of branches in the logic tree, against external reference data and engineering judgment. These strategies to resolve bias in the logic tree are recommended for use in future case studies that leverage the National Structure Inventory (NSI) data. Once bias has been resolved and the logic tree has been collapsed to its better-supported branches when sufficient information is available, the remaining uncertainty resulting from the reduced logic tree is defined as residual uncertainty and is recommended to be propagated in regional seismic risk assessments to more fully characterize exposure uncertainty. This study also presents a statistical analysis of residual uncertainty and compares various exposure model uncertainties with the epistemic uncertainty due to ground motion model selection.

The analysis presented in this study has several limitations. First, all results are shown for a single case study city, Hayward, CA, which may not be representative of all communities. Additional analysis is needed to assess whether similar patterns emerge for other test case cities. Furthermore, the study focuses on footprint-level characterization of building features needed to conduct a HAZUS-style assessment. If more detailed vulnerability models were adopted or alternate hazards considered, different building features would be needed to inform downstream analysis. In addition, we narrow the scope of exposure model uncertainty to only what is contained within the National Synthesis Workflow outlined in Figure 3, meaning all inventories resulting from the logic tree originate from many of the same input data sources. Given that, the exposure model uncertainty estimated in this study is likely an underestimation relative to if all possible input data sources were considered. We constrained the analysis to the National Synthesis Workflow because different locations have different local data availability and quality, so leveraging only national data sources makes results more broadly applicable. Furthermore, all logic tree branches were weighted equally, as we considered all reasonable exposure modeling decisions; however, if the logic tree incorporates higher-quality local data as well as

more generalized national data, judgments on data quality and logic tree branch weights would be required. This could present compelling future work.

The following takeaways from the case study presented in this paper are expected to be relevant to exposure uncertainty more broadly.

Exposure model uncertainty is non-negligible relative to other components of risk. This study compares total and residual exposure model uncertainty to the epistemic uncertainty due to ground motion model selection and finds that exposure model uncertainty is non-negligible. Across downstream risk metrics, including various damage and loss metrics, total exposure model uncertainty exceeds epistemic uncertainty from ground motion model selection. Furthermore, residual uncertainty, which this study recommends propagating through future regional seismic risk assessments, is also comparable to ground motion model uncertainty. This illustrates the relevance and importance of considering exposure uncertainty in risk assessments.

Exposure uncertainty should be propagated in regional seismic risk assessments, and considering only within-model uncertainty is insufficient. Due to the non-negligible impacts of exposure model uncertainty, this study recommends that future studies quantify and propagate exposure model uncertainty in regional seismic risk assessments. More specifically, residual uncertainty, which characterizes exposure uncertainty once reasonable efforts to resolve bias have been implemented, should be propagated. This is particularly important if stakeholders intend to use the results of a regional study to inform decisions. In addition, several recent studies have assessed within-model uncertainty, where gaps in the data are sampled probabilistically rather than filled deterministically. While this is a step towards capturing exposure uncertainty, it results in an underestimation of uncertainty in downstream damage and loss metrics. In other words, between-model exposure uncertainty drives much of the variation in seismic risk metrics at the city scale. Thus, it is important to go beyond just probabilistic sampling of gaps if the goal is to robustly capture exposure uncertainty.

In the context of regional building inventories using the National Structure Inventory, significant biases can be resolved through robust exposure modeling methods. These are recommended for future use. Total exposure model uncertainty, which captures the spread in metrics due to reasonable modeling decisions, can be greatly reduced through robust exposure modeling methods to resolve bias. This is promising for future studies as it enables higher quality exposure models without requiring additional manual data collection. There are several exposure modeling steps that drive uncertainty, such as building footprint source and data attribution method selection. Using robust methods for these steps dramatically reduces bias relative to other commonly used methods. In addition, manually correcting NSI's inaccurate mobile home characterization, assigning 2-4 unit buildings to structural systems of single-family homes rather than larger multiunit buildings, and using the ACS-based sampling method to assign year of construction rather than using NSI median all substantially improve the exposure model. Future users of NSI data should adopt these decisions to reduce bias in their inventories, though other biases not present in Hayward may arise.

Remaining uncertainty is not distributed uniformly across inventory characteristics, downstream risk metrics, occupancy classes, and geospatial locations. We observe that certain inventory characteristics, such as full structural replacement value, are impacted by the remaining exposure uncertainty to a greater extent than others, such as number of buildings in the inventory. This occurs in the damage metrics targeted as results of the risk assessment, where the magnitude of uncertainty in certain damage states such as $RB_{2,80}$ (4 in 5 chance of damage) is less than that in states such as $RB_{3,20}$ (1 in 5 chance of extensive damage). Additionally, through the use of comparative figures (i.e. heatmaps and plots of footprint data points with underlying geospatial data), we observe that different occupancy classes and different neighborhoods throughout the Hayward inventory experience different amounts of uncertainty. Such disparate impacts of the remaining exposure uncertainty are important for a stakeholder interested in implications of uncertainty for their case study region.

As sources of uncertainty differ depending on the risk metric under consideration, context is essential to determine the optimal direction of resources for further reducing exposure uncertainty. The preliminary histogram analysis and numerical one-way ANOVA demonstrate that different exposure modeling decisions are responsible for uncertainty in different risk metrics. Thus, understanding which risk metrics are more valuable to stakeholders helps to determine which decisions are most consequential. For example, if a stakeholder were interested in the number of families facing displacement following the M7 earthquake, the damage metric $RB_{3,20}$, which measures number of residential units with a 1 in 5 chance of high, displacement-inducing damage, is of interest. The decision of whether to use the NSI-reported building type or use Hazus occupancy-class specific ratios drives the most uncertainty in that metric and thus is of concern to the stakeholder. This is not the case for every damage metric selected for this case study and/or of interest to stakeholders. Thus, in order to fully determine where resources should be directed, the objective of the study and the context surrounding the risk assessment must be taken into account. If these are not

available, we generally observe that feature-level decisions (i.e. data source prioritization, imputation methodology, and inference methodology) create larger amounts of uncertainty than inventory-level modeling decisions. Resolution of these decisions should be a focus of future work.

This study captures the significance of building inventories in accurately and meaningfully understanding regional risk. We demonstrate that exposure model uncertainty propagates through risk assessments and produces uncertainty in risk metrics essential for decision-making at a regional scale. Further attention should be directed towards the footprint-level building inventories underpinning regional risk assessments, as methods used and decisions made in their development impact how we inform actions to protect communities exposed to natural hazards.

7.1 Software and data availability

All Python functions, example scripts, and data used in the National Synthesis Workflow are available on GitHub at https://github.com/mlochhead/Building_Inventory_Generation. The core Python functions are designed to support inventory generation beyond the Hayward case study with minimal modification.

7.2 Declaration of Generative AI and AI-assisted technologies in the writing process

During the preparation of this work, the authors used Claude to edit sections of the manuscript for language and clarity. After using this tool, the authors reviewed and edited the text and take full responsibility for the content of the published article.

7.3 Acknowledgments

This material is based upon work supported by the National Science Foundation Graduate Research Fellowship Program under Grant No. DGE-2146755, the Stanford Graduate Fellowship, the SimCenter (supported by the National Science Foundation, United States under Grants CMMI-1612843 and 2131111), the Sustainability, Engineering, and Science Undergraduate Research Program, and the Stanford Vice Provost for Undergraduate Education Conference Grant. Any opinions, findings, and conclusions expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation. The authors would also like to thank Dr. Jack Baker for his support and feedback throughout the project, as well as his expertise in building inventories and regional risk assessment.

S1 SUPPLEMENTAL MATERIALS

S1.1 Exposure Modeling Decisions Included in Total Uncertainty

Table S1 lists each exposure modeling decision that was considered when evaluating total exposure model uncertainty, given the National Synthesis Workflow shown in Figure 3. Each decision represents a branch point in the logic tree that characterizes total exposure model uncertainty, and each option represents an individual branch. Decisions are grouped into *spatial modeling decisions* and *feature modeling decisions*, as described in Figure 2. The *branch after resolving bias* column indicates the preferred branch that is subsequently used in residual uncertainty evaluation. When an option is listed, efforts to resolve bias showed that branch was preferred, meaning it was most effective at reducing bias in the resulting inventory. When no option is listed, all options remain in residual uncertainty evaluation and the decision is unresolved. Detailed descriptions of each option and the flag name used in the code repository follow the table.

TABLE S1. Exposure modeling decisions characterizing total exposure model uncertainty.

#	Decision	Options	Reason for inclusion	Branch after resolving bias
Spatial modeling decisions				
S1	Footprint source	Microsoft OpenStreetMap Overture USA Structures Local (City of Hayward)	All listed footprint sources are used for regional risk assessments in the literature. Varying sources capture differences in coverage and geometry, and no single source is currently complete or most accurate everywhere. Based on the bias evaluation in Figures 7 and 6, the local footprint source is preferred. In general, it is important to evaluate the quality and completeness of a footprint source prior to use in a new case study location.	Local
S2	Minimum footprint area	0 ft^2 450 ft^2	USA Structures lists a minimum footprint size of 450 ft^2, as small polygons may be sheds or garages. However, filtering may remove some real dwellings.	–
S3	Minimum NSI floor area	0 ft^2 450 ft^2	NSI contains points with floor area values below what is plausible for a habitable structure. Filtering these points out assumes they are not meaningful structures, consistent with the footprint filter (S2). However, this removes buildings, value, and people from the inventory, which may underestimate exposure.	–
S4	Merge NSI points with duplicate building IDs	Yes No	NSI documentation lists building ID (BID) as the north-axis-aligned "bounding box" of its footprint. BID values can be duplicated across NSI points, and the documentation is ambiguous about the relationship between NSI points with the same BID. When merged, additive fields (e.g., number of units) are summed, and categorical fields (e.g., occupancy class) are retained as lists.	Yes
S5	Augment NSI using HIFLD data	Yes No	As shown in Lochhead et al. Lochhead et al. [2], HIFLD data can be used to improve the characterization of schools and emergency response facilities relative to NSI data. The option to not include HIFLD is considered reasonable as many studies adopt raw NSI data, but is neither recommended nor included in residual uncertainty. Only the HIFLD integration method of Lochhead et al. Lochhead et al. [2] is considered, though HIFLD could be incorporated in alternative ways.	Yes
S6	Point-to-footprint attribution method	Enhanced spatial join Simple spatial join Nearest neighbor	Several methods to attribute points to footprints are used in the literature and are evaluated here. Based on the bias evaluation in Figure 6, the enhanced spatial join method proposed in Lochhead et al. Lochhead et al. [2] is preferred.	Enhanced

Continued on next page

Table S1 continued

#	Decision	Options	Reason for inclusion	Branch after resolving bias
S7†	Allow points to move between census blocks during attribution	Yes No	NSI points located near block boundaries may belong to footprints in an adjacent block. However, some steps in the development of NSI are likely conducted to be consistent with census totals, so moving points across block boundaries shifts units, population, and value between blocks.	–
S8†	Estimate footprint capacity using height	Yes No	Footprint capacity is used in S9 to determine if a footprint is larger than typical for its assigned occupancy. Approximating the number of stories from footprint height (available in many footprint data sources) introduces error, but using footprint area alone treats every building as single-story.	–
S9†	Prioritize partially full footprints	Yes No	Beyond empty footprints, points can be preferentially attributed to footprints larger than typical for their assigned occupancy over footprints of typical size for their assigned occupancy (i.e., partially full footprints in Lochhead et al. Lochhead et al. [2]). A partially full footprint may contain uses not captured by the NSI points already attributed to it, but its size may also reflect geometry error.	–
S10†	Allow attribution to any footprint	Yes No	Beyond empty and partially full footprints (S9), NSI points located outside of footprints can be attributed to the nearest footprint regardless of whether it already contains data. Assigning additional points to already-occupied footprints retains more NSI data but may overload buildings or combine unrelated uses.	–
S11†	Final attribution distance bound	10 m 100 m	Most NSI points fall relatively close to building footprints. However, imprecise or inaccurate geolocation or missing building footprints mean this is not always the case. Allowing points to be moved farther retains more NSI data, though it may lead to inaccuracies in assignment.	–
Feature modeling decisions				
F1	Plan area source	NSI Footprint	NSI and the footprint geometry may provide different plan area values for the same building. NSI plan area (total floor area divided by number of stories) is not tied to the selected footprint geometry, while footprint area depends on the footprint source (S1). Plan area is one input used to compute full structural replacement value if the Hazus method is adopted (F4), which scales value by floor area.	–
F2	Residential unit source	NSI units Sampled within occupancy range Census unit scaling method	Available sources and methods may provide different residential unit counts for the same building. All three methods are retained in both total and residual uncertainty because each correlates most closely with the local reference inventory, either at the individual building level or when aggregated to the census block or tract level.	–
F3	Year of construction source	NSI median Lognormal sampling about NSI median ACS year sampling	NSI provides only the median year of construction across the census tract, rather than building-specific values. To better approximate variability in year of construction within a census tract, two sampling methods are considered, though neither recreates building-level data on year of construction. Based on the bias evaluation in Figure 11, ACS year sampling is preferred.	ACS
F4	Replacement value source	Hazus NSI	Hazus and NSI are both widely used sources of replacement value for estimating the financial implications of an earthquake. Hazus replacement value further depends on the plan area source (F1) and number of stories. In both cases, this value is the full (undepreciated) value of the structure, excluding content and vehicle value.	–

Continued on next page

Table S1 continued

#	Decision	Options	Reason for inclusion	Branch after resolving bias
F5	Use NSI building material	Yes No	NSI building material (wood, masonry, concrete, steel, or manufactured housing) can be used to constrain Hazus structure-type assignment (Lochhead et al. Lochhead et al. [2], Figure 7, Potential Modification 1). However, NSI building materials are themselves inferred, so constraining structure type with them may add information or propagate error.	–
F6	Region-specific structural system inference modifications	Yes No	Hazus uses a default mapping schema to assign structural systems from a combination of other features. Lochhead et al. [2] adopts modifications to modulate structural systems using weights (see Lochhead et al. Figure 7, Potential Modification 2). This decision determines whether or not to adopt these proposed modulations or to use the default Hazus mapping.	Yes
F7	Fill structural system with no valid Hazus match	Yes No	Some combinations of occupancy, height, and era have no compatible Hazus structural system. Removing these buildings from the inventory loses exposure, while retaining them requires an arbitrary structure-type assignment.	–
F8	Manual mobile home polygons	Yes No	National datasets can misclassify mobile home parks, which are particularly vulnerable to earthquake damage. Manually delineated polygons correct these occupancies (following the methodology in Lochhead et al. [2]), but require local effort; many studies adopt national data without such corrections.	Yes
F9‡	Scale university population to HIFLD	Yes No	NSI daytime population on university campuses may not match campus population reported in HIFLD. Scaling assumes HIFLD better represents campus population, while using NSI population unchanged is consistent with studies that adopt raw NSI data. As with S5, not scaling is considered reasonable, but is neither recommended nor included in residual uncertainty.	–

†Considered only when S6 = enhanced spatial join.
‡Considered only when S5 = Yes.

Additional detail on spatial modeling decisions

S1. Footprint source [`footprint_name`]

- **Microsoft**: Microsoft's machine-learning-derived footprints, retrieved with the BRAILS Microsoft scraper [22].
- **OpenStreetMap**: All object names ("ways") tagged `building`, retrieved through the Overpass API using the BRAILS OpenStreetMap scraper (data timestamp 7 October 2026).
- **Overture**: Buildings from Overture Maps Foundation release 2026-09-23.1 (buildings theme).
- **USA Structures**: The FEMA USA Structures feature service, retrieved with the BRAILS USA Structures scraper.
- **Local (City of Hayward)**: City of Hayward Building Footprints feature service (last edited 7 August 2026).

All sources were downloaded in October 2026 for the 2020 Census place boundary of Hayward (BRAILS scrapers from [28]). In each source, when one footprint overlaps another by at least 70% of its own area, only the largest footprint in each connected group of overlapping footprints is kept.

S2. Minimum footprint area [`footprint_name`]

- **0 ft^2**: No footprints are removed based on area. This assumes that all footprints included in the datasets are candidates for point attribution.

- **450 ft^2**: Footprints smaller than 450 ft^2 are removed during preprocessing, matching the threshold used in USA Structures. This assumes that small footprints are not habitable (e.g., sheds, garages) and thus should be removed.

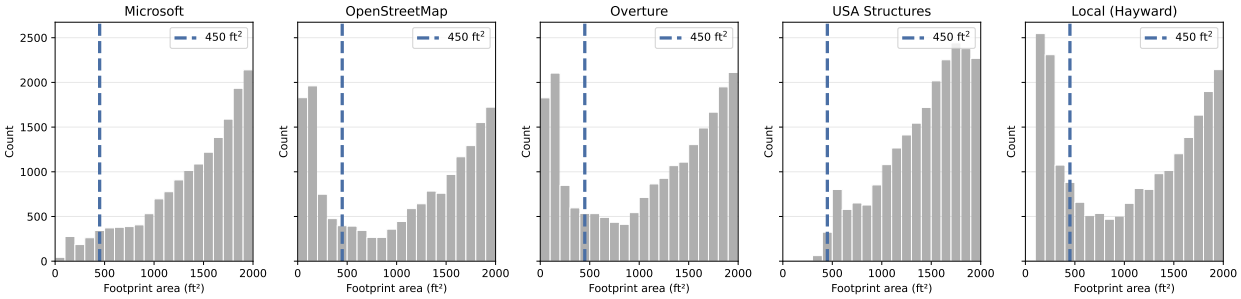


Fig. S1. Histograms of building footprint areas between 0 and 2,000 ft^2 for each footprint source, relative to optional minimum footprint area filter of 450 ft^2.

S3. Minimum NSI floor area [`nsi_filter_limit`]

- **0 ft^2**: No NSI points are removed. This assumes that small floor areas reflect errors in area assignment of habitable buildings rather than small, uninhabitable buildings.
- **450 ft^2**: Points with total floor area below 450 ft^2 are removed during preprocessing. This assumes these points do not represent real habitable buildings.

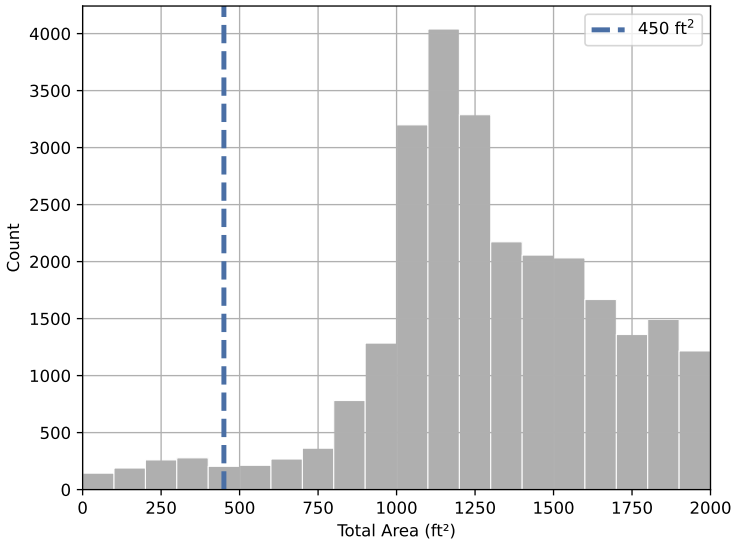


Fig. S2. Histogram of total area of NSI points between 0 and 2,000 ft^2, relative to optional minimum area filter. Approximately 1000 points fall below 450 ft^2

S4. Merge NSI points with duplicate building IDs [`merge_duplicate_bid`]

- **Yes**: Before attribution, NSI points sharing a BID are combined into one record located at the first point in the group. Additive fields (e.g., population, value, floor area, units) are summed. Other fields (e.g., building material, stories, year of construction) are kept as a single value if identical and as a list otherwise. Residential occupancy classes (RES1, RES3) are combined into one class from the sum of each class's midpoint unit count (e.g., two RES1 points become RES3A). Any non-residential classes are kept alongside,

producing a mixed-use record. Points with no BID are not merged. The NSI documentation is not clear on the relationship between points with the same BID.

- **No**: Each NSI point enters attribution as a separate record.

S5. Augment NSI using HIFLD data [`include_hifld`]

- **Yes**: Three groups of HIFLD data are integrated with NSI during preprocessing, following the occupancy class naming conventions of NSI.
  - *K–12 schools:* HIFLD public and private schools are added as EDU1 records, with the HIFLD population as daytime population and zero nighttime population. Each HIFLD school is paired with the nearest NSI EDU1 point within 50 m. The paired record keeps the NSI point's location and attributes, takes the HIFLD daytime population, and keeps the NSI nighttime and over-65 daytime population. Remaining NSI EDU1 points more than 150 m from every HIFLD school are removed, and NSI GOV1 points within 150 m of a HIFLD school are reclassified as EDU1.
  - *Colleges and universities:* HIFLD university points and campus polygons within Hayward's Census blocks are used. University points outside every campus polygon, and university points inside campus polygons that contain no NSI EDU1 or GOV1 points, are added as EDU2 records. Their daytime population comes from HIFLD (split 99.5% under 65 and 0.5% over 65), and their nighttime population is 0.5% of daytime population, following the Hazus inventory convention FEMA [27]. NSI GOV1 and EDU1 points inside campus polygons are reclassified as EDU2.
  - *Emergency response facilities:* HIFLD fire and EMS stations, local law enforcement locations, and local and state emergency operations centers are added as GOV2 records. Each HIFLD facility is paired with the nearest NSI GOV2 point within 50 m. The paired record keeps the NSI point's location and attributes with the HIFLD occupancy class, and the paired NSI point is removed. Remaining NSI GOV2 points more than 50 m from every HIFLD facility are reclassified as GOV1, and NSI GOV1 points within 10 m of a HIFLD facility are removed.
- **No**: NSI points are used without these additions or modifications.

S6. Point-to-footprint attribution method [`point_attribution_method`]

- **Enhanced spatial join**: Points are attributed to footprints per the method described in Lochhead et al. [2].
- **Simple spatial join**: Each footprint is joined with the points that intersect it. When more than one point intersects a footprint, only the first point in join order is kept, and the others are discarded. Points outside all footprints are discarded.
- **Nearest neighbor**: Each point is assigned to its nearest footprint (distance is zero for points inside a footprint), with no distance limit. When more than one point is assigned to the same footprint, only the closest point is kept, and the others are discarded.

S7. Allow points to move between census blocks during attribution [`use_surrounding_bgs`]

- **Yes**: Candidate footprints are those in the point's census block and all adjacent blocks.
- **No**: Candidate footprints are those in the point's census block only.

S8. Estimate footprint capacity using height [`estimate_stories`]

Following Lochhead et al. [2], capacity determines whether a footprint is partially full, which partially controls how attribution is conducted (S9). For each NSI occupancy class, a size limit is set at the mean plus two standard deviations of capacity among footprints containing exactly one point of that class (excluding BID-merged records, S4). A footprint is partially full when its capacity exceeds the largest limit among its record's occupancy classes. This is re-evaluated whenever the footprint receives additional points.

- **Yes**: Footprint capacity is used to identify partially full footprints, per the enhanced spatial join method in Lochhead et al. [2] (S9). The number of stories is estimated from footprint height: one story below 19 ft, plus one story for each additional 10 ft (e.g., two stories for 19–29 ft). Footprints with missing or negative height are assigned one story. Capacity is the product of footprint area by the estimated number of stories.

- **No**: Capacity is footprint area.

S9. Prioritize partially full footprints [`prioritize_partial_footprints`]

- **Yes**: Within each block, after points have been attributed to intersecting footprints and to empty footprints within the distance bound, each remaining point within the bound of a partially full footprint is merged into the record of its nearest partially full footprint (using the S4 field rules). This repeats until no remaining point is within the bound of a partially full footprint not yet updated in this stage.
- **No**: This stage is skipped.

S10. Allow attribution to any footprint [`use_full_ftpt`]

- **Yes**: Within each block, after the preceding stages, each remaining point within the distance bound of any candidate footprint (empty, partially full, or occupied) is merged into the record of its nearest footprint in a single pass.
- **No**: This stage is skipped.

S11. Final attribution distance bound [`final_distance_bound`]

Distance is measured from a point to the nearest edge of a footprint. Points still unattributed after the final round are removed from the inventory.

- **10 m**: Only the 10 m round of distance-limited attribution is run.
- **100 m**: After the 10 m round, attribution following S9 and S10 is repeated with a 100 m bound for points that are still unattributed.

Additional detail on feature modeling decisions

When multiple points are attributed to a single footprint, fields kept as lists (e.g., number of stories, median year of construction, building type) are reduced to their most frequent value, with ties broken at random. After applying the decisions below, missing values for number of stories, year of construction, number of units, and daytime and nighttime population are filled using the BRAILS k-nearest-neighbor imputer.

F1. Plan area source [`area_source`]

- **NSI**: NSI total area divided by NSI number of stories.
- **Footprint**: Footprint polygon area.

Plan area is used to compute Hazus replacement vallue (F4) and sampled RES3F unit counts (F2).

F2. Residential unit source [`unit_source`]

- **NSI units**: The number of units reported in NSI, summed across points attributed to each footprint.
- **Sampled within occupancy range**: An integer drawn uniformly from the record's unit range, which is implied by occupancy class (RES1: 1; RES3A: 2; RES3B: 3–4; RES3C: 5–9; RES3D: 10–19; RES3E: 20–50) and summed across points attributed to each footprint. RES3F has no upper bound, so its units are total floor area divided by a floor area per unit drawn from a lognormal distribution (median 797 ft^2, $\sigma = 0.34$), with a minimum of 51 units per RES3F point. The parameters were estimated from RES3F points in the Hayward NSI data, excluding values below 200 ft^2 per unit.
- **Census unit scaling method**: The method developed in Lochhead et al. [2].

In all options, RES2 records have one unit and non-residential records have none. After units are assigned, each RES1 and RES3 record's occupancy class is reset from its unit count (RES1: ≤1; RES3A: 2; RES3B: 3–4; RES3C: 5–9; RES3D: 10–19; RES3E: 20–50; RES3F: >50), and mixed-use classes are reset the same way.

F3. Year of construction source [`sample_year`]

- **NSI median**: Year of construction is the NSI median year, which NSI reports at the census tract level.

- **Lognormal sampling about NSI median**: To estimate year of construction, an age is drawn from a lognormal distribution with median equal to the median age of the building (2026 minus the NSI median year) and $\sigma = 0.20$. Year of construction is 2026 minus the sampled age.
- **ACS year sampling**: Years are drawn using the ACS year-built sampler method included in BRAILS (*acs_year_built_sampler*) [28].

F4. Replacement value source [`cost_source`]

Throughout this study, replacement value is defined as the full (undepreciated) structure value, excluding content and vehicle value. Both sources below characterize this value.

- **Hazus**: Full structural replacement value is computed using the Hazus assumed cost per square foot for the record's occupancy class [27] multiplied by floor area, which is plan area (F1) times the NSI number of stories (one story if missing).
- **NSI**: Full structural replacement value is taken from NSI's *fullrep* field in the 2026 NSI release.

F5. Use NSI building material [`use_nsi_bldg_type`]

- **Yes**: Candidate structural systems are limited to those consistent with the record's NSI building material (Lochhead et al. [2], Figure 7, Potential Modification 1). Records without an NSI material are sampled per the Hazus Building Stock Mapping Scheme Tables [27]. The 2022 NSI release had significant material inaccuracies (e.g., over-representation of masonry in California housing), but the 2026 release has fewer clear biases.
- **No**: The structural system is sampled from the Hazus Building Stock Mapping Scheme Tables FEMA [27] without considering NSI material (Lochhead et al. [2], Figure 7, without Potential Modification 1.)

F6. Region-specific structural system inference modifications [`bldg_type_modifications`]

- **Yes**: Three changes are made to the default mapping (Lochhead et al. [2], Figure 7, Potential Modification 2), with weights renormalized: (i) mobile home (MH) is removed as a candidate for all occupancy classes except RES2; (ii) unreinforced masonry (URM) is removed for all occupancy classes, reflecting retrofit efforts in Hayward; and (iii) RES3A and RES3B use the RES1 mapping, as their construction is assumed to resemble single-family homes more than larger multi-unit buildings.
- **No**: The default Hazus mapping is used (Lochhead et al. [2], Figure 7, without Potential Modification 2).

F7. Fill structure system with no valid Hazus match [`fill_in_error_structure_types`]

This decision applies to buildings for which the mapping (after F5 and F6) gives no possible structural system.

- **Yes**: A structural system is drawn uniformly at random from W1, W2, S1–S5, C1–C3, PC1, PC2, RM1, and RM2 (excluding URM and MH). URM is excluded due to retrofit efforts in Hayward, and MH is excluded because the mobile home structural system (MH) is reserved for the mobile home occupancy (RES2).
- **No**: These buildings are removed from the inventory.

F8. Manual mobile home polygons [`include_manual_mh_polygons`]

- **Yes**: Lochhead et al. [2] found significant NSI bias for mobile and manufactured housing and recommended manually drawn mobile home park polygons as a feasible city-level correction. Footprints whose centroids fall within these polygons are assigned RES2.
- **No**: No mobile home corrections are made.

F9. Scale university population to HIFLD [`scale_edu2_pop`]

- **Yes**: Lochhead et al. [2] observed that the total daytime population of NSI points on university campuses was lower than the HIFLD campus population. The daytime population of every point inside each HIFLD campus polygon is therefore scaled by the ratio of HIFLD campus population to the points' summed daytime population. EDU2 daytime population is then split 99.5% under 65 and 0.5% over 65, and nighttime population is set to 0.5% of daytime population, following the Hazus convention FEMA [27].
- **No**: NSI populations are not modified.

S1.2 Exposure Modeling Decisions characterizing Exposure Models A and B

To evaluate resolvable biases in inventory generation, this study uses two exposure models (each defined as distinct outcomes from a suite of decisions) that are intentionally developed to represent significant differences in the remaining decisions. As described in the main text, *Exposure Model A* selects options informed by the original NSI data when possible: it retains NSI points by attributing them broadly to footprints and takes plan area, residential units, replacement value, and building material from NSI. *Exposure Model B* selects options that are not informed by the original NSI data: it removes small NSI points and footprints, discards NSI points that cannot be attributed under the most limited attribution options, and uses footprint plan area, units sampled within the occupancy range, Hazus replacement value, and the Hazus structure-type mapping without NSI building material. Table S2 lists the option selected for each decision in both models.

Decisions evaluated in the bias assessment (S1, S4, S5, S6, F3, F6, and F8) are not part of either model definition and are marked with an asterisk in Table S2. The options for these decisions are evaluated with both Exposure Models A and B to assess whether one branch resolves bias in both.

TABLE S2. Decisions characterizing Exposure Models A and B. Decisions marked * are evaluated as part of the bias assessment and thus are not included in model definitions (see Section 5.1).

Decision	Exposure Model A	Exposure Model B
Spatial modeling decisions		
S1. Footprint source	*	*
S2. Minimum footprint area	0 ft^2	450 ft^2
S3. Minimum NSI floor area	0 ft^2	450 ft^2
S4. Merge NSI points with duplicate building IDs	*	*
S5. Augment NSI using HIFLD data	*	*
S6. Point-to-footprint attribution method	*	*
S7. Allow points to move between census blocks during attribution†	Yes	No
S8. Estimate footprint capacity using height†	Yes	No
S9. Prioritize partially full footprints†	Yes	No
S10. Allow attribution to any footprint†	Yes	No
S11. Final attribution distance bound†	100 m	10 m
Feature modeling decisions		
F1. Plan area source	NSI	Footprint
F2. Residential unit source	NSI units	Sampled within occupancy range
F3. Year of construction source	*	*
F4. Replacement value source	NSI	Hazus
F5. Use NSI building material	Yes	No
F6. Region-specific structure-type modifications	*	*
F7. Fill structure types with no valid Hazus match	Yes	No
F8. Manual mobile home polygons	*	*
F9. Scale university population to HIFLD‡	Yes	No

†Considered only when S6 = enhanced spatial join.
‡Considered only when S5 = Yes.

S1.3 Comprehensive Bias Results

Section 5.1 demonstrates bias resolution using either Exposure Model A or B (Figures 6 through 10). In this section, the figures are repeated for both exposure models A and B to demonstrate that the preferred branch resolves bias in both exposure models (i.e., across residual uncertainty). The magnitude of the bias and the impact of bias resolution are not equivalent across Exposure Models A and B, but they do show consistent effects.

Footprint Source and Attribution Method

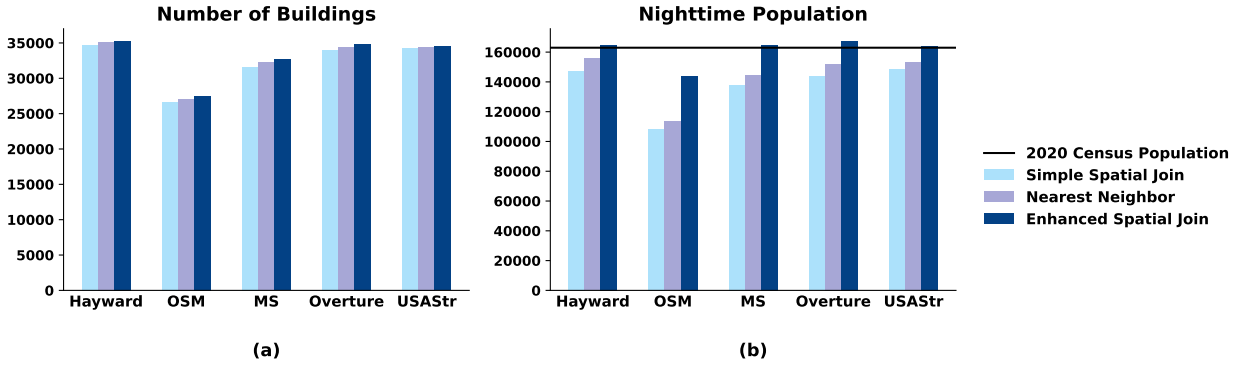


Fig. S3. (a) Building counts and (b) nighttime population for combinations of footprint source and attribution method using Exposure Model A, compared with a reference inventory with local data.

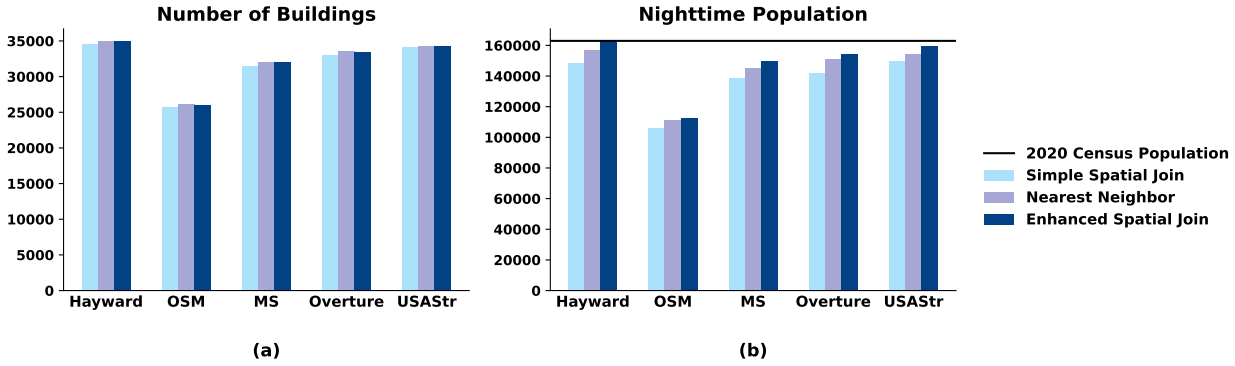


Fig. S4. (a) Building counts and (b) nighttime population for combinations of footprint source and attribution method using Exposure Model B, compared with a reference inventory with local data.

Mobile and Manufactured Homes

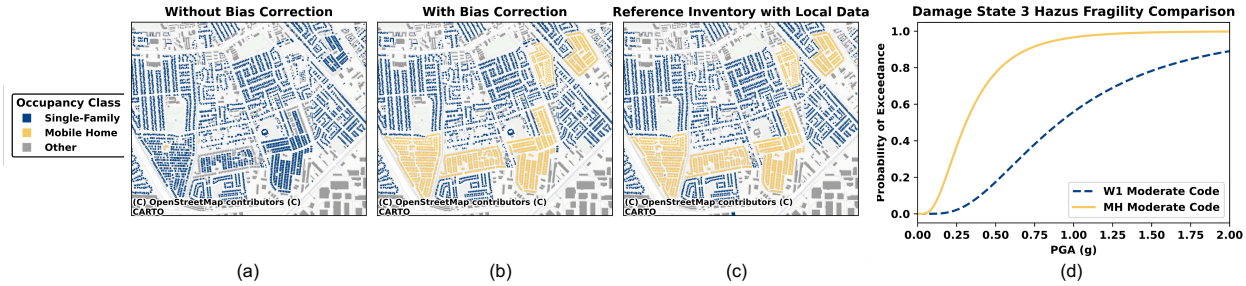


Fig. S5. Resolving bias using manual mobile home polygons, illustrated using Exposure Model A. (a) Inventory without bias correction, (b) inventory with bias correction, which better aligns with (c) reference inventory with local data. (d) Fragility curves showing resulting difference in estimated seismic risk.

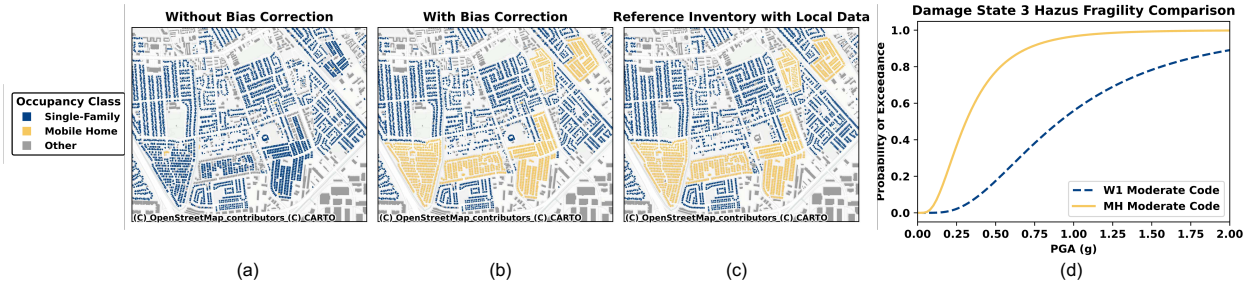


Fig. S6. Resolving bias using manual mobile home polygons, illustrated using Exposure Model B. (a) Inventory without bias correction, (b) inventory with bias correction, which better aligns with (c) reference inventory with local data. (d) Fragility curves showing resulting difference in estimated seismic risk.

Structural System Inference

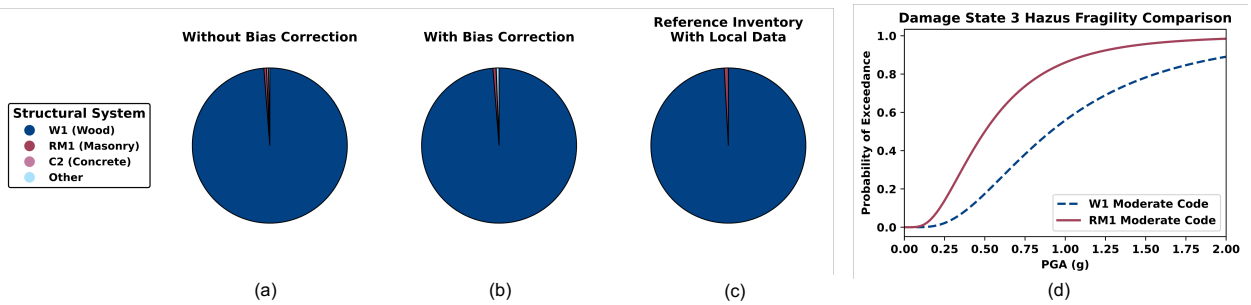


Fig. S7. Resolving bias using adjusted structural system inference, illustrated using Exposure Model A. (a) Inventory without bias correction, (b) inventory with bias correction, which better aligns with (c) reference inventory with local data. (d) Fragility curves showing resulting difference in estimated seismic risk.

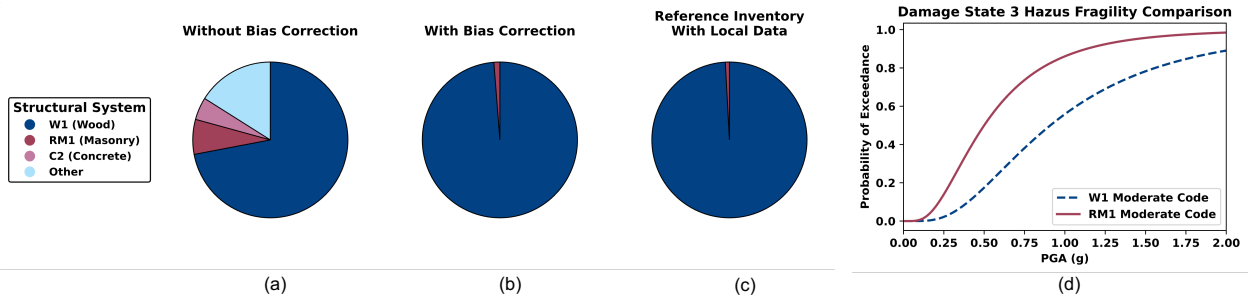


Fig. S8. Resolving bias using adjusted structural system inference, illustrated using Exposure Model B. (a) Inventory without bias correction, (b) inventory with bias correction, which better aligns with (c) reference inventory with local data. (d) Fragility curves showing resulting difference in estimated seismic risk.

Year of Construction

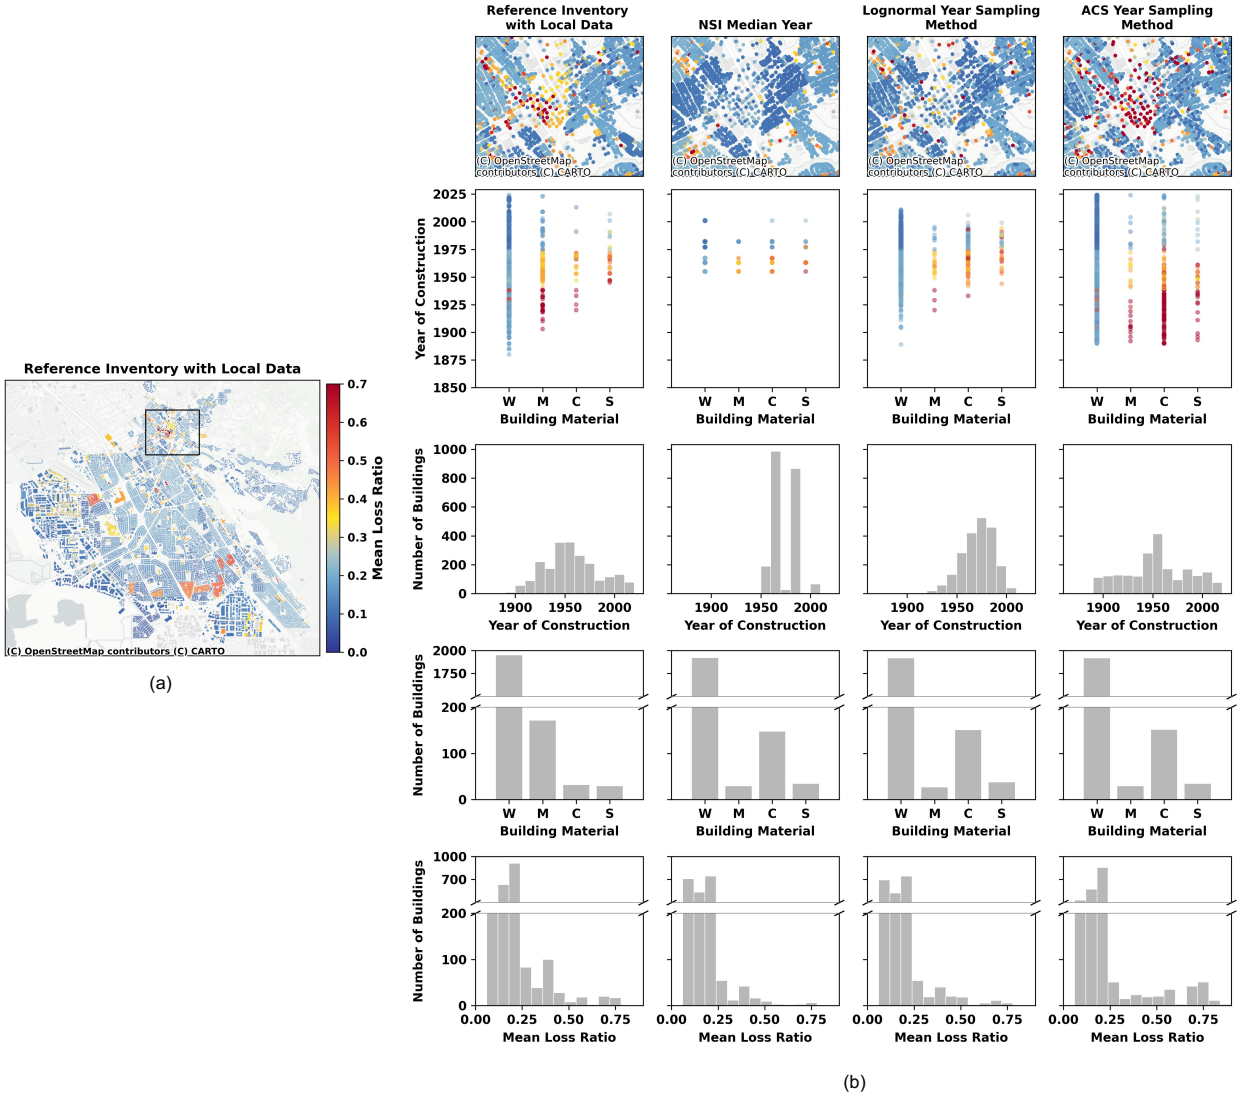


Fig. S9. Year of construction branch comparison for Exposure Model A. (a) Mean loss ratio across Hayward for the reference inventory, with the selected neighborhood boxed. (b) Neighborhood results for the reference inventory and three candidate branches, with rows (in descending order) showing mean loss ratio by footprint, construction year by building material colored by mean loss ratio, histograms of year of construction, bar charts of building material, and histograms of mean loss ratio. Mean loss ratio (map and scatter) share the color scale in (a). W, M, C, and S denote wood, masonry, concrete, and steel.

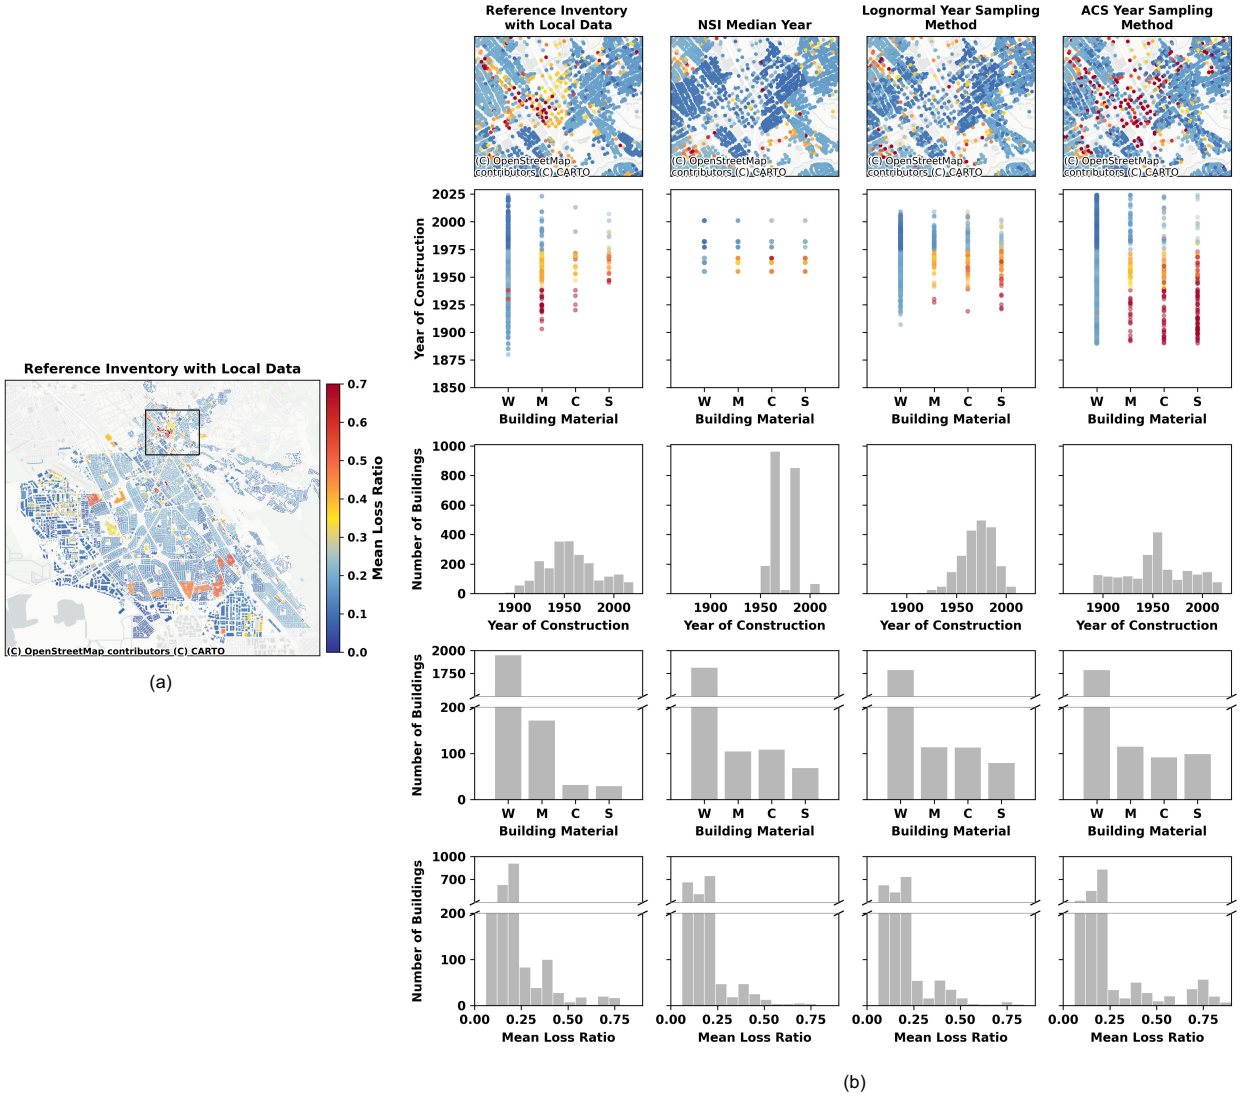


Fig. S10. Year of construction branch comparison for Exposure Model B. (a) Mean loss ratio across Hayward for the reference inventory, with the selected neighborhood boxed. (b) Neighborhood results for the reference inventory and three candidate branches, with rows (in descending order) showing mean loss ratio by footprint, construction year by building material colored by mean loss ratio, histograms of year of construction, bar charts of building material, and histograms of mean loss ratio. Mean loss ratio (map and scatter) share the color scale in (a). W, M, C, and S denote wood, masonry, concrete, and steel.

S1.4 Comprehensive Damage and Loss Results across Metrics

Using the methodology outlined in Equations 2 and 3, a suite of potential damage metrics was generated, with each representing a unique combination of exceedance-probability thresholds and damage-state thresholds. The authors then selected three damage metrics for use in the study. The following figures display the values for the following metrics across inventory versions of Hayward, allowing for a comparison of the total number of buildings or units falling into each damage metric and the level of residual uncertainty. Metrics in blue were omitted from the study due to redundancy, insufficient coverage of the Hayward inventory, or relatively lower impact on decision-making for regional disaster mitigation or resilience measures.

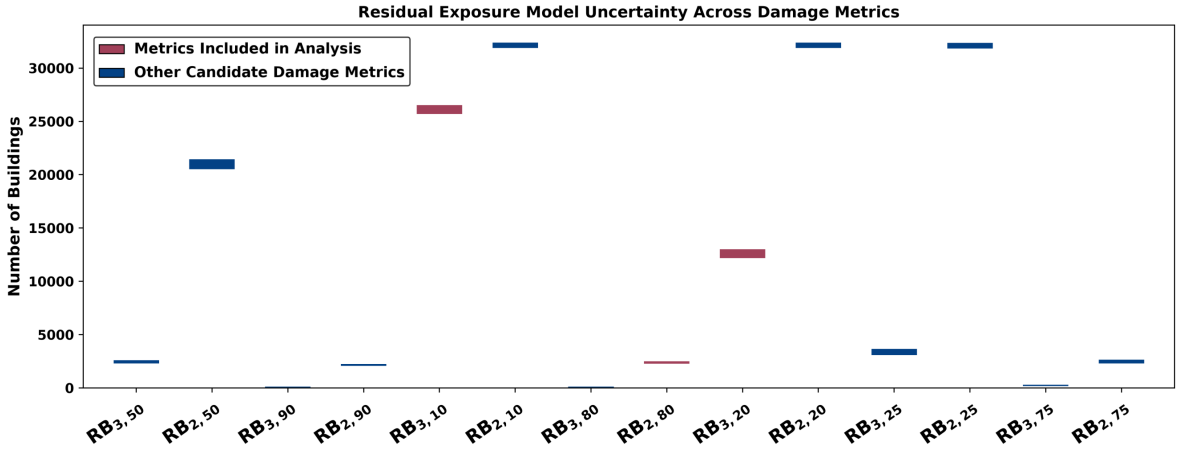


Fig. S11. Complete suite of damage metrics to describe the number of buildings in the Hayward inventory that were generated in the study. The three selected for analysis are shown in red.

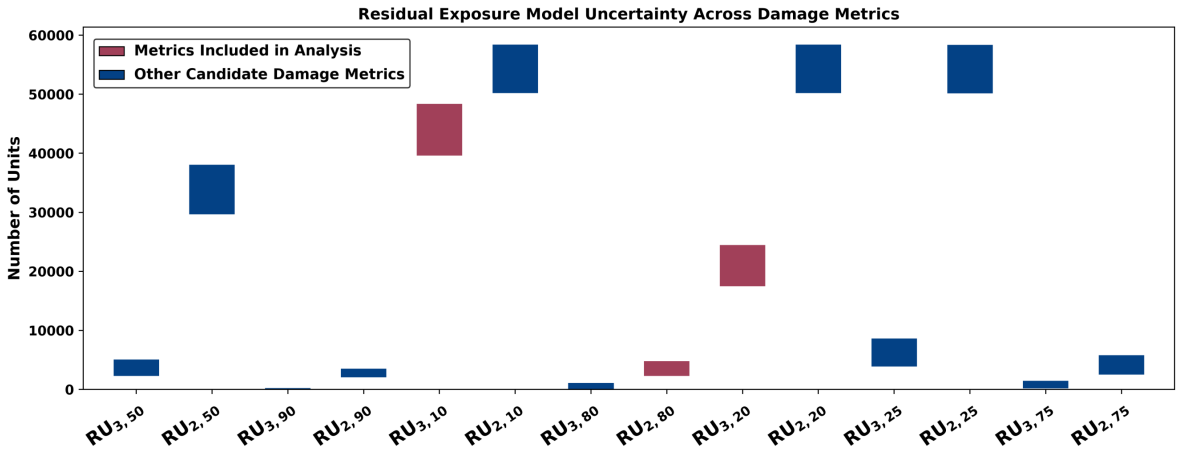


Fig. S12. Complete suite of damage metrics to describe the number of residential units in the Hayward inventory that were generated in the study. The three selected for analysis are shown in red.

REFERENCES

[1] FEMA. Hazus 7.0 Inventory Technical Manual. Technical report, Federal Emergency Management Agency, 2025.

[2] Meredith Lochhead, Ádám Zsarnóczay, and Gregory G Deierlein. Exposure matters: A synthesis framework for high-resolution building inventory development. *International Journal of Disaster Risk Reduction*, page 106148, 2026.

[3] Charles K Huyck, Zhenghui Hu, Michael Eguchi, Georgiana Esquivias, Paul Amyx, Kay Smith, and Colm Jordan. Characterizing uncertainty of general building stock exposure data. *Earthquake Spectra*, 38(3):2008–2025, 2022.

[4] Dylan Sanderson and Daniel Cox. Comparison of national and local building inventories for damage and loss modeling of seismic and tsunami hazards: From parcel-to city-scale. *International journal of disaster risk reduction*, 93:103755, 2023.

[5] Petros Kalakonas, Vitor Silva, Amaryllis Mouyiannou, and Anirudh Rao. Exploring the impact of epistemic uncertainty on a regional probabilistic seismic risk assessment model. *Natural Hazards*, 104(1):997–1020, 2020.

[6] Nirmal Jayaram and Jack W Baker. Correlation model for spatially distributed ground-motion intensities. *Earthquake Engineering & Structural Dynamics*, 38(15):1687–1708, 2009.

[7] Julian J Bommer, Frank Scherbaum, Hilmar Bungum, Fabrice Cotton, Fabio Sabetta, and Norman A Abrahamson. On the use of logic trees for ground-motion prediction equations in seismic-hazard analysis. *Bulletin of the Seismological Society of America*, 95(2):377–389, 2005.

[8] Morgan P Moschetti, Brad T Aagaard, Sean K Ahdi, Jason Altekruse, Oliver S Boyd, Arthur D Frankel, Julie Herrick, Mark D Petersen, Peter M Powers, Sanaz Rezaeian, et al. The 2023 us national seismic hazard model: Ground-motion characterization for the conterminous united states. *Earthquake Spectra*, 40(2):1158–1190, 2024.

[9] Keith Porter, Robert Kennedy, and Robert Bachman. Creating fragility functions for performance-based earthquake engineering. *Earthquake spectra*, 23(2):471–489, 2007.

[10] Adam Zsarnóczay, Gregory G Deierlein, Frank McKenna, Matthew Schoettler, Sang-Ri Yi, Barbaros Cetiner, Aakash Bangalore Satish, Jinyan Zhao, Justin Bonus, Abiy F Melaku, et al. An open-source simulation platform to support and foster research collaboration in natural hazards engineering. *Frontiers in Built Environment*, 11: 1590479, 2025.

[11] Federal Emergency Management Agency. Seismic performance assessment of buildings, volume 2 – implementation guide. Technical Report FEMA P-58-2, Federal Emergency Management Agency, Washington, D.C., December 2018. Prepared by the Applied Technology Council for the Federal Emergency Management Agency.

[12] Vitor Silva, Sinan Akkar, Jack Baker, Paolo Bazzurro, José Miguel Castro, Helen Crowley, Matjaz Dolsek, Carmine Galasso, Sergio Lagomarsino, Ricardo Monteiro, et al. Current challenges and future trends in analytical fragility and vulnerability modeling. *Earthquake Spectra*, 35(4):1927–1952, 2019.

[13] Chenhao Wu and Henry Burton. Quantifying exposure information uncertainty in regional risk assessment. *arXiv preprint arXiv:2605.08272*, 2026.

[14] Juan Camilo Gómez Zapata, Massimiliano Pittore, Fabrice Cotton, Henning Lilienkamp, Simantini Shinde, Paula Aguirre, and Hernán Santa María. Epistemic uncertainty of probabilistic building exposure compositions in scenario-based earthquake loss models. *Bulletin of Earthquake Engineering*, 20(5):2401–2438, 2022.

[15] Zihao Song, Wenjun Jiang, Sang-ri Yi, and Jize Zhang. Reliable building inventory imputation for regional-scale risk assessment: An uncertainty-guided framework using spatially-enhanced transformers. *Reliability Engineering & System Safety*, 274:112436, 2026.

[16] Saeid Ghasemi, Milad Roohi, Dylan Sanderson, Omar Sediek, and John W van de Lindt. Integrated data imputation of national inventories for bridging information gaps and outcome risk uncertainty in community resilience modeling. *International Journal of Disaster Risk Reduction*, page 106077, 2026.

[17] U.S. Army Corps of Engineers. National structure inventory, 2022. URL https://www.hec.usace.army.mil/confluence/nsi/technicalreferences/latest/technical-documentation.

[18] David M Boore, Jonathan P Stewart, Emel Seyhan, and Gail M Atkinson. Nga-west2 equations for predicting pga, pgv, and 5% damped psa for shallow crustal earthquakes. *Earthquake Spectra*, 30(3):1057–1085, 2014.

[19] Eric M Thompson. An updated vs30 map for california with geologic and topographic constraints. *US Geological Survey (USGS) Data Release*, page 75, 2018.

[20] A. Zsarnoczay, T. Lee, and I. Vouvakis Manousakis. NHERI-SimCenter/DamageAndLossModelLibrary: v3.1, 2026. URL https://doi.org/10.5281/zenodo.7951657.

[21] Adam Zsarnoczay, Ioannis Vouvakis Manousakis, and Jinyan Zhao. Nheri-simcenter/pelicun: v3.10 (version v3.10), 2026. URL https://doi.org/10.5281/zenodo.21812989.

[22] Microsoft. Global ML building footprints. GitHub repository, 2026. URL https://github.com/microsoft/GlobalMLBuildingFootprints. Version: 2026-02-03.

[23] OpenStreetMap Contributors. Building Footprints. https://www.openstreetmap.org. Accessed: 2026-10-07.

[24] Federal Emergency Management Agency (FEMA) and Oak Ridge National Laboratory (ORNL). USA Structures. https://gis-fema.hub.arcgis.com/pages/usa-structures. Version: 2025-09-23.

[25] Overture Maps Foundation. Overture Maps Buildings Theme. https://docs.overturemaps.org/. Version: 2026-09-23.1.

[26] Hayward Open Data. Building Footprints. City of Hayward Open Data (ArcGIS Hub). URL https://opendata.hayward-ca.gov/datasets/building-footprints. Accessed: 2026-10-07.

[27] FEMA. Hazus 6.1 Inventory Technical Manual. Technical report, Federal Emergency Management Agency, 2024.

[28] Barbaros Cetiner, Frank McKenna, Sang-ri Yi, Bingyu Wang, and Ioannis V. Manousakis. BRAILS++, 2026. URL https://doi.org/10.5281/zenodo.20350119. Version v4.3.0.

[29] Norman A. Abrahamson, Walter J. Silva, and Ronnie Kamai. Summary of the ASK14 ground motion relation for active crustal regions. *Earthquake Spectra*, 30(3):1025–1055, 2014. doi: 10.1193/070913EQS198M.

[30] Kenneth W. Campbell and Yousef Bozorgnia. NGA-West2 ground motion model for the average horizontal components of PGA, PGV, and 5% damped linear acceleration response spectra. *Earthquake Spectra*, 30(3): 1087–1115, 2014. doi: 10.1193/062913EQS175M.

[31] Brian S.-J. Chiou and Robert R. Youngs. Update of the Chiou and Youngs NGA model for the average horizontal component of peak ground motion and response spectra. *Earthquake Spectra*, 30(3):1117–1153, 2014. doi: 10.1193/072813EQS219M.