Internalizing Fall Risk in Housing Design: A Probabilistic Optimization Framework with Medicare Advantage Implications

Authors: Nilesh Salvi, Subrahmanya Chandra Bhamidipati, Alisha Johnson, Prasad Calyam

Abstract

Falls among adults aged 65 and older are a leading source of injury, hospitalization, and Medicare expenditure in the United States. Although Universal Design (UD) and building codes reduce environmental hazards, they do not minimize fall risk within the feasible space of compliant housing configurations. This paper develops a tractable probabilistic optimization framework that embeds fall risk directly into housing design decisions and translates risk reductions into expected payer cost savings. Focusing on a high-frequency exposure pathway, the nighttime bedroom-to-bathroom route, we construct a lightweight environmental digital twin, estimate mechanism-based per-traversal fall probability (slip, trip, misstep), and evaluate 200 feasible configurations under parameter uncertainty.

The optimization incorporates a construction and assembly cost layer, creating an economically relevant frontier between safety gains and capital intensity. In the central calibration, the optimized configuration reduces per-traversal fall probability by approximately 20 percent, lowering annual nighttime fall probability from roughly 24 percent to 19–20 percent. At an assumed $25,000 cost per injurious fall, this implies approximately $1,100–$1,200 in median expected annual savings per resident attributable to a single exposure path, with positive net savings after amortizing incremental build cost over 10 years.

Illustrative path configurations demonstrate that substantial risk reductions can be achieved through lighting continuity, friction upgrades, and threshold removal without major geometric reconfiguration, while structural layout changes become economically justified only when environmental levers are exhausted. Scaled to facility portfolios, projected savings exceed $0.1 million annually for 100 units and $1 million for 1,000 units. The framework provides a modular, AI-compatible mechanism for internalizing fall-related externalities in residential capital investment without relying on high-fidelity architectural digital twins.

1. Introduction

Falls among older adults represent one of the most persistent and economically consequential sources of injury in the United States. Approximately one in four adults aged 65 and older experiences a fall each year, and falls are a leading cause of injury-related emergency department visits and hospitalizations in this population (CDC 2024; Haddad et al. 2019; Reider et al. 2024). Beyond acute injury, falls frequently trigger a cascade of adverse outcomes, including functional decline, fear of falling, reduced mobility, increased dependency, and transition to institutional care(Mathon et al., 2017; Meyer et al., 2020). These downstream effects amplify long-term healthcare expenditures and reduce quality of life (Reider et al. 2024; Haddad et al. 2019; Jiang et al. 2024).

From a payer perspective—particularly within Medicare Advantage (MA)—fall-related injuries represent a material, partially preventable cost center. Medicare and MA plans finance a substantial share of acute care, post-acute rehabilitation, and long-term services associated with fall injuries (Haddad et al. 2019; Florence et al. 2018). As MA enrollment expands and value-based payment contracts increasingly reward reductions in avoidable utilization, insurers have growing incentives to reduce environmental contributors to injury risk rather than solely financing downstream treatment. Yet most fall-prevention policy remains focused on behavioral interventions, monitoring technologies, pharmaceuticals, or post-injury rehabilitation(Mathon et al., 2017). Structural determinants embedded in housing design are rarely treated as part of the insurer's optimization problem.

Environmental features play a measurable and often underappreciated role in fall risk, particularly within recurrent, high-frequency exposure contexts. Lighting discontinuities, surface transitions, low-friction flooring, stair geometry, and path complexity increase the probability that a momentary mobility perturbation results in injury. These risks are particularly salient during nighttime bedroom-to-bathroom trips, which combine low illumination, sleep inertia, urgency, and transient instability (Thölking et al. 2020). Because such exposures occur repeatedly, often nightly, even small reductions in per-traversal fall probability may generate meaningful annual reductions in injury incidence and payer cost.

Housing design decisions are durable capital investments with long-lived safety implications. Once a layout is constructed, geometry, path length, and transition structure remain fixed for decades. This durability creates both opportunity and inertia. If risk-reducing design features are embedded at the construction stage, they may yield long-term cost savings that accrue to insurers and society. If they are not, structural risk persists. Universal Design (UD) and building codes establish minimum standards for safety and accessibility, but they do not rank or minimize risk within the compliant design space. Multiple UD-compliant layouts may differ substantially in lighting continuity, path complexity, and threshold frequency, yet all satisfy regulatory requirements(Carr et al., 2013; Mustaquim, 2015).

This paper asks whether fall risk can be internalized into the housing design objective function itself and aligned with payer incentives. We develop a tractable probabilistic framework that maps environmental design parameters to per-traversal fall probability, annualizes risk through repeated exposure, and translates risk into expected Medicare-relevant cost(McDonough et al., 2017; Shumway-Cook et al., 2009). The central economic insight is that fall risk is partially determined by durable capital configuration, while injury costs are largely internalized by Medicare Advantage plans. This separation creates a classic externality: builders optimize construction cost and market demand, whereas payers internalize downstream injury expenditure.

The framework is intentionally AI-compatible in a limited and pragmatic sense. The environmental representation is structured and machine-readable, the probabilistic model

functions as a replaceable surrogate, and the optimization problem is solver-agnostic. These properties allow technical refinement—improved estimation, calibration, uncertainty modeling, or search efficiency—without requiring construction of complex, high-fidelity architectural digital twins or end-to-end deep learning systems. The environmental encoding, risk model, and optimization engine are modular components, enabling future extensions (e.g., graph-based spatial encoding or physics-informed refinement) without altering the underlying economic objective.

2. Related Literature

2.1 Environmental Determinants of Falls

A substantial body of public health and gerontology research identifies environmental contributors to fall incidence, including inadequate lighting, irregular flooring, raised thresholds, cluttered pathways, and stair design (Centers for Disease Control & Prevention, 2024; Haddad et al., 2019; Lu et al., 2019; Reider et al., 2024) Interventions targeting these factors, such as improved illumination or removal of tripping hazards, have been shown to reduce fall rates in specific contexts (Grant et al., 2022; Thölking et al., 2020). However, this literature typically evaluates discrete interventions rather than continuous optimization across a multidimensional design space. As a result, the question of how to systematically minimize risk within feasible architectural configurations remains underexplored.

2.2 Universal Design and Regulatory Standards

UD and accessibility regulations aim to improve safety and usability through minimum standards for spatial dimensions, lighting, handrails, and maneuverability. From an economic standpoint, UD operates as a constraint set: it defines what is permissible but does not optimize among permissible alternatives. Two UD-compliant bedroom-to-bathroom configurations may differ meaningfully in path length, transition frequency, and lighting gradients. The existing literature rarely formalizes fall risk as a function over this feasible design space or embeds injury cost into the architectural decision problem.

2.3 AI in Fall Prevention and Monitoring

Artificial intelligence has been extensively applied to fall detection and fall-risk prediction using wearable sensors, ambient monitoring systems, and machine learning algorithms(González-Castro et al., 2024). These systems classify individuals by risk level or detect fall events in real time based on motion, posture transitions, or gait characteristics. Ambient sensing approaches, including depth cameras and radar-based monitoring deployed in residential settings, have demonstrated the feasibility of continuous, unobtrusive fall detection and longitudinal mobility assessment (Skubic et al., 2014; Stone & Skubic, 2015).

These methods have improved early detection, response time, and risk stratification. However, they operate primarily at the level of individual monitoring rather than environmental

optimization. The architectural configuration itself—such as path geometry, lighting gradients, and transition frequency—is typically treated as fixed rather than as a decision variable(Albukhari, 2025; González-Castro et al., 2024). As a result, AI applications in fall prevention have largely focused on detecting or predicting falls after environmental risk has already been embedded in the built environment.

This downstream orientation leaves a critical upstream opportunity unaddressed: the use of AI to minimize fall risk through optimization of architectural layout during the design phase. Treating environmental structure as an optimizable variable, rather than a static backdrop, represents a necessary extension of current AI-based fall prevention frameworks.

2.4 Generative Design and Economic Objectives

Generative and parametric design tools are increasingly used in architecture to optimize energy efficiency, structural performance, and construction cost(Aman & Matisziw, 2025; Suphavarophas et al., 2024). Yet health-risk probabilities are rarely incorporated as explicit objective functions. Moreover, translation of predicted risk reductions into payer-aligned economic metrics is largely absent. The integration of probabilistic risk modeling, constrained generative search, and Medicare Advantage cost translation remains underdeveloped.

This paper contributes to this intersection by embedding expected fall-related injury cost directly into the housing design optimization problem and framing the intervention explicitly in terms of incentive alignment between builders and payers.

3. Conceptual Framework

We formalize fall-related injury cost as a function of environmental design. Let D denote a vector of housing design parameters (lighting continuity, friction class, threshold presence, turn severity, path geometry). Expected annual injury-related cost (C) per resident is:

$$E[C \mid D] = P_{annual}(D)E[C_{event}] \quad \textbf{(1)}$$

Where,

$$P_{annual}(D) = 1 - (1 - P_{fall}(D))^N$$

The exponent $N = 365\lambda_{night}$ where λ_{night} is the average number of nighttime trips per night, $P_{fall}(D)$ is per-traversal fall probability conditional on design, and $E[C_{event}]$ is expected Medicare-relevant cost per injurious fall. For small P_{fall}, $P_{annual} \approx NP_{fall}$.

3.1 Private vs. Social Objective

A builder choosing design D typically minimizes construction and compliance cost: $min_D\ C_{build}(D)$ subject to $D \in$ feasible design space (UD- and code-compliant). By contrast, a Medicare Advantage plan internalizes downstream injury expenditure and would prefer to

minimize: $min_D\ E[C \mid D]$. If construction cost and injury cost are borne by different economic actors, the privately optimal design may not coincide with the cost-minimizing design from the payer's perspective. The socially efficient objective internalizes both:

$$min_D\ C_{build}(D)\ +\ E[C \mid D] \qquad \textbf{(2)}$$

This formulation makes explicit the externality embedded in durable housing capital: environmental design influences injury incidence, yet builders may not directly capture the financial benefits of reduced downstream healthcare utilization.

3.2 Constrained Optimization Within Compliance

Universal Design and regulatory standards define a feasible set of permissible configurations but do not rank alternatives within that set. The relevant problem is therefore Equation 2 subject to $D \in \mathcal{D}_{UD}$, where $\mathcal{D}_{UD}$ represents the UD- and code-compliant design space.

By expressing fall risk as an explicit function over $\mathcal{D}_{UD}$, the framework converts injury prevention from a qualitative architectural goal into a quantitative capital allocation problem. This provides a formal basis for contracting, certification, or shared-savings mechanisms that align builder incentives with Medicare Advantage cost containment objectives.

3.3 Computational Formulation of Design Problem

In this paper we approximate the objective landscape by evaluating a structured sample of 200 feasible configurations; Section 4.5 discusses how this extends to scalable search. The housing configuration vector D can be interpreted as a discrete, structured design object defined over a finite but combinatorially large feature space comprising segment-level attributes (lighting class, friction class, width proxy) and transition-level attributes (threshold class, turn severity, lighting continuity), subject to UD and code-compliance constraints. The feasible design set $\mathcal{D}_{UD}$ therefore defines a constrained combinatorial search space rather than a continuous convex domain. Even under modest feature cardinalities, the number of admissible configurations grows exponentially in the number of segments and transitions, rendering exhaustive enumeration computationally impractical as dimensionality increases. The current evaluation of 200 candidate configurations can be viewed as a structured sampling of this feasible manifold to approximate the underlying objective landscape. More generally, the problem may be cast as a solver-agnostic discrete optimization task in which the objective function combines expected injury cost and amortized construction cost, and may be approached using enumerative search, heuristic or evolutionary algorithms, Bayesian optimization, or mixed-integer programming formulations depending on scale and structural constraints. This computational framing clarifies that the framework is not limited to fixed candidate sets but is extensible to scalable design-space exploration under formal optimization paradigms.

4. Model and Methods

4.1 Study Design and Simulation Overview

This study implements a simulation-based proof-of-concept framework to evaluate whether fall risk can be reduced through structured optimization of environmental design features. The analysis focuses on a single high-frequency exposure path *the nighttime bedroom-to-bathroom route* and evaluates multiple UD–compliant configurations using a mechanism-based probabilistic model.

The workflow proceeds in four stages:

1. Structured encoding of candidate path configurations.
2. Mechanism-level risk estimation (slip, trip, misstep).
3. Aggregation into total per-traversal fall probability.
4. Translation of annualized risk into expected economic cost.

The framework is modular: environmental representation, risk modeling, and economic translation are separable components.

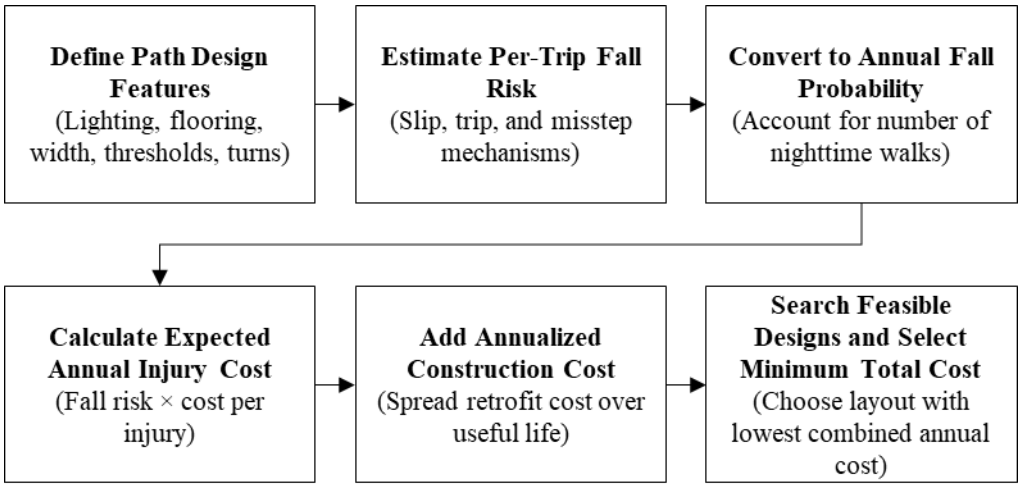


Figure 1. Workflow from Environmental Encoding to Expected Cost Minimization.

This figure illustrates the modeling pipeline: structured encoding of bedroom-to-bathroom path features; mechanism-specific hazard estimation (slip, trip, misstep); aggregation into per-traversal fall probability; annualization via repeated exposure; translation into expected annual injury cost; and constrained search over feasible designs.

Figure 1 summarizes the end-to-end pipeline and clarifies that the optimization operates on a low-dimensional, machine-readable abstraction rather than a high-fidelity geometric twin, reinforcing the model's tractability and extensibility.

4.2 Environmental Representation

The bedroom-to-bathroom route is modeled as an ordered sequence of segments (flat walking regions) and transitions (doorways, thresholds, turns). Each candidate configuration is encoded using a structured, machine-readable abstraction rather than a high-fidelity geometric twin.

Segment-Level Features

For each segment i:

- Lighting level (low, medium, high)
- Surface friction class (low, standard, high)
- Width proxy (narrow vs compliant)

Transition-Level Features

For each transition j:

- Threshold class (flush vs raised)
- Turn severity (straight, 90°, multiple turns)
- Lighting continuity (presence or absence of light drop)

A typical path includes three segments and two transitions. All sampled configurations are constrained to satisfy UD and basic code compliance requirements (e.g., minimum width, allowable threshold height limits).

4.2.1 Graph-Based Spatial Encoding

While the baseline representation models the bedroom-to-bathroom route as an ordered sequence of segments and transitions, this structure can be equivalently expressed as a graph $G = (V, E)$, where nodes V represent spatial regions (e.g., bedroom area, hallway segments, bathroom entry zone) and edges Erepresent feasible transitions between regions. In the present implementation, the route corresponds to a simple path graph with fixed ordering. Node attributes encode environmental features such as lighting level, surface friction class, and width proxy, while edge attributes encode transition-specific features including threshold presence, turn severity, and lighting discontinuity. Mechanism-specific hazard indices are computed as functions of these node and edge attributes, and per-traversal fall probability is obtained by aggregating hazard contributions along the path. This graph-based interpretation clarifies the structural organization of the environmental encoding and provides a transparent abstraction that could accommodate more complex layouts without altering the probabilistic or economic formulation used in this study.

4.3 Mechanism-Based Fall Risk Model

Instead of modeling fall risk as a single undifferentiated probability, we decompose it into three mechanisms:

1. Slip (loss of traction),
2. Trip (foot obstruction),
3. Misstep (navigation or foot placement error).

4.3.1 Hazard Indices

For each configuration, mechanism-specific hazard indices are computed as linear combinations of relevant environmental features.

- Slip hazard (H_s) is driven by low friction, low lighting, and their interaction across segments.
- Trip hazard (H_t) is driven by raised thresholds, lighting discontinuities at transitions, and their interaction.
- Misstep hazard (H_m) is driven by turn severity, lighting discontinuity, and overall path complexity.

These hazard indices are additive across segments and transitions.

4.3.2 Conversion to Probabilities

Each hazard index is mapped to a bounded per-traversal probability using a logistic link function: $p_k = sigmoid(\alpha_k + H_k), for\ k \in \{s, t, m\}$. The overall per-traversal fall probability is computed using a union approximation:

$$P_{fall} = 1 - (1 - p_s)(1 - p_t)(1 - p_m) \quad \textbf{(3)}$$

This expression is mathematically exact under the assumption that slip, trip, and misstep mechanisms occur independently; it constitutes a modeling approximation only insofar as real-world mechanisms may exhibit statistical dependence. The logistic link serves as an interpretable surrogate model; its parametric form can be replaced by nonparametric or neural function approximators without altering the optimization structure.

4.4 Calibration and Parameter Uncertainty

Model coefficients are estimated through a constrained likelihood-matching calibration procedure that aligns predicted per-traversal fall probability with epidemiological incidence benchmarks while preserving the mechanism-level structure of the model. In the reference UD–compliant but non-optimized configuration, the calibrated model yields a baseline per-traversal fall probability of approximately 0.0005. Assuming an average of 1.5 nighttime trips per night (≈550 exposures annually), this implies an annual nighttime fall probability near 0.24 attributable

to this single recurrent exposure path. This quantity reflects the probability of any fall during nighttime traversal and does not imply that all such events are injurious. The calibration targets consistency with broad population-level annual fall incidence among older adults, recognizing that the model isolates one high-frequency environmental context rather than the full set of fall environments.

To characterize model uncertainty, we implement a stochastic parameter framework. Hazard coefficients are modeled as random variables drawn from prior distributions centered at their calibrated values, with dispersion reflecting structural uncertainty (10 percent coefficient of variation in the central specification). For each sampled parameter vector, intercept terms are re-estimated to preserve baseline mechanism probabilities and maintain coherence with the annual incidence anchor.

For each configuration (baseline, optimized, and candidate designs), multiple parameter realizations are sampled and the resulting per-traversal fall probability, annualized fall probability, and expected annual cost are recomputed. Reported outcomes reflect the posterior predictive distribution induced by parameter uncertainty and are summarized using medians and interquartile ranges.

This formulation propagates uncertainty in marginal hazard effects through the full decision pipeline, yielding a stochastic design-evaluation framework that preserves interpretability while maintaining calibration discipline.

4.4.1 Uncertainty Formulation

The stochastic parameter framework implemented in this study treats mechanism-level hazard coefficients as random variables centered on calibrated baseline values. Let the vector of coefficients governing slip, trip, and misstep hazards be denoted θ. Each coefficient is sampled from a truncated normal distribution centered at its calibrated baseline value,

$$\theta_j \sim \mathcal{N}(\theta_{0j}, (CV \cdot | \theta_{0j} |)^2), \theta_j \geq 0, \quad \textbf{(4)}$$

where θ_{0j} denotes the calibrated baseline coefficient and CV specifies the relative uncertainty level (10 percent in the central specification).

For each sampled parameter vector θ, intercept terms α are recalibrated so that the baseline environmental configuration reproduces the epidemiologically anchored mechanism-specific fall probabilities. This procedure preserves consistency with annual incidence benchmarks while allowing uncertainty in marginal feature effects.

Predictive quantities viz., per-traversal fall probability, annualized fall probability, and expected injury cost, are then computed via Monte Carlo integration over sampled coefficient draws. This calibrated sampling framework propagates structural parameter uncertainty through the full

optimization pipeline, generating dispersion estimates and interquartile ranges for risk reductions and expected cost savings rather than relying on single-point estimates.

4.5 Optimization Procedure

We construct a set of 200 candidate path configurations by systematically varying lighting continuity, surface friction class, threshold presence, turn severity, width proxies, and, in selected cases, limited geometric adjustments within regulatory constraints. This sampled set is designed to approximate the feasible region of the compliant design space rather than exhaustively enumerate all admissible configurations. The objective is to characterize the structure of the constrained stochastic optimization problem and identify economically efficient solutions, rather than guarantee global optimality. In contrast to a purely risk-minimizing search, the optimization incorporates a construction and assembly cost layer, enabling evaluation of tradeoffs between safety gains and capital intensity within a unified objective framework.

For each configuration D, we compute:

$$C_{total}(D) = C_{annual}(D) + \frac{C_{build}(D)}{T} \quad \textbf{(5)}$$

where $C_{annual}(D)$ is expected annual injury cost, $C_{build}(D)$ is incremental construction cost relative to the baseline configuration, and T is an amortization horizon (10 years in the central specification).

Incremental build cost includes:

- Lighting upgrades (per segment),
- Surface friction upgrades (per segment),
- Threshold removal,
- Minor turn-severity reductions,
- A coordination/complexity penalty per change,
- A substantially larger penalty for geometry/layout modification.

Geometry changes (segment-length or routing alterations) are assigned materially higher cost weights than lighting or surface upgrades to reflect structural reconfiguration burden. The selected optimized configuration minimizes $C_{total}(D)$ rather than $P_{fall}(D)$ alone.

Because expected annual injury cost is monotonic in P_{fall} under fixed event cost assumptions, pure risk minimization would coincide with cost minimization in the absence of build cost. Introducing $C_{build}(D)$ creates an economically relevant frontier between safety gains and capital intensity.

4.6 Annualization and Economic Translation

Per-traversal fall probability is converted into annual nighttime fall probability using:

$$P_{annual} = 1 - \left(1 - P_{fall}\right)^{N} \quad \textbf{(5)}$$

where $N = 365\,\lambda_{night}$.

Expected annual injury cost per resident is then:

$$C_{annual} = P_{annual}\,E[C_{event}] \quad \textbf{(6)}$$

We assume $E[C_{event}]$ = \$25,000 for the central case; alternative values can be examined in supplementary robustness checks without altering the model structure. Facility-level projections are computed by scaling per-resident expected cost reductions linearly across unit counts. This scaling assumes independence of fall events across residents and no spillover or behavioral equilibrium effects; it should be interpreted as a first-order approximation rather than a general equilibrium prediction.

4.7 Scope and Modeling Limitations

This paper presents a structural modeling framework rather than a reduced-form empirical estimate of environmental fall risk. Hazard coefficients are simulated and calibrated to reproduce epidemiological benchmarks for annual fall incidence among older adults; they are not estimated from claims-linked architectural variation. As such, the quantitative magnitudes should be interpreted as calibrated projections under stated assumptions rather than causal estimates.

The mechanism aggregation formula assumes conditional independence across slip, trip, and misstep channels. While these mechanisms may be correlated in practice, the union approximation provides a tractable upper-bound aggregation that preserves interpretability and allows transparent decomposition of risk drivers.

The analysis is partial equilibrium. It does not model behavioral adaptation (e.g., residents turning on lights, installing assistive devices), endogenous sorting across housing types, or general equilibrium responses in insurance pricing. Event cost is represented by a central-case expected value (\$25,000), abstracting from the heavy-tailed distribution of fall severity.

Construction cost weights are stylized and intended to represent relative capital intensity (e.g., geometry changes versus lighting upgrades). These weights can be replaced with empirically derived cost schedules without altering the optimization structure. Likewise, projected savings scale proportionally with exposure frequency and assumed event cost over plausible ranges of these parameters, providing transparency regarding how alternative calibration choices would affect magnitudes.

The purpose of the framework is to demonstrate that fall risk can be embedded into a constrained capital-design objective and translated into Medicare-relevant economic terms. Empirical validation using facility-level architectural data linked to claims would strengthen external validity and represents an important direction for future research.

5. Results

5.1 Simulated Design Set, Calibration, and Uncertainty

We generated 200 feasible bedroom-to-bathroom path configurations by varying lighting continuity, minimum lighting level, friction class (with emphasis on the bathroom segment), threshold presence, turn severity, width proxy, and limited geometric adjustments. All sampled configurations satisfy UD and basic code compliance constraints.

A representative baseline configuration includes: two low-light segments and one medium-light segment; one raised threshold at bathroom entry; one 90° turn; standard hallway friction with lower friction in the bathroom segment; and moderate path length (three segments, two transitions).

Under the calibrated mechanism-based model, the baseline yields a median per-traversal fall probability of approximately 0.00050 across parameter draws. With λ_{night} = 1.5 trips per night (≈550 annual exposures), the implied median annual nighttime fall probability is approximately 0.24, with interquartile dispersion reflecting parameter uncertainty.

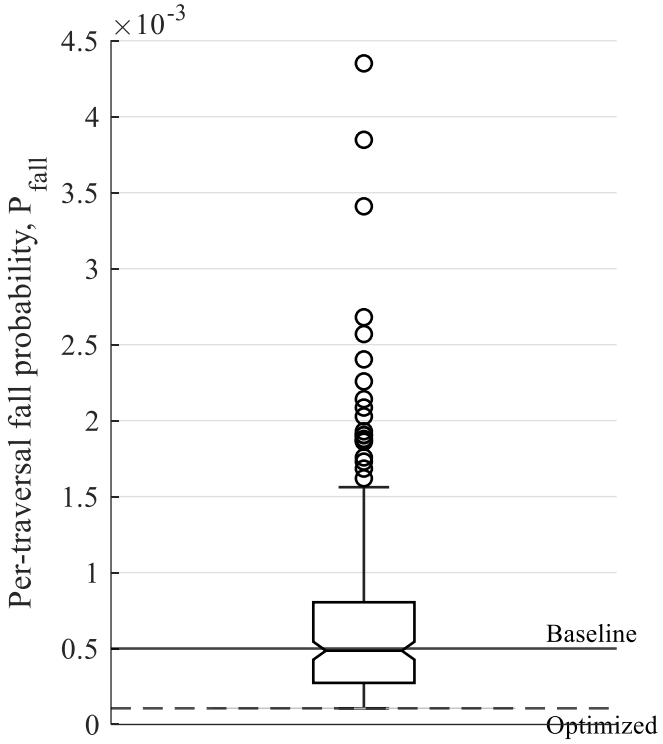


Figure 2. Dispersion in Per-Traversal Fall Probability Across Feasible Designs.

Box plots display the distribution of per-traversal fall probability across UD-compliant candidate designs (medians across parameter draws). Solid black lines denote the median baseline estimate; dashed black lines denote the median optimized configuration. Shaded bands represent interquartile ranges under parameter uncertainty.

Figure 2 shows that substantial risk dispersion exists within the compliant design space. The baseline configuration is not risk-minimal, and meaningful reductions are achievable without violating UD constraints.

5.2 Optimized Configuration

The constrained optimization identifies configurations that reduce raised thresholds, improve lighting continuity, and upgrade friction in the bathroom-adjacent segment while remaining fully UD-compliant. Importantly, the selected design minimizes total expected cost (injury cost plus amortized build cost), not fall probability alone.

Under parameter uncertainty, the optimized configuration yields a median per-traversal fall probability of approximately 0.00040—about a 20 percent reduction relative to the 0.00050 baseline. Annualized under λ_{night} = 1.5 trips per night (≈550 exposures), the median annual nighttime fall probability declines from roughly 0.24 to approximately 0.195, an absolute reduction of 4–5 percentage points.

The interquartile ranges under coefficient uncertainty remain strictly separated, indicating that the improvement is not driven by a narrow subset of parameter realizations. The result is economically meaningful because compounding under repeated exposure magnifies modest per-traversal improvements into material annual differences.

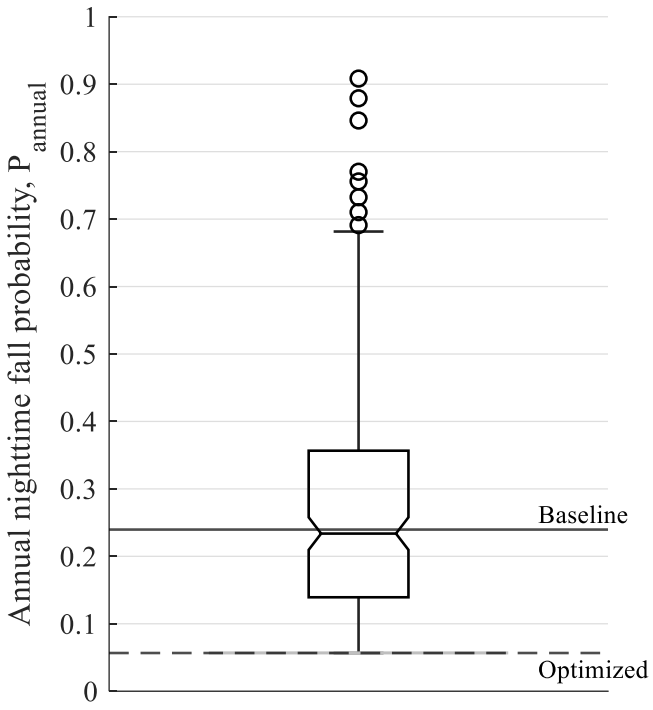


Figure 3. Annual Nighttime Fall Probability Under Repeated Exposure.

The figure annualizes per-traversal risk assuming λ_{night} = 1.5 trips per night (≈550 exposures annually). The box plot reflects dispersion across feasible designs; solid and dashed lines indicate baseline and optimized medians, respectively. Shaded bands represent interquartile ranges under parameter uncertainty.

5.3 Mechanism Decomposition

The mechanism-based structure allows attribution of risk reduction across slip, trip, and misstep channels.

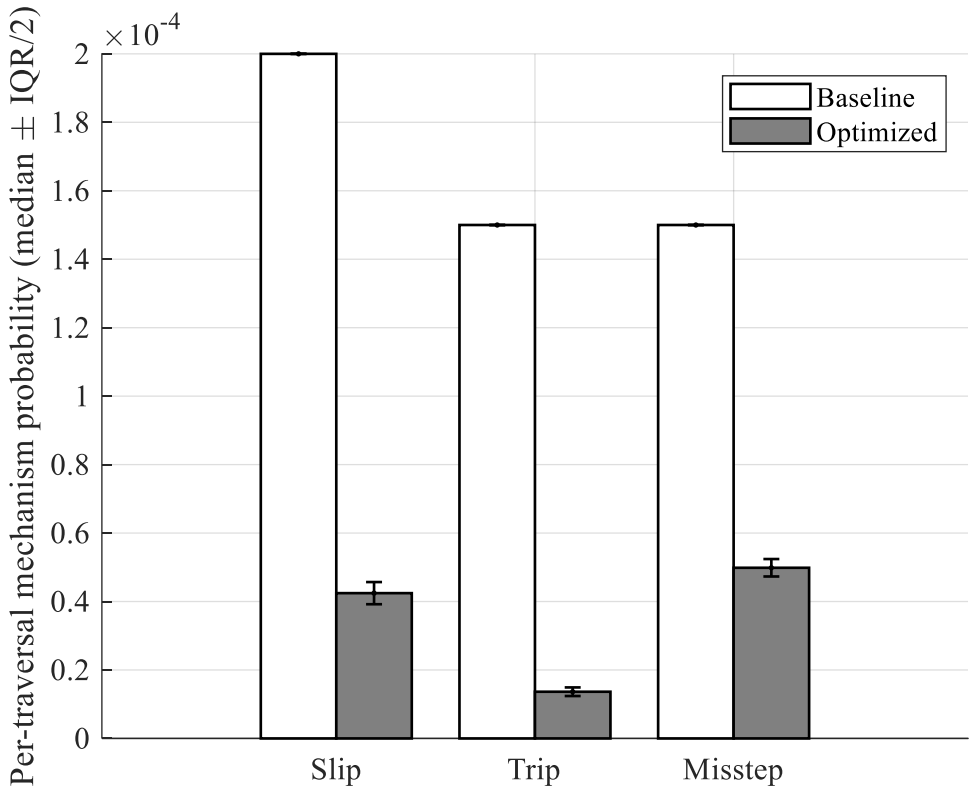


Figure 4. Mechanism Decomposition of Per-Traversal Fall Risk.

Bars report median per-traversal probabilities for slip, trip, and misstep mechanisms under baseline and optimized configurations. Error bars denote half the interquartile range across parameter draws.

Under the calibrated model, baseline per-traversal probabilities are approximately p_s = 0.00020, p_t = 0.00015, and p_m = 0.00015. In the optimized configuration, these decline to approximately 0.00016, 0.00011, and 0.00011, respectively. The decomposition confirms that lighting continuity and threshold removal primarily reduce trip and misstep risk, while friction upgrades reduce slip risk near the bathroom segment.

The additive hazard structure allows the model to identify which environmental levers drive aggregate reduction rather than treating fall probability as an opaque scalar outcome.

5.4 Economic Impact per Resident

Translating risk into economic terms clarifies the capital-allocation implications. With $E[C_{event}]$ = \$25,000, expected annual cost is proportional to annual fall probability.

Under parameter uncertainty, the baseline configuration yields a median expected annual injury cost of approximately \$6,000 per resident attributable to the nighttime path. The optimized

configuration yields approximately $4,875, implying median annual savings of about $1,125 per resident (interquartile range strictly positive).

Incorporating construction cost preserves economic viability. Incremental build cost for the selected configuration is on the order of several thousand dollars; amortized over 10 years, the annualized build burden is materially smaller than the median injury-cost reduction. Net annual savings remain positive across the interquartile range of parameter draws, indicating that the result is not sensitive to moderate coefficient variation.

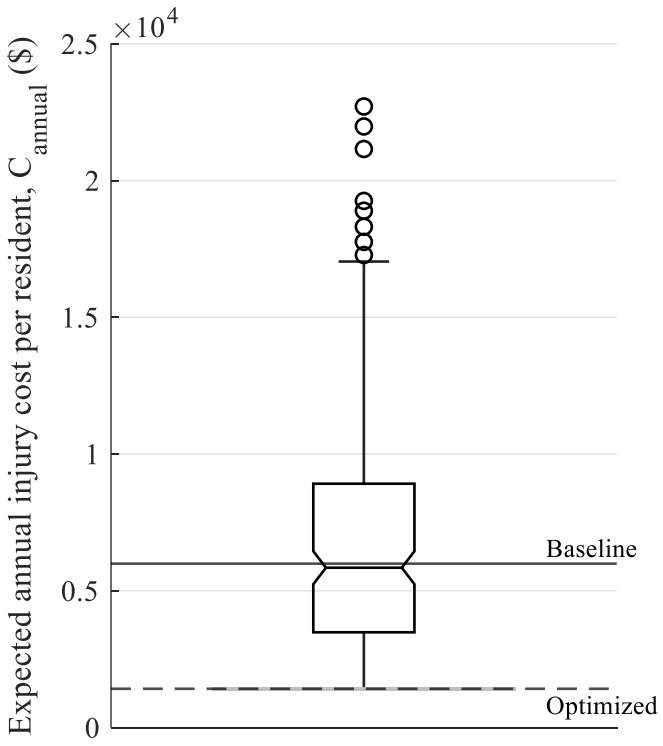


Figure 5. Expected Annual Injury Cost per Resident.

The box plot displays dispersion in expected annual cost across feasible designs (medians across parameter draws). Solid and dashed lines denote baseline and optimized medians. Shaded bands represent interquartile ranges under parameter uncertainty.

Figure 5 demonstrates that embedding injury cost into the design objective yields economically meaningful and robust savings at the per-resident level, providing the foundation for portfolio-level scaling.

5.5 Facility-Level Scaling

Scaling linearly across units, the central-case estimate of approximately $1,125 in per-resident annual savings implies roughly $0.11 million in annual savings for a 100-unit facility and approximately $1.1 million for a 1,000-unit portfolio.

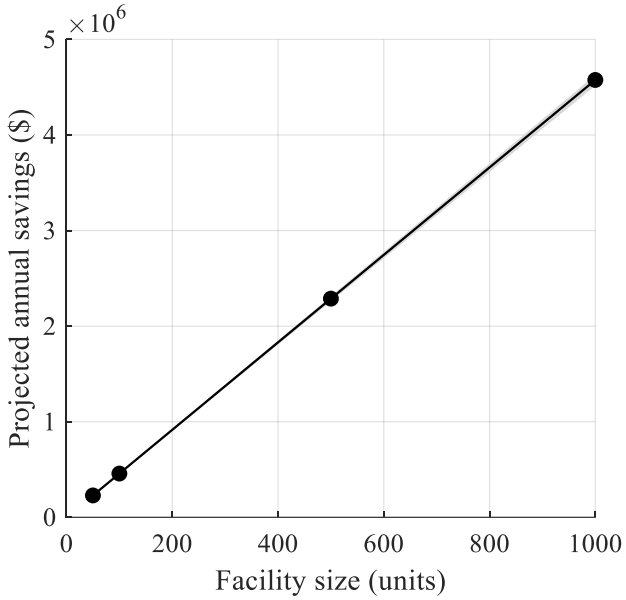


Figure 6. Scaling of Expected Annual Savings with Facility Size.

The solid line represents median projected annual savings as facility size increases, computed by multiplying per-resident median savings by the number of units. The shaded band indicates the interquartile range implied by parameter uncertainty.

These projections pertain only to the nighttime bedroom-to-bathroom path. Incorporating additional high-frequency exposure routes (e.g., kitchen access, stair transitions) would proportionally increase projected savings.

5.6 Environmental and Geometric Tradeoffs

To illustrate that geometry is not the sole lever for risk reduction, we constructed nine representative path configurations spanning: baseline; four non-geometric variants (lighting-only, surface-only, threshold-only, combined non-geometry); two slight geometry modifications combined with environmental upgrades; one significant geometry change alone; and one significant geometry change combined with environmental upgrades.

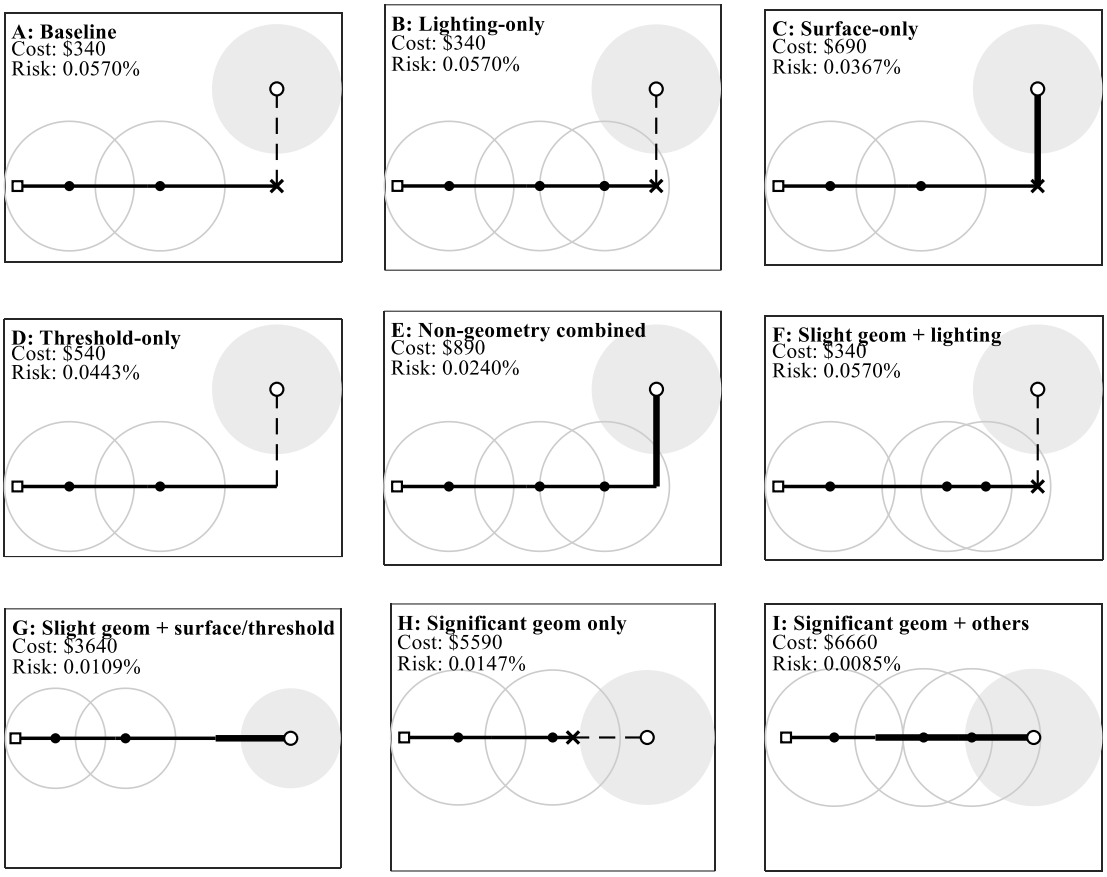


Figure 7. Nine Illustrative Path Configurations with Environmental and Geometric Tradeoffs: Each panel shows a bedroom-to-bathroom path (square → circle) with friction encoded by line style (dashed = low, solid = standard, thick = high), light sources shown as black dots with gray illumination ranges, raised thresholds marked by "×," and a gray bathroom-adjacent low-friction zone. Panels A–E vary lighting, surface, and threshold features without changing layout geometry. Panels F–G combine modest geometric adjustments with environmental improvements. Panels H–I depict substantial route redesign, with Panel-I combining geometry and environmental upgrades. Reported values indicate incremental construction cost relative to baseline and modeled per-traversal fall probability. The figure highlights tradeoffs between risk reduction and implementation cost across alternative intervention strategies.

Panels display fixed-radius illumination zones, bathroom-adjacent low-friction region, incremental build cost relative to baseline, and nominal per-traversal fall probability for each configuration. Geometry-modifying cases are assigned substantially higher incremental build costs.

The figure demonstrates three key patterns. First, measurable reductions in risk (P_{fall}) can be achieved through lighting continuity, friction upgrades, and threshold removal without altering geometry. Second, slight geometry refinements combined with environmental changes can further reduce risk at moderate incremental cost. Third, substantial geometry reconfiguration yields the lowest risk but at markedly higher capital intensity.

These configurations span the principal intervention classes within the UD-compliant space and illustrate how different combinations of environmental and geometric adjustments map into distinct cost–risk profiles. This structural variation provides intuition for the broader safety–cost frontier examined formally in Figure 8.

5.7 Safety–Cost Frontier

Figure 8 plots annualized incremental build cost, C_{build}/T, against median expected annual injury cost, C_{annual}, for the UD-compliant candidate configurations. Each point represents a feasible design evaluated under the calibrated probabilistic model with uncertainty propagated through Monte Carlo sampling.

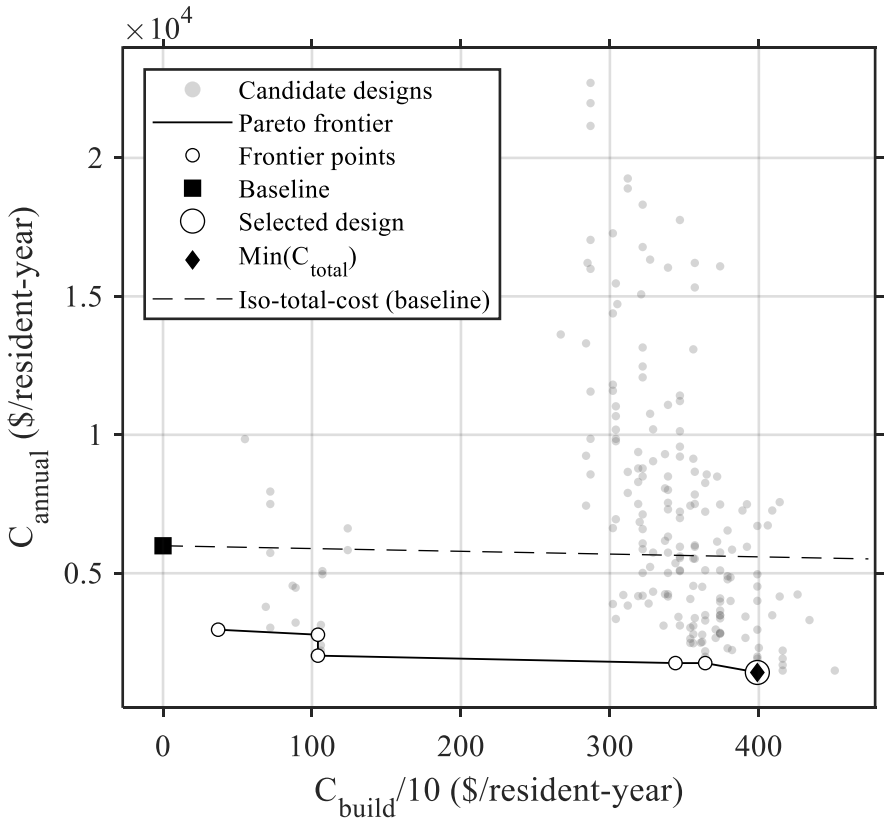


Figure 8. Safety–Cost Frontier Across UD-Compliant Configurations. The horizontal axis reports annualized incremental build cost relative to baseline. The vertical axis reports median expected annual injury cost computed under parameter uncertainty. Each point represents a feasible configuration. The solid line traces the Pareto-

efficient frontier. The square denotes the baseline configuration, the hollow circle the selected design, and the diamond the global minimizer as per Equation 5.

The dispersion of points indicates that regulatory compliance does not uniquely determine expected injury cost or associated economic burden. The dashed line represents combinations yielding total expected cost equal to that of the baseline configuration. Notably, the baseline does not lie on the lower envelope of the feasible set; alternative configurations achieve lower expected annual injury cost at comparable or only modestly higher annualized construction cost. Points below this line therefore represent configurations that yield net economic improvements relative to baseline.

The solid curve traces the Pareto-efficient frontier. Points above this frontier are strictly dominated. The configuration selected as per Equation 9 lies on the frontier. The objective-minimizing design coincides with the selected configuration under structured sampling, indicating that the optimization procedure identifies a cost-efficient solution rather than a risk-minimizing one.

The dashed iso-total-cost line through the baseline denotes combinations yielding equal total expected cost as baseline. Frontier configurations below this line represent net economic improvements relative to baseline. Together, Figures 7 and 8 demonstrate that within the UD-compliant space, risk-aware housing modifications can jointly reduce modeled fall risk and total expected economic cost.

6. Discussion

This study demonstrates that fall risk can be embedded directly into the housing design objective using a tractable probabilistic framework. Even with a simplified environmental digital twin and interpretable surrogate model, the approach identifies lower-risk configurations within UD-compliant spaces. The dispersion results (Figures 2, 3 and 5) show that the relative ordering of baseline and optimized designs is preserved under parameter uncertainty, indicating that the economic gains are not driven by a single calibration. Because exposure events occur repeatedly, modest per-traversal improvements translate into economically meaningful annual savings.

The central economic contribution is to formalize the capital-design externality. Builders optimize construction cost subject to regulatory compliance and market demand. Medicare Advantage plans internalize a substantial portion of downstream injury expenditure. When durable architectural features influence injury incidence, but the financial consequences accrue to a different actor, underinvestment in safety-enhancing design is predictable. Embedding expected injury cost into the objective function makes this misalignment explicit and quantifiable.

The heterogeneity of intervention pathways is illustrated most clearly in Figure 7. The nine path configurations demonstrate that lighting continuity, friction upgrades, and threshold removal frequently deliver favorable risk-reduction-per-dollar ratios relative to major geometric

reconfiguration. Panels B–E show that meaningful reductions can be achieved without altering layout geometry, while Panels H–I illustrate that substantial geometric redesign can further reduce risk but at materially higher capital cost. The figure makes clear that geometry is not the only lever—and often not the first-best lever—when environmental improvements remain available. This distinction matters for capital allocation: it enables prioritization of high–marginal-return interventions rather than uniform redesign.

The framework also clarifies feasible contracting mechanisms and channels for incentive realignment. Insurers could incorporate environmental risk scores into preferred-network criteria, certification programs, or shared-savings arrangements. Developers could differentiate "fall-optimized" units analogous to energy-efficiency ratings. Because the model produces Medicare-relevant expected cost metrics, it translates architectural decisions into payer-aligned financial terms without requiring continuous behavioral monitoring.

Methodologically, the framework is structured as a modular probabilistic pipeline in which environmental encoding, hazard estimation, uncertainty propagation, and economic translation operate as separable layers (Figure 1). In the present study, optimization is performed over a structured sample of 200 feasible configurations and uncertainty is propagated via calibrated Monte Carlo sampling of hazard coefficients. The modular structure clarifies which components are implemented here and which represent potential extensions.

6.1 Scalable AI Extensions

The present study implements a calibrated parametric surrogate model and evaluates a structured sample of 200 configurations. This sampling approach is sufficient to demonstrate the objective function, uncertainty propagation, and capital tradeoffs central to the paper.

The optimization problem itself is discrete and potentially high-dimensional. In larger design spaces, formal discrete optimization methods could be introduced to explore feasible modifications while preserving regulatory constraints. Such extensions would operate on the same objective function and uncertainty structure defined in Sections 3 and 4. The contribution of the current study is to define that objective transparently and demonstrate its economic implications; computational scaling can proceed without altering its interpretation or policy relevance.

Several limitations warrant consideration. First, the hazard model is a reduced-form surrogate rather than a physics-based simulation; coefficient values are calibrated to plausible targets and subjected to uncertainty propagation, but they are not yet estimated from large-scale longitudinal housing data. Second, the independence assumption across fall mechanisms simplifies aggregation and may understate interaction effects. Third, the build-cost proxy—while sufficient to illustrate capital tradeoffs in Figure 7—does not capture full contractor heterogeneity, regional price variation, or dynamic renovation constraints. Finally, the environmental digital twin

abstracts from resident heterogeneity; results reflect expected risk under representative exposure patterns rather than individualized trajectories.

Future extensions may incorporate a resident-level digital twin capturing gait variability or physiological state. Such integration would allow interaction between individual vulnerability and environmental hazard surfaces. Importantly, the economic logic remains unchanged: quantified reductions in expected injury cost can be priced, contracted, and embedded into capital investment decisions without continuous surveillance.

6.2 Application to Existing Housing Stock

Although the primary analysis focuses on new construction, the economic logic is arguably even more relevant for existing housing stock. In most prebuilt homes or established senior-living facilities, environmental risk is already embedded in durable capital. In these settings, the relevant decision is not layout from scratch but targeted retrofit.

A baseline environmental digital twin can be generated using one-time spatial capture methods (e.g., stereoscopic camera scans, LiDAR-enabled devices, or photogrammetry). Reconstructed geometry and object classification can populate the same structured feature representation used in the core model. The probabilistic framework then estimates baseline fall probability along high-frequency exposure paths and simulates counterfactual modifications such as lighting upgrades, friction enhancements, threshold removal, or limited geometric adjustment.

From a Medicare Advantage perspective, the retrofit problem can be expressed as a return-on-investment (ROI) calculation:

$$ROI \approx \frac{Expected\ Annual\ Injury - Cost\ Reduction}{Retrofit\ Cost} \quad \textbf{(6)}$$

Using the central-case estimate of approximately \$1,100–\$1,200 in expected annual savings per resident attributable to a single nighttime path, a targeted retrofit costing \$3,000–\$4,000 would imply a breakeven horizon of roughly 3–4 years. Because many MA beneficiaries remain enrolled for multiple years and facilities operate on longer horizons, such retrofits may be economically rational under shared-savings or preferred-network contracting structures. Under average Medicare Advantage tenure on the order of three to five years, the central-case breakeven horizon remains within a plausible enrollment window, although realized returns will depend on plan-specific turnover and discounting assumptions.

This framing creates a concrete pathway for insurer participation. MA plans could (i) subsidize certified retrofits in high-risk units, (ii) incorporate environmental risk scoring into care management programs, or (iii) structure shared-savings agreements with facility operators. Unlike continuous monitoring technologies, environmental retrofit is a one-time capital intervention with persistent risk effects, potentially yielding multi-year cost offsets.

Technically, the retrofit extension does not require altering the probabilistic model or the economic objective function. The same mechanism-based risk decomposition and cost translation apply. What changes is the decision margin—from design selection among new builds to capital allocation across retrofit options.

Limitations include reliance on simulated hazard coefficients, simplified representation of resident frailty, and assumptions regarding exposure frequency and cost per event. Empirical validation using claims-linked facility data would strengthen external validity and enable plan-specific ROI calibration.

7. Conclusion

Fall risk among older adults is partly embedded in durable housing capital. Because Medicare Advantage plans internalize much of the downstream injury cost while builders do not capture the financial returns to safer design, underinvestment in risk-reducing architecture is economically predictable.

Using a calibrated probabilistic framework, we demonstrate that within a Universal Design–compliant bedroom-to-bathroom path, per-traversal fall probability declines from 0.00050 to 0.00040 (20% reduction) under cost-aware optimization. Annual nighttime fall probability decreases from 0.240 to 0.195. At an assumed $25,000 cost per injurious fall, this yields median expected savings of approximately $1,125 per resident-year from a single high-frequency exposure path.

Scaled linearly, this implies roughly $0.11 million in annual savings for a 100-unit facility and $1.1 million for a 1,000-unit portfolio. After amortizing incremental build costs over 10 years, optimized configurations remain net-beneficial under central assumptions.

The contribution is economic and structural: fall risk is formalized as an objective function over a constrained design space, translated into Medicare-relevant cost, and embedded directly into capital allocation decisions. The framework is modular and extensible while remaining interpretable. Embedding quantified injury risk into housing design provides a tractable mechanism for aligning residential capital investment with value-based healthcare cost containment.

References

Albukhari, I. N. (2025). The role of artificial intelligence (AI) in architectural design: a systematic review of emerging technologies and applications. *Journal of Umm Al-Qura University for Engineering and Architecture*, *16*(4), 1457–1476. https://doi.org/10.1007/s43995-025-00186-1

Aman, J., & Matisziw, T. C. (2025). Urban sentiment mapping using language and vision models in spatial analysis. *Frontiers in Computer Science*, *7*, 1504523. https://doi.org/10.3389/fcomp.2025.1504523

Carr, K., Weir, P. L., Azar, D., & Azar, N. R. (2013). Universal Design: A Step toward Successful Aging. *Journal of Aging Research*, *2013*, 324624. https://doi.org/10.1155/2013/324624

Centers for Disease Control, & Prevention. (2024). *Facts About Falls*. https://www.cdc.gov/falls/data-research/facts-stats/index.html

González-Castro, A., Leirós-Rodríguez, R., Prada-García, C., & Benítez-Andrades, J. A. (2024). The Applications of Artificial Intelligence for Assessing Fall Risk: Systematic Review. *Journal of Medical Internet Research*, *26*, e54934. https://doi.org/10.2196/54934

Grant, L. K., St. Hilaire, M. A., Heller, J. P., Heller, R. A., Lockley, S. W., & Rahman, S. A. (2022). Impact of Upgraded Lighting on Falls in Care Home Residents. *Journal of the American Medical Directors Association*, *23*(10), 1698–1704.e2. https://doi.org/10.1016/j.jamda.2022.06.013

Haddad, Y. K., Bergen, G., & Florence, C. S. (2019). Estimating the Economic Burden Related to Older Adult Falls by State. *Journal of Public Health Management and Practice*, *25*(2), E17–E24. https://doi.org/10.1097/PHH.0000000000000816

Lu, X., Park, N.-K., & Ahrentzen, S. (2019). Lighting Effects on Older Adults' Visual and Nonvisual Performance: A Systematic Review. *Journal of Housing for the Elderly*, *33*(3), 298–324. https://doi.org/10.1080/02763893.2018.1562407

Mathon, C., Beaucamp, F., Roca, F., Chassagne, P., Thevenon, A., & Puisieux, F. (2017). Post-Fall Syndrome: Profile and Outcomes. *Annals of Physical and Rehabilitation Medicine*, *60*, e50–e51. https://doi.org/10.1016/j.rehab.2017.07.102

McDonough, C. M., Colla, C. H., Carmichael, D., Tosteson, A. N. A., Tosteson, T. D., Bell, J. E., Cantu, R. V., Lurie, J. D., & Bynum, J. P. W. (2017). Falling Down on the Job: Evaluation and Treatment of Fall Risk Among Older Adults With Upper Extremity Fragility Fractures. *Physical Therapy*, *97*(3), 280–289. https://doi.org/10.1093/ptj/pzx009

Meyer, M., Constancias, F., Vogel, T., Kaltenbach, G., & Schmitt, E. (2020). Gait Disorder among Elderly People, Psychomotor Disadaptation Syndrome: Post-Fall Syndrome, Risk Factors and Follow-Up – A Cohort Study of 70 Patients. *Gerontology*, *67*(1), 17–24. https://doi.org/10.1159/000511356

Mustaquim, M. M. (2015). A Study of Universal Design in Everyday Life of Elderly Adults. *Procedia Computer Science*, *67*(4), 57–66. https://doi.org/10.1016/j.procs.2015.09.249

Reider, L., Falvey, J. R., Okoye, S. M., Wolff, J. L., & Levy, J. F. (2024). Cost of U.S. Emergency Department and Inpatient Visits for Fall Injuries in Older Adults. *Injury*, *55*(2). https://doi.org/10.1016/j.injury.2023.111199

Shumway-Cook, A., Ciol, A. M., Hoffman, J., Dudgeon, J. B., Yorkston, K., & Chan, L. (2009). Falls in the Medicare Population: Incidence, Associated Factors, and Impact on Health Care. *Physical Therapy*, *89*(4), 324. https://doi.org/10.2522/ptj.20070107

Skubic, M., Stone, E. E., Rantz, M., & Abbott, L. (2014). Automated Health Alerts Using In-Home Sensor Data for Embedded Health Assessment. *IEEE Journal of Translational Engineering in Health and Medicine*, *2*, 1–11. https://doi.org/10.1109/JTEHM.2014.2312180

Stone, E. E., & Skubic, M. (2015). Fall Detection in Homes of Older Adults Using the Microsoft Kinect. *IEEE Journal of Biomedical and Health Informatics*, *19*(1), 290–301. https://doi.org/10.1109/JBHI.2014.2312180

Suphavarophas, P., Wongmahasiri, R., Keonil, N., & Bunyarittikit, S. (2024). A Systematic Review of Applications of Generative Design Methods for Energy Efficiency in Buildings. *Buildings 2024, Vol. 14, 14*(5). https://doi.org/10.3390/buildings14051311

Thölking, T. W., Lamers, E. C. T., & Olde Rikkert, M. G. M. (2020). A Guiding Nightlight Decreases Fear of Falling and Increases Sleep Quality of Community-Dwelling Older People: A Quantitative and Qualitative Evaluation. *Gerontology*, *66*(3), 295–303. https://doi.org/10.1159/000504883