GREEN TECHNOLOGY & INNOVATION
Volume 2 (2026) · ISSN 2979-1456

TECHNICAL ARTICLE

Instability-Constrained Lifecycle Optimisation of Low-Carbon Buildings under Structural–Envelope Degradation and Energy Feedback

Seyed Reza Samaei[1*], Assed Naked Haddad[2], James Riffat[1]

[1] World Society of Sustainable Energy Technologies, Nottingham, United Kingdom
[2] Environmental Engineering Program, Escola Politécnica, Universidade Federal do Rio de Janeiro (UFRJ), Rio de Janeiro, Brazil

RECEIVED 1 August 2026 · REVISED 25 August 2026 · ACCEPTED 21 September 2026 · PUBLISHED 13 October 2026

ABSTRACT

Lifecycle retrofit optimisation can favour interventions that perform well in conventional cost–carbon terms while leaving coupled deterioration–energy dynamics insufficiently controlled. This study introduces an instability-constrained lifecycle decision framework that jointly optimises intervention timing, deterioration reset, envelope restoration, and HVAC coefficient-of-performance improvement against expected risk-adjusted lifecycle cost, lifecycle carbon, and the 95% conditional value-at-risk (CVaR95) of cumulative instability exposure. Instability is quantified using a coupled-state eigenvalue indicator (CSEI), which represents positive local modal amplification within the reduced-order deterioration–envelope–energy model and is not a measure of structural failure, code compliance, or physical safety. The framework is demonstrated using a 12,000 m^2 office-building reference case anchored to Building Data Genome Project 2 electricity data and evaluated under uncertain degradation and climate conditions using separate training and independent 5,000-scenario out-of-sample evaluation sets. Relative to no intervention, the selected instability-aware strategy reduced expected risk-adjusted lifecycle cost by 23.41%, lifecycle carbon by 6.28%, and mean instability exposure by 99.799%, while persistent instability decreased from 96.38% to 1.60% of scenarios. Cost- and carbon-oriented strategies achieved superior performance in their respective conventional objectives but retained materially greater tail instability, demonstrating that cost and carbon alone do not recover the full lifecycle risk ordering. Across nine deterministic parameter perturbations, the fixed selected package consistently reduced operational carbon and produced no positive nominal CSEI interval through year 25. These findings establish instability-related risk as a distinct and decision-relevant dimension of lifecycle retrofit optimisation, revealing vulnerabilities that conventional cost–carbon assessment can overlook while defining a clear pathway toward monitoring-informed, asset-specific application.

KEYWORDS Lifecycle optimisation, Building retrofit, Instability constraint Risk-adjusted, lifecycle cost Lifecycle carbon, Chance-constrained decision-making Digital twin

1. Introduction

Building retrofit decisions are commonly framed as a trade-off between capital cost, operational savings and carbon reduction. This framing is necessary but incomplete when deterioration and energy demand evolve as a coupled system: envelope degradation increases thermal demand, sustained operation intensifies thermal and moisture exposure, and the resulting feedback can accelerate further deterioration. Consequently, an intervention that appears attractive in conventional cost–carbon space may still leave the asset close to a dynamically amplifying regime. In this study, instability is characterised through the coupled-state eigenvalue indicator (CSEI), which captures positive local modal amplification within the reduced-order deterioration–envelope–energy system. This formulation introduces dynamic stability as an explicit lifecycle decision dimension,

* Corresponding author. Email address: seyedreza.samaei1984@gmail.com (Seyed Reza Samaei)

© 2026 The Author(s).
Published by Caravel Press.

doi.org/10.65582/gti.2026.017

enabling retrofit strategies to be distinguished not only by their economic and carbon performance but also by the evolution of the coupled system they are intended to control.

Multi-objective building optimisation has matured from deterministic energy minimisation to lifecycle formulations that include capital cost, operating cost and emissions (Evins 2013; Hamdy *et al.* 2013). Standards provide consistent principles for life-cycle costing and environmental assessment (European Committee for Standardization 2011; International Organization for Standardization 2017). Yet a conventional Pareto front does not identify whether small state perturbations decay. Risk must be represented explicitly, particularly when uncertain degradation parameters create a long tail of early instability. Building digital-twin research has established the need for connected asset representations, probabilistic updating and decision support (Boje *et al.* 2020; Deng *et al.* 2021; Kapteyn *et al.* 2021; Rasheed *et al.* 2020). Recent smart-building studies similarly emphasise the role of AI-enabled monitoring, building automation and deployable decision-support interfaces in translating model-based energy analysis into operational practice (Onweh *et al.* 2025; Razak *et al.* 2026). Recent retrofit optimisation studies increasingly account for model and climate uncertainty rather than treating uncertainty as a post-processing exercise. Hybrid urban building-energy models have shown that uncertainty can change the number, composition and ranking of Pareto-optimal retrofit solutions. This evidence strengthens the case for robust lifecycle assessment, but existing cost–energy–carbon formulations do not explicitly test whether a coupled degradation–energy state remains locally non-amplifying (Luo *et al.* 2025).

Recent research has also moved from uncertainty-aware ranking toward explicit robustness assessment of retrofit measures across climatic, behavioural and model conditions. These studies show that a solution with strong nominal energy or cost performance may not remain preferred when uncertainty and robustness are evaluated directly (Shu *et al.* 2025; Zhang *et al.* 2025). Applied building-decarbonisation studies likewise demonstrate the value of evaluating retrofit packages against measured asset conditions and interacting energy measures rather than relying only on isolated nominal savings (Taileb and Sherzad 2026).

Recent work has treated stability preservation as an operational-control problem in building-energy models (Samaei and Riffat 2026). The present study addresses a distinct lifecycle planning problem: the timing and depth of a one-time intervention are selected over a 25-year horizon under coupled structural–envelope–energy uncertainty. The present formulation therefore concerns long-horizon intervention planning rather than short-horizon operational control.

Four research questions organise the analysis.

RQ1/H1: Do representative strategies with attractive expected cost or carbon performance exhibit materially different mean, tail and persistence-related instability exposure? It is hypothesised that cost and carbon coordinates alone do not recover the complete instability-risk ranking.

RQ2/H2: Can a risk-balanced intervention improve expected risk-adjusted lifecycle cost, lifecycle carbon and instability exposure simultaneously relative to no intervention? It is hypothesised that a compromise strategy improves all three outcomes without minimising every individual objective.

RQ3/H3: How does intervention timing affect the joint cost–carbon–instability response for a fixed intervention package? It is hypothesised that later action reduces discounted intervention expenditure but increases lifecycle carbon and persistence risk, while very early action may retain tail exposure because the HVAC benefit gradually decays.

RQ4/H4: Does the direction of benefit remain present under structured changes in deterioration, envelope, operational-feedback and climate coefficients? It is hypothesised that the numerical magnitudes and onset times change, while operational carbon and nominal instability exposure remain lower under the fixed selected package in the tested perturbation cases.

The contribution is a three-objective, risk-aware lifecycle decision formulation that propagates parameter and climate uncertainty, represents tail instability through $CVaR_{95}$, evaluates persistence through confidence-bounded chance constraints and separates training scenarios from out-of-sample strategy evaluation (Figure 1). The study additionally reports fixed-decision component ablation,

discrete timing sensitivity, search-budget diagnostics, preference-weight sensitivity and nine deterministic construction–climate parameter perturbations. The state model is included to define the decision engine; the primary contribution of this paper is the lifecycle optimisation and decision-level robustness analysis.

The framework should be interpreted as a simulation-based lifecycle decision model informed by digital-twin principles rather than as a deployed operational digital twin. It does not implement live data assimilation, a bidirectional asset connection or online state updating.

2. Decision model

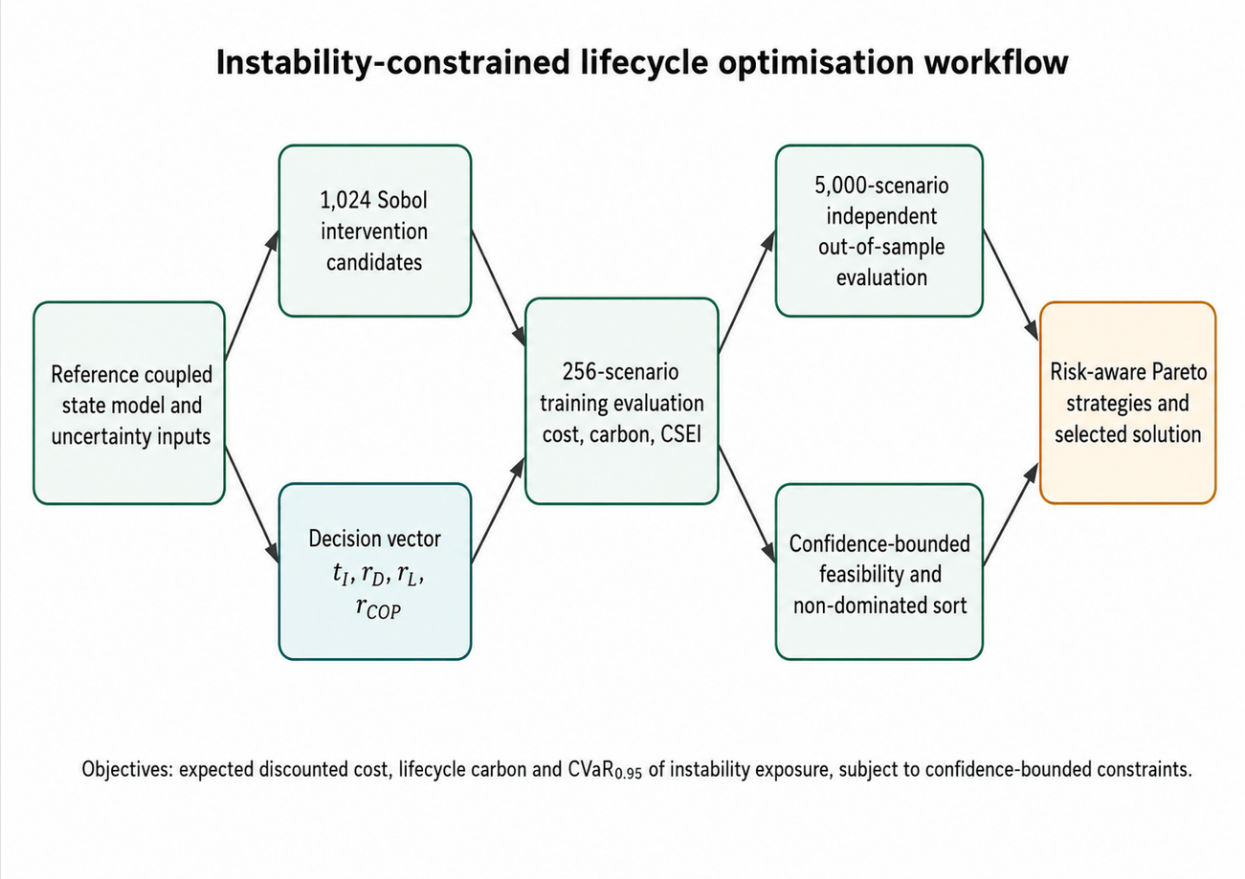

Figure 1. Instability-constrained lifecycle decision workflow. A 1,024-vector scrambled Sobol design is evaluated under 256 common-random training scenarios before feasibility filtering and non-dominated sorting. Frozen representative strategies are then assessed on an independent 5,000-scenario out-of-sample evaluation set.

2.1. Reference building and state model

The reference case is a standardised 12,000 m^2 office-building model. Its electricity anchor was derived from the Building Data Genome Project 2-meter Hog_office_Napoleon, which has a reported floor area of 12,365.4 m^2 and complete electricity records for 2016 and 2017. Annual electricity consumption was 1,140,900.84 kWh in 2016 and 1,123,697.12 kWh in 2017, corresponding to energy-use intensities of 92.27 and 90.87 kWh m^{-2} $year^{-1}$, respectively. The two-year mean energy-use intensity of 91.5699 kWh m^{-2} $year^{-1}$ was rescaled to the 12,000 m^2 reference case, giving an initial annual electricity anchor of 1,098,838.8 kWh $year^{-1}$. The empirical monthly allocation shares are reported in Table A1.

The continuous state vector is $x(t) = [D(t), L(t), H(t)]^T$, where D is the dimensionless deterioration state, L is dimensionless envelope loss and H is normalised thermal-load memory. The governing equations are:

$$dD/dt = k_D(1 - D) + q_D D^3 + c_H s(H - 1)(1 - D) \tag{1}$$

$$dL/dt = k_L(1 - L) + c_D D^2(1 - L) + c_{LH} s(H - 1)(1 - L) \tag{2}$$

$$dH/dt = \{C_m\left[1 + u_1 L + u_2 L^2\right]/COP_{rel} - H\}/\tau \tag{3}$$

where

$$s(z) = \ln\left[1 + \exp(8z)\right]/8 \tag{4}$$

and the monthly climate multiplier is

$$C_m = 1 + 0.15 \sin(2\pi m/12) + \varepsilon_m, \tag{5}$$

with m = 0 representing January. The stochastic deviation ε_m follows a stationary Gaussian AR(1) process with lag-one correlation 0.65 and stationary standard deviation 0.04. The empirical monthly electricity shares allocate the annual reference electricity across calendar months, whereas C_m acts inside the thermal-load-memory state equation and is not applied directly for a second time to the monthly electricity vector.

For calendar month q(m), monthly electricity is calculated as

$$E_m = E_{ref}\, w_q(m)\, M_{E,m}, \tag{6}$$

where E_{ref} = 1,098,838.8 kWh year^{-1}, $w_q(m)$ is the empirical monthly share and

$$M_{E,m} = (1 - 0.58) + 0.58 H_m. \tag{7}$$

The coupled-state evolution therefore modifies the HVAC-related share of whole-building electricity while retaining the empirically anchored monthly allocation.

The coupled-state eigenvalue indicator is defined as

$$CSEI(t) = \max_i \operatorname{Re}\left[\lambda_i(J(x(t)))\right] \tag{8}$$

where J is the analytical Jacobian of Equations (1–3). For the reported CSEI calculation, the Jacobian is evaluated at the scenario-specific state vector and the direct climate term is fixed at C_J = 1.0. Climate uncertainty therefore affects CSEI indirectly through the evolving states, but not through an instantaneous climate perturbation of the Jacobian.

CSEI is the dominant frozen-time modal rate. A negative value indicates asymptotic decay of the dominant frozen eigenmode, not absence of finite-time transient amplification in a non-normal time-varying system. Throughout this study, “instability” refers only to positive local modal amplification in the reduced-order deterioration–envelope–energy model. It does not denote structural buckling, collapse, code non-compliance or an experimentally validated safety threshold.

The reference parameters were prespecified to define a transparent computational test case and were not calibrated to longitudinal degradation measurements from a specific building. Numerical intervention years and package depths should therefore be interpreted as scenario-conditioned decision results rather than asset-specific prescriptions.

2.2. Intervention variables

An intervention vector $z = [t_I, r_D, r_L, r_{COP}]$ comprises intervention time, fractional reset of deterioration, fractional restoration of envelope performance and initial fractional improvement in HVAC coefficient of performance. Intervention time is restricted to the monthly grid

$t_I \in \{5, 5 + 1/12, \ldots, 22\}$ years,

and the admissible decision ranges are

$0.15 \le r_D \le 0.55$,

$0.15 \le r_L \le 0.50$,

$0.10 \le r_{COP} \le 0.35$.

The state resets are applied as

$$D(t_I^+) = (1 - r_D)D(t_I^-) \tag{9}$$

$$L(t_I^+) = (1 - r_L)L(t_I^-) \tag{10}$$

The post-intervention HVAC multiplier is

$$COP_{rel}(t) = 1 + r_{COP} \exp\left[- \lambda_{COP}(t - t_I)\right], \qquad t \ge t_I \tag{11}$$

where $\lambda_{COP} = 0.005$ year^{-1}. Before intervention, $COP_{rel} = 1$.

All state simulations use a fixed integration step $h = 1/12$ year. Intervention times are generated directly on the admissible monthly grid. At the intervention month, the pre-intervention state is recorded, the state-reset equations are applied instantaneously and fourth-order Runge–Kutta integration resumes from the post-intervention state.

The intervention window excludes immediate commissioning-stage correction and preserves at least three years of post-intervention evaluation within the 25-year lifecycle horizon. The selected intervention depths lie close to the upper limits of the adopted decision ranges; the reported depths are therefore conditional on these bounds and should not be interpreted as unconstrained physical optima.

The model treats the reference decision as a coordinated package. The component-ablation analysis subsequently deactivates individual components while holding the selected timing and remaining component depths fixed; it does not re-optimise the reduced component subsets.

2.3. Risk-adjusted cost and lifecycle-carbon accounting

All monetary quantities are scenario values in constant 2026 US dollars and are not market quotations. The cost objective is a risk-adjusted accounting measure composed of discounted electricity expenditure, discounted intervention expenditure, a nonlinear corrective/disruption reserve and a monetised instability reserve.

Monthly electricity price is

$$p_m = 0.16(1 + 0.01)^{m/12}\, USD\, kWh^{-1}, \tag{12}$$

and the real discount rate is 3.5% year^{-1}. Discounted electricity expenditure in million USD is

$$J_{elec} = 10^{-6} \sum_{m=0}^{299} \left[p_m E_m / (1 + 0.035)^{m/12} \right]. \tag{13}$$

The undiscounted intervention-package cost in million USD is

$$C_I(z) = (0.42 + 1.10 r_D) + (0.55 + 1.80 r_L) + (0.30 + 1.20 r_{COP}), \tag{14}$$

and its discounted value is $C_I(z)/ (1 + 0.035)$ ^t_I.

The complete risk-adjusted lifecycle-cost objective is

$$J_{cost} = J_{elec} + C_I(z)/(1 + 0.035)^{t_I} + 3.8(D_{25}/0.35)^4 + 45 J_{inst}, \tag{15}$$

where J_{cost} is expressed in million USD. The final two terms are present-value scenario weights rather than observed invoices. They represent a nonlinear corrective/disruption reserve and a monetised instability reserve, respectively.

The grid-emission factor is

$$EF_m = 0.42(1 - 0.012)^{m/12}\, kgCO_2e\, kWh^{-1}. \tag{16}$$

Intervention embodied carbon is

$$C_{embodied} = (45 + 220 r_D) + (70 + 420 r_L) + (30 + 180 r_{COP})\, tCO_2e. \tag{17}$$

Lifecycle carbon is calculated with explicit unit conversion as

$$J_{carbon}\left[ktCO_2e\right] = 10^{-6} \sum_{m=0}^{299} EF_m \left[kgCO_2e\, kWh^{-1}\right] E_m \left[kWh\right] + 10^{-3}\, C_{embodied}\left[tCO_2e\right]. \tag{18}$$

Operational carbon includes grid-related emissions from whole-building electricity over the 25-year horizon. Intervention carbon includes only the structural, envelope and HVAC package terms defined above. Transport, construction waste, refrigerant leakage, recurring replacement and end-of-life modules are outside the adopted boundary.

The instability exposure is

$$J_{inst} = \int_0^{25} \max\left[CSEI(t),\ 0\right] dt. \tag{19}$$

The cost objective is deliberately risk-adjusted and is not independent of J_{inst} because it includes a monetised instability reserve. The third optimisation objective retains J_{inst} in physical, non-monetary form so that the instability result can be inspected separately from the assumed reserve coefficient. All monetary rankings are conditional on the disclosed scenario coefficients and reserve functions (Table 1).

In the fixed-decision component-ablation analysis, both fixed and variable cost and embodied-carbon terms of an inactive component are set to zero. Only active components contribute their corresponding package costs and embodied carbon.

Table 1. Reference scenario assumptions used in lifecycle cost and carbon accounting.

Assumption	Reference scenario value	Basis and interpretation
Electricity price at year 0	USD 0.16 kWh^{-1}	Scenario assumption
Real electricity escalation	1.0% $year^{-1}$	Scenario assumption
Real discount rate	3.5% $year^{-1}$	Scenario assumption
Initial grid-emission factor	0.42 $kgCO_2e\ kWh^{-1}$	Scenario assumption; not a named-grid forecast
Grid-factor decline	1.2% $year^{-1}$	Scenario assumption
Structural package cost	USD [0.42 + 1.10r_D] million	Scenario cost; not a contractor quotation
Envelope package cost	USD [0.55 + 1.80r_L] million	Scenario cost; not a contractor quotation
HVAC package cost	USD [0.30 + 1.20r_{COP}] million	Scenario cost; not a supplier quotation
Structural embodied carbon	45 + 220r_D tCO_2e	Scenario coefficient; not a product-specific EPD
Envelope embodied carbon	70 + 420r_L tCO_2e	Scenario coefficient; not a product-specific EPD
HVAC embodied carbon	30 + 180r_{COP} tCO_2e	Scenario coefficient; not a product-specific EPD
Corrective/disruption reserve	USD 3.8(D_{25}/0.35)^4 million	Present-value decision-model weight
Instability reserve	USD 45J_{inst} million	Present-value decision-model weight
COP-benefit decay	0.005 $year^{-1}$	Scenario assumption

2.4. Robust objectives and constraints

The three optimisation objectives are expected risk-adjusted lifecycle cost, expected lifecycle carbon and the 95% conditional value-at-risk of cumulative instability exposure. For a sample of N instability-exposure values, $CVaR_{95}$ is evaluated using the Rockafellar–Uryasev formulation

$$CVaR_{0.95}(J_{inst}) = \min_{\eta} \{\eta + \left[1/(0.05N)\right]\sum_{i=1}^{N} \max\left[J_{inst,i} - \eta,\ 0\right]\}. \tag{20}$$

For the 5,000-scenario evaluation set, the empirical tail contains 250 scenario values.

The deterioration-screening score is

$$S_D(D_{25}) = \left[1 + \exp\{-35(D_{25} - 0.55)\}\right]^{-1}. \tag{21}$$

A deterioration-screening violation is recorded when $S_D(D_{25}) > 0.05$, which is equivalent to approximately $D_{25} > 0.46586$. S_D is a dimensionless decision-filtering surrogate. It is not a structural probability of failure, a reliability index or a code-based limit-state assessment.

A thermal-load-memory violation is recorded when

$$\max_{0 \le t \le 25} H(t) > 1.80 \tag{22}$$

This threshold represents sustained amplification of the modelled load-memory state and is not an occupant-comfort criterion.

A persistent-instability event is recorded when CSEI is positive at 13 consecutive monthly evaluation points. This definition is used identically in the Abstract, Methods, Results, tables and figure captions.

The final scenario tolerances are 5% for the deterioration-screening score, 5% for thermal-load memory and 10% for persistent instability. Training uses stricter internal limits of 3%, 3% and 5%, respectively. These margins were prespecified relative to the final tolerances to reduce the

probability that sampling error in the 256-scenario training set would admit candidates close to a final feasibility boundary.

2.5. Search, strategy selection and out-of-sample evaluation

A scrambled Sobol design generated 1,024 space-filling intervention vectors in the four-dimensional decision space using seed 20260729. Every candidate was evaluated under the same 256 common-random parameter–climate training scenarios generated with seed 20260730. Expected objectives and training chance-constraint frequencies were calculated before feasibility filtering and non-dominated sorting; no candidate was removed solely on the basis of its nominal trajectory.

Training feasibility required no more than seven deterioration-screening violations, seven thermal-load-memory violations and 12 persistent-instability violations among the 256 scenarios. These integer limits correspond to the prespecified 3%, 3% and 5% training margins.

Representative non-dominated strategies were subsequently assessed on an independent 5,000-scenario out-of-sample evaluation set generated with seed 20260728. This set was not used for candidate generation, feasibility screening, non-dominated sorting or representative-strategy selection. Final feasibility was accepted only when the one-sided 95% Clopper–Pearson upper bound for each violation probability remained below its prescribed final tolerance.

Expected-value confidence intervals were calculated using Student-t intervals. $CVaR_{95}$ confidence intervals were obtained from 2,000 percentile bootstrap resamples using seed 20260802. Each bootstrap resample treated one complete 25-year scenario trajectory as the resampling unit, and the empirical VaR threshold and $CVaR_{95}$ were recomputed in every resample.

The instability-aware representative strategy minimised weighted normalised Euclidean distance to the sampled robust utopia point using weights of 0.28, 0.32 and 0.40 for cost, carbon and instability, respectively. Candidate generation, training-set feasibility screening, objective normalisation, construction of the sampled utopia point and selection of the representative candidate IDs were completed before evaluation-set outcomes were examined.

Non-dominance refers to the finite sampled decision set. The method does not claim a continuous Pareto frontier or a globally optimal solution. Search-budget diagnostics were conducted using nested prefixes of the same scrambled Sobol sequence; these diagnostics assess stability within that sequence and do not quantify variability across independent scrambles.

3. Results

3.1. No-intervention risk

Across the independent 5,000-scenario out-of-sample evaluation set, no intervention produced expected risk-adjusted lifecycle cost of USD 8.108022 million and expected lifecycle carbon of 11.551830 $ktCO_2e$. Mean lifecycle electricity was 32.091415 GWh, mean cumulative instability exposure was 0.01758101 and mean final deterioration was 0.337072.

At least one positive monthly CSEI evaluation occurred in 4,918 of 5,000 scenarios (98.36%). Positive CSEI persisted for 13 consecutive monthly evaluation points in 4,819 of 5,000 scenarios, corresponding to 96.38%, with an exact two-sided 95% Clopper–Pearson interval of 95.82–96.88%.

The 95% confidence intervals were USD 8.053674–8.162370 million for expected risk-adjusted lifecycle cost, 11.543838–11.559822 $ktCO_2e$ for expected lifecycle carbon and 0.336028–0.338117 for mean final deterioration. These quantities define the disclosed reference decision scenario and are not market forecasts or asset-specific predictions.

3.2. Representative robust strategies

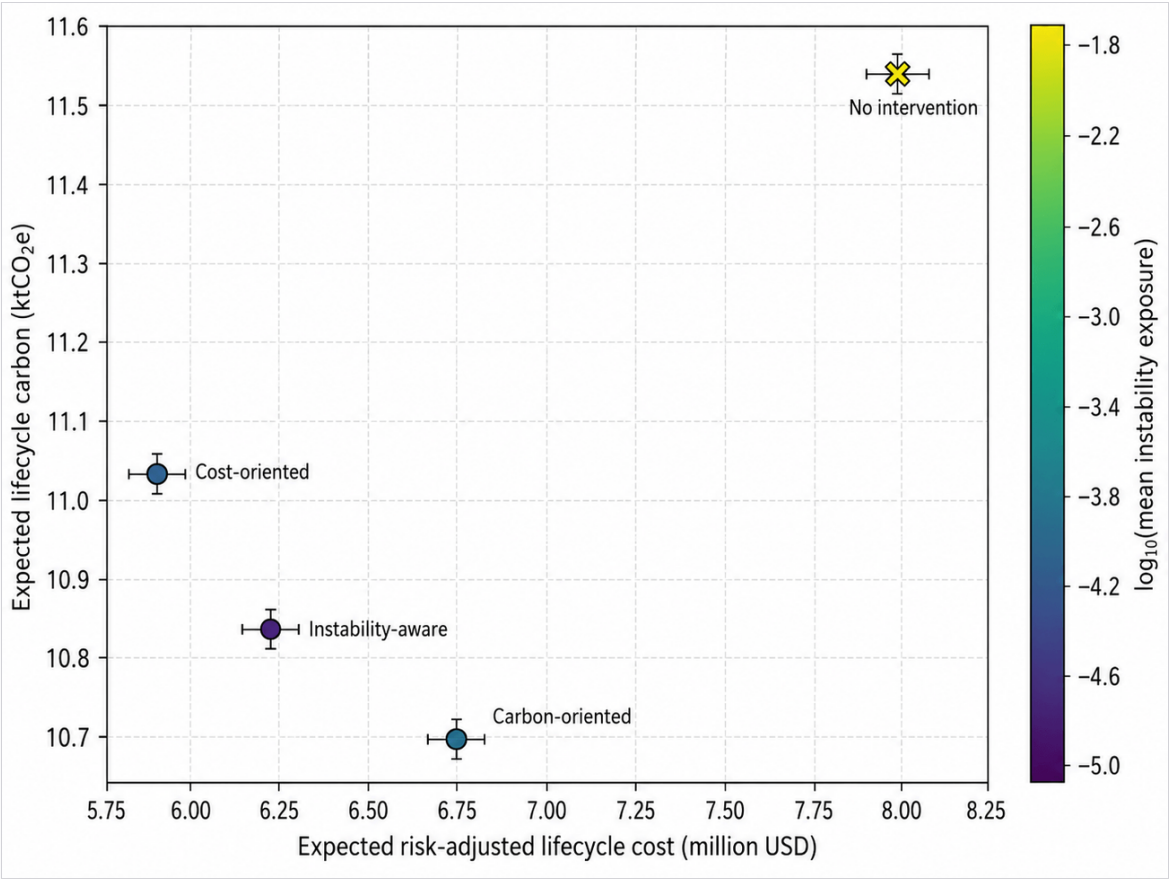

Figure 2. Representative non-dominated strategies evaluated on the independent 5,000-scenario out-of-sample set. Points show expected risk-adjusted lifecycle cost and expected lifecycle carbon; error bars show 95% confidence intervals, and colour denotes log10 of mean instability exposure. The no-intervention case is shown separately for reference.

The three intervention strategies are representative non-dominated solutions from the finite sampled decision set, not a visualisation of the complete sampled front. Figure 2 supports comparison of the selected preferences but does not establish a continuous or globally optimal Pareto frontier.

Table 2A. Decision variables of the representative robust strategies.

Strategy	Intervention year	r_D	r_L	r_{COP}
Cost-oriented	14.1667	0.538864	0.254886	0.304866
Carbon-oriented	9.7500	0.496111	0.415159	0.331019
Instability-aware	12.9167	0.543905	0.467926	0.339490

Table 2B. Out-of-sample performance of the representative robust strategies.

Metric	Cost-oriented	Carbon-oriented	Instability-aware
Expected risk-adjusted cost (USD M)	5.882358	6.706634	6.210242
95% CI	5.871903–5.892813	6.689002–6.724265	6.198803–6.221681
Change relative to no intervention	−27.45%	−17.28%	−23.41%
Expected lifecycle carbon ($ktCO_2e$)	11.026888	10.624812	10.826927

Table 2B. (continued)

Metric	Cost-oriented	Carbon-oriented	Instability-aware
95% CI	11.021256–11.032520	10.620116–10.629509	10.822248–10.831606
Change relative to no intervention	−4.54%	−8.02%	−6.28%
Mean instability exposure, J_{inst}	0.00012027	0.00017369	0.00003534
$CVaR_{95}$	0.00226089	0.00278958	0.00070674
$CVaR_{95}$ 95% CI	0.001903–0.002610	0.002454–0.003116	0.000549–0.000868
Persistent-instability events	228/5,000	226/5,000	80/5,000
Estimated persistence probability	4.56%	4.52%	1.60%
One-sided 95% upper bound	5.0753%	5.0333%	1.9238%
At least one positive monthly CSEI evaluation	9.98%	16.78%	4.90%

The cost-oriented strategy achieved the lowest expected risk-adjusted lifecycle cost but retained $CVaR_{95}$ of 0.00226089. The carbon-oriented strategy achieved the lowest lifecycle carbon but had the largest $CVaR_{95}$ among the three representative strategies. The instability-aware strategy did not minimise either conventional objective, but it produced the lowest mean and tail instability exposure and the lowest persistence probability. Relative to no intervention, it reduced expected risk-adjusted lifecycle cost by 23.41%, lifecycle carbon by 6.28% and mean instability exposure by 99.799%. These comparisons reveal risk differences among representative non-dominated strategies; they do not isolate the causal effect of adding or removing individual optimisation objectives or constraints.

Table 3. Out-of-sample chance-constraint evidence for the representative strategies. A deterioration-screening violation denotes $S_D(D_{25}) > 0.05$, a thermal-load-memory violation denotes max H(t) > 1.80, and a persistent-instability event denotes positive CSEI at 13 consecutive monthly evaluation points.

Strategy	Constraint	Violations	Estimate	One-sided 95% upper bound	Tolerance	Decision
Cost-oriented	Deterioration-screening score	0/5,000	0%	0.0599%	5%	Pass
Cost-oriented	Thermal-load memory	0/5,000	0%	0.0599%	5%	Pass
Cost-oriented	Persistent instability	228/5,000	4.56%	5.0753%	10%	Pass
Carbon-oriented	Deterioration-screening score	0/5,000	0%	0.0599%	5%	Pass
Carbon-oriented	Thermal-load memory	0/5,000	0%	0.0599%	5%	Pass
Carbon-oriented	Persistent instability	226/5,000	4.52%	5.0333%	10%	Pass
Instability-aware	Deterioration-screening score	0/5,000	0%	0.0599%	5%	Pass
Instability-aware	Thermal-load memory	0/5,000	0%	0.0599%	5%	Pass
Instability-aware	Persistent instability	80/5,000	1.60%	1.9238%	10%	Pass

3.3. State trajectories and timing

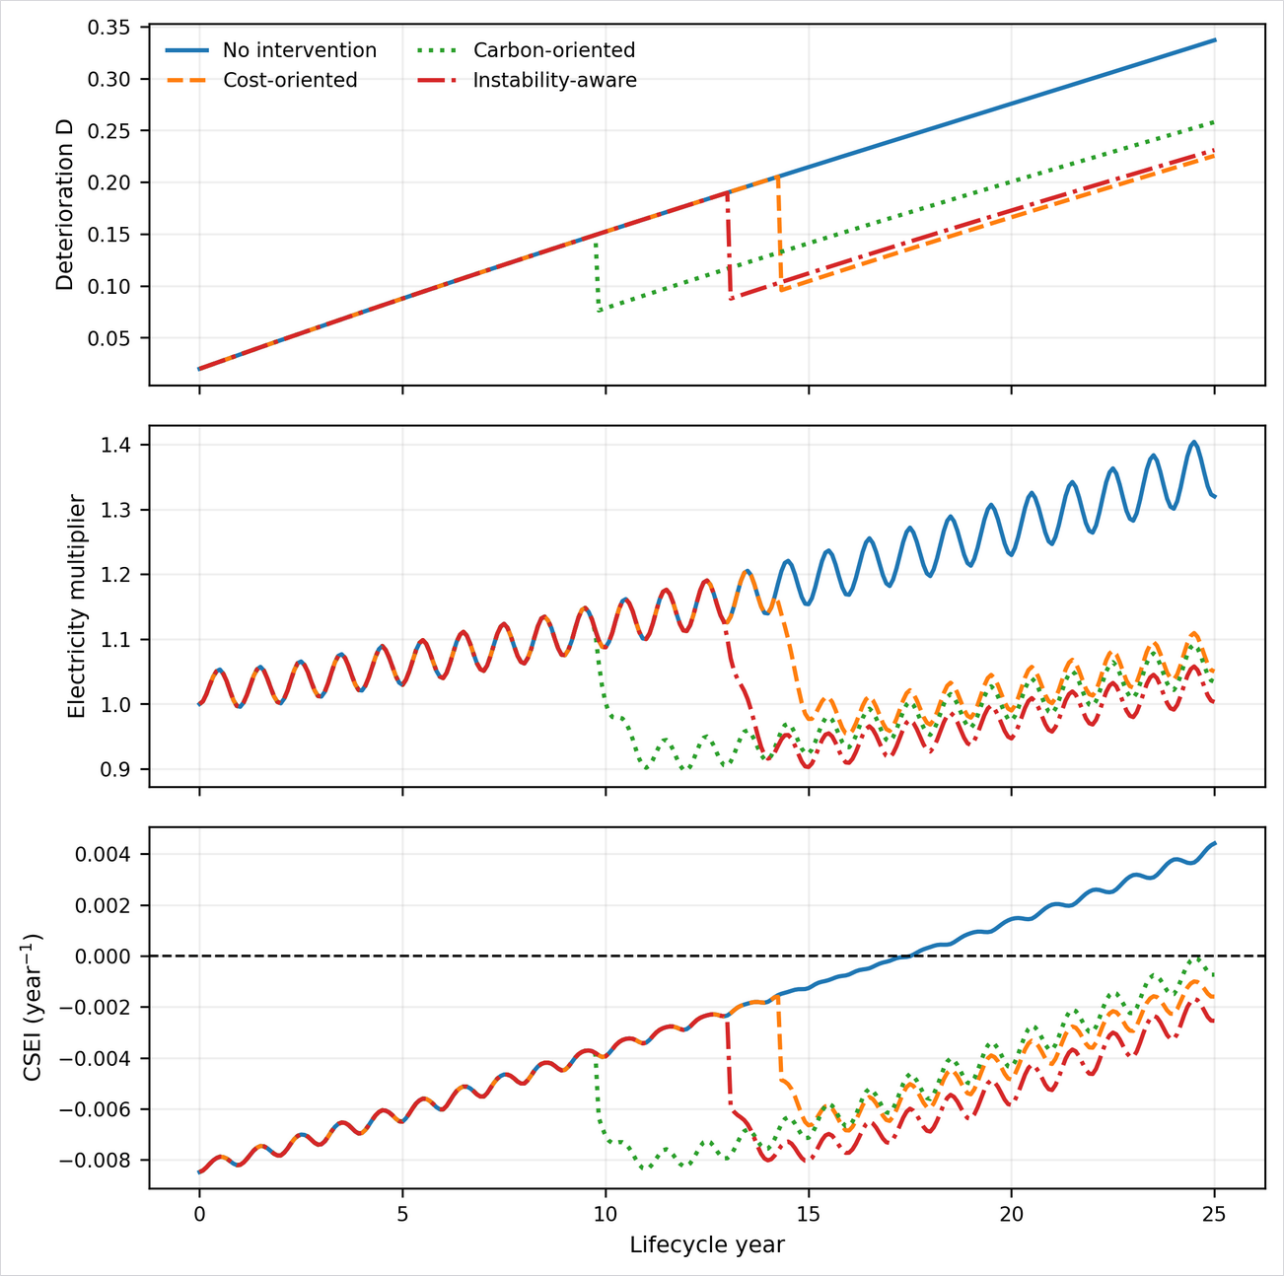

Figure 3. Nominal deterioration, electricity-multiplier and CSEI trajectories for no intervention and the three representative strategies using the final monthly-grid intervention decisions. CSEI denotes the dominant frozen-time modal rate and should not be interpreted as a structural safety margin.

Figure 3 illustrates why expected cost and carbon alone do not describe the complete dynamic response. All three intervention strategies reduce the post-intervention electricity multiplier, but their timing and state resets produce different CSEI trajectories. The instability-aware strategy yields no positive nominal CSEI interval through year 25, whereas the cost-oriented strategy permits greater pre-intervention exposure because later expenditure is favoured by discounting. These nominal trajectories support mechanism interpretation but do not replace the out-of-sample probabilistic results in Tables 2 and 3.

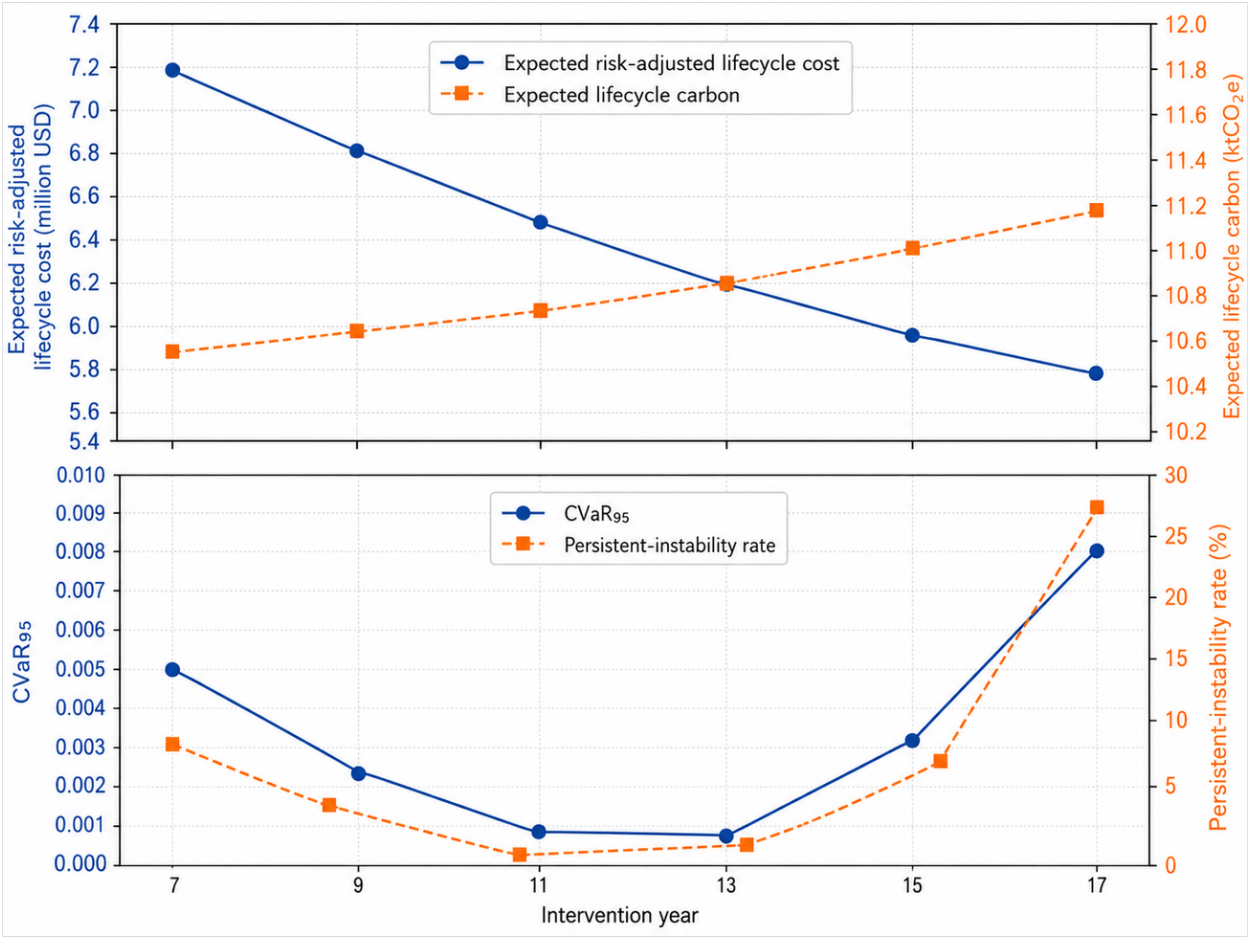

Figure 4. Out-of-sample timing sensitivity of the selected intervention package. Package depths are fixed at r_D = 0.543905, r_L = 0.467926 and r_{COP} = 0.339490; only intervention timing changes. Each tested timing is evaluated using the same 5,000-scenario out-of-sample set.

Timing produced a non-monotonic risk response among the six tested intervention years (Figure 4). Years 11 and 13 produced the lowest tail-instability measures in this discrete sweep. Later action reduced discounted risk-adjusted lifecycle cost but increased lifecycle carbon, $CVaR_{95}$ and persistence. The tested points support the selected year-12.92 decision, but they do not establish a continuous optimal interval.

3.4. Fixed-decision intervention-component ablation

Table 4. Fixed-decision component ablation at the selected instability-aware intervention time and component depths. Reduced configurations are not re-optimised.

Configuration	Expected cost (USD M)	Expected carbon ($ktCO_2e$)	Mean J_{inst}	Persistence
No intervention	8.108022	11.551830	0.01758101	96.38%
Structural only	5.506176	11.689017	0.00172986	43.38%
Envelope only	8.474309	11.373032	0.01372157	95.10%
HVAC only	7.416906	10.744590	0.00694651	76.80%
Structural + envelope	6.137398	11.513171	0.00090570	29.34%
Structural + HVAC	5.456821	10.888966	0.00012619	3.98%
Envelope + HVAC	7.975926	10.680431	0.00379830	51.60%
Full package	6.210242	10.826927	0.00003534	1.60%

This analysis isolates component mechanisms at the timing and depths of the selected full-package decision (Table 4). It is not a re-optimisation of the seven reduced intervention classes and therefore does not identify the optimum within each component subset. Structural repair produced the largest reduction in the deterioration-related reserve, HVAC improvement produced the largest direct operational-carbon benefit and the full package yielded the lowest mean instability exposure and persistence among the tested fixed-depth configurations. Structural plus HVAC had the lowest risk-adjusted cost among the low-persistence fixed-depth combinations. The selected full package should therefore be interpreted as the risk-balanced representative decision, not as the minimum-cost component combination.

3.5. Structured construction–climate parameter stress tests

The framework was examined using nine deterministic parameter-perturbation cases formed from three construction-related parameter sets and three climate-related coefficient conditions. The RC-educational-inspired set applies $k_D \times 0.85$, $k_L \times 1.10$, $c_H \times 1.15$ and an HVAC share of 0.62. The hybrid-residential-inspired set applies $k_D \times 0.75$, $k_L \times 1.20$, $c_H \times 0.80$ and an HVAC share of 0.45. The hot-humid condition applies a climate-severity multiplier of 1.18, $c_{LH} \times 1.25$ and $u_2 \times 1.10$. The cold condition applies a climate-severity multiplier of 1.12, $u_1 \times 1.20$ and $k_L \times 1.10$.

The nine cases are deterministic parameter-perturbation probes. The selected reference package is applied unchanged and is not re-optimised for each case. The hot-humid and cold labels are shorthand for structured coefficient changes rather than location-specific weather files. The results therefore assess directional transfer of a fixed decision, not archetype-specific optimality or statistical generalisation to building stocks.

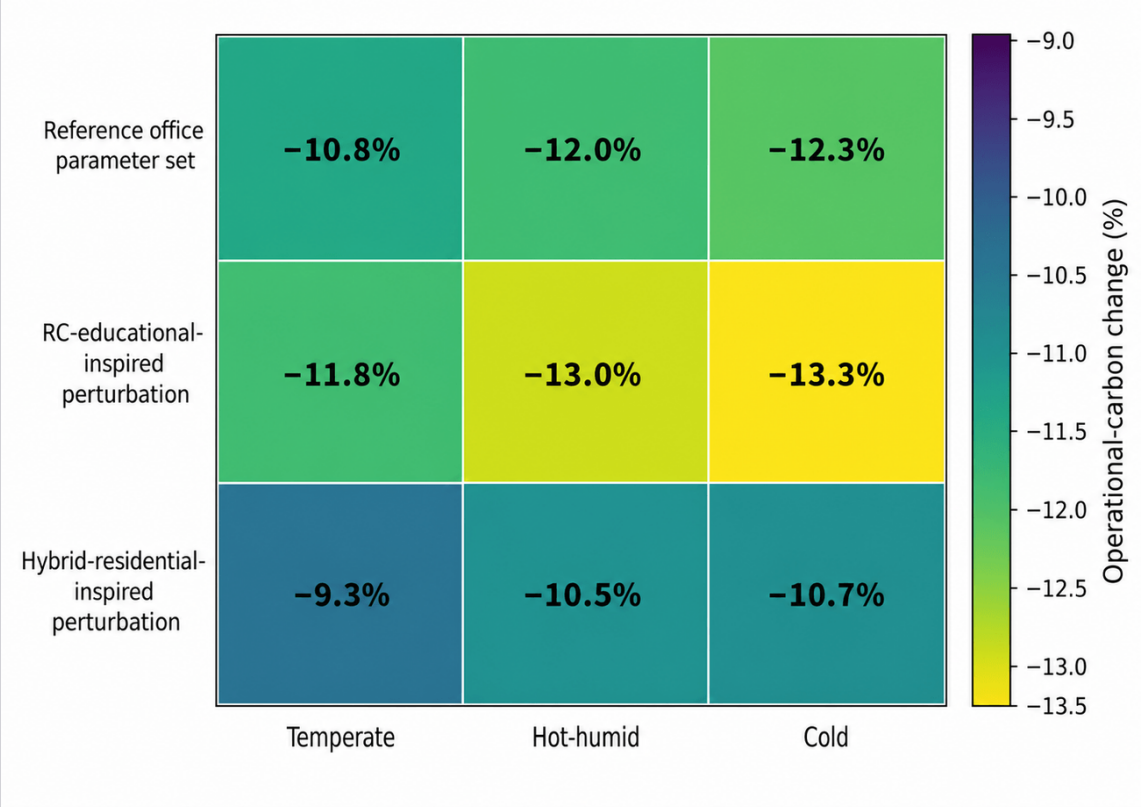

Figure 5. Operational-carbon change, excluding intervention embodied carbon, under nine deterministic construction–climate parameter perturbations. The fixed selected package is applied without re-optimisation, and the construction labels do not represent calibrated building-stock archetypes.

Table 5. Structured construction–climate stress-test results for the selected instability-aware strategy.

Parameter set	Climate	Baseline onset	Onset under fixed package	Final D: baseline→package	Operational-carbon change	Nominal CSEI result
Reference office parameter set	Temperate	17.58	>25	0.337→0.239	−10.8%	No positive interval
Reference office parameter set	Hot-humid	15.92	>25	0.354→0.246	−12.0%	No positive interval
Reference office parameter set	Cold	17.00	>25	0.353→0.245	−12.3%	No positive interval
RC-educational-inspired perturbation	Temperate	16.92	>25	0.304→0.210	−11.8%	No positive interval
RC-educational-inspired perturbation	Hot-humid	15.75	>25	0.325→0.220	−13.0%	No positive interval
RC-educational-inspired perturbation	Cold	16.92	>25	0.324→0.219	−13.3%	No positive interval
Hybrid-residential-inspired perturbation	Temperate	20.83	>25	0.270→0.188	−9.3%	No positive interval
Hybrid-residential-inspired perturbation	Hot-humid	19.83	>25	0.286→0.195	−10.5%	No positive interval
Hybrid-residential-inspired perturbation	Cold	21.75	>25	0.285→0.195	−10.7%	No positive interval

Application of the fixed selected package reduced operational carbon by 9.3–13.3% across the nine parameter perturbations, excluding intervention embodied carbon from this comparison (Figure 5 and Table 5). No nominal trajectory under the fixed package contained a positive CSEI interval through year 25. These results are directional stress-test outcomes within the reference model and should not be interpreted as calibrated predictions for educational, residential, hot-humid or cold-climate building stocks.

3.6. Search-budget, preference-weight and fixed-strategy sensitivity

Nested prefixes of the same scrambled Sobol sequence showed that the selected representative decision changed between 256 and 512 candidates but remained candidate 492 at budgets of 512, 1,024 and 2,048 candidates. Normalised hypervolume increased and inverted generational distance decreased as sampled-front coverage improved. All five tested preference-weight vectors selected candidate 492. These diagnostics establish stability of the representative decision within one nested scrambled sequence; they do not demonstrate independence from the Sobol scramble or prove global optimality.

Table 6. Fixed-strategy one-at-a-time post-optimality sensitivity of the three representative strategies. Decision vectors were held fixed; the analysis was not a re-optimisation.

Assumption	Tested levels	Strategy with minimum risk-adjusted cost	Strategy with minimum lifecycle carbon	Strategy with lowest $CVaR_{95}$	Any rank reversal?
Corrective-reserve coefficient	USD 1.90, 2.85, 3.80, 4.75 and 5.70 million	Cost-oriented	Carbon-oriented	Instability-aware	No
Instability-reserve coefficient	USD 22.50, 33.75, 45.00, 56.25 and 67.50 million	Cost-oriented	Carbon-oriented	Instability-aware	No
Real discount rate	2.0%, 3.5% and 5.0%	Cost-oriented	Carbon-oriented	Instability-aware	No
Electricity price at year 0	USD 0.12, 0.16 and 0.20 kWh^{-1}	Cost-oriented	Carbon-oriented	Instability-aware	No
Intervention-cost multiplier	0.75, 1.00 and 1.25	Cost-oriented	Carbon-oriented	Instability-aware	No
Embodied-carbon multiplier	0.75, 1.00 and 1.25	Cost-oriented	Carbon-oriented	Instability-aware	No

No ranking reversal occurred among the three fixed representative strategies within the tested one-at-a-time ranges: the cost-oriented strategy retained the lowest risk-adjusted cost, the carbon-oriented strategy retained the lowest lifecycle carbon and the instability-aware strategy retained the lowest $CVaR_{95}$ (Table 6). Because the decision vectors were held fixed, this table tests post-optimality ranking stability only. It does not represent re-optimisation and does not examine alternative functional forms for the nonlinear corrective reserve. Monetary rankings remain conditional on the adopted fourth-power function and the D_{25} reference value of 0.35.

4. Discussion

4.1. Evaluation of the research hypotheses

The simulation results are consistent with H1 within the adopted model and scenario assumptions. The cost-oriented strategy achieved expected risk-adjusted lifecycle cost of USD 5.882 million but retained $CVaR_{95}$ of 0.00226089. The carbon-oriented strategy achieved the lowest lifecycle carbon, 10.625 $ktCO_2e$, but had $CVaR_{95}$ of 0.00278958. The instability-aware strategy had higher cost and carbon than the respective single-objective representatives but reduced $CVaR_{95}$ to 0.00070674 and persistence to 1.60%. Cost and carbon coordinates alone therefore did not recover the complete risk ordering of the three representative strategies.

The results are consistent with H2 under the stated risk-adjusted cost accounting. Relative to no intervention, the instability-aware strategy reduced expected risk-adjusted lifecycle cost by 23.41%, lifecycle carbon by 6.28% and mean instability exposure by 99.799%. It therefore improved all three reported objectives relative to the reference case without minimising either conventional objective.

The discrete timing results are consistent with H3. Later action reduced discounted risk-adjusted cost but increased lifecycle carbon, $CVaR_{95}$ and persistent-instability probability. Years 11 and 13 produced the lowest tail-instability measures among the six tested timings. The analysis does not establish a continuous optimal timing interval.

The results are consistent with H4 only in directional terms within the nine deterministic parameter-perturbation cases. The fixed selected package reduced operational carbon and retained

negative nominal CSEI through year 25 in every tested case. These outcomes do not constitute statistical generalisation to different building stocks or climates.

4.2. Decision-level robustness checks within the reference model

The analyses in this study evaluate decision robustness within a prespecified reference model; they do not constitute external validation of the state equations against longitudinal observations from a real building. The independent evaluation scenarios protect the reported strategy comparison from direct reuse of the training scenarios, but they remain simulations generated from the same model family and uncertainty assumptions. Model calibration and prospective post-intervention validation remain necessary before asset-specific application.

The decision framework was evaluated independently of candidate selection. All 1,024 intervention vectors were assessed under the same common-random training scenarios before feasibility filtering and non-dominated sorting, thereby avoiding nominal pre-screening bias. Final performance was estimated using a non-overlapping 5,000-scenario out-of-sample evaluation set. Student-t confidence intervals for expected values, bootstrap confidence intervals for $CVaR_{95}$ and one-sided exact Clopper–Pearson upper bounds distinguished point estimates from statistical evidence of feasibility. Within the single scrambled Sobol sequence, nested candidate budgets up to 2,048 vectors showed that the selected representative decision remained unchanged from 512 candidates onward, while five alternative preference-weight vectors selected the same compromise candidate. Fixed-decision component ablation, discrete timing sensitivity, structured construction–climate parameter perturbations and fixed-strategy post-optimality economic sensitivity provided additional evidence of decision stability within the adopted model and scenario assumptions. These checks strengthen confidence in the internal robustness of the reported decision but do not establish global optimality, external model validity or asset-level predictive accuracy.

4.3. Asset-management implications

The results indicate that intervention scope should be selected according to the mechanism that drives lifecycle risk. In the fixed-decision component analysis, structural repair reduced the deterioration state and the associated corrective/disruption reserve but increased lifecycle carbon when implemented alone because its embodied-carbon burden was not offset by a substantial reduction in operational energy demand. HVAC improvement produced the largest direct operational-carbon benefit but retained considerable persistent-instability exposure when structural deterioration and envelope loss continued to develop. Envelope restoration weakened the deterioration-to-load pathway but was insufficient by itself to control the coupled state response.

Among the tested fixed-depth configurations, the coordinated full package produced the lowest mean instability exposure and persistence probability. The structural-plus-HVAC configuration produced a lower risk-adjusted lifecycle cost but retained greater residual instability than the full package. Because these reduced component configurations were not independently re-optimised, the comparison should be interpreted as a mechanism-focused fixed-decision ablation rather than evidence that one component combination is universally preferable.

Asset owners should therefore report at least four decision outputs for each candidate intervention: expected risk-adjusted lifecycle cost, expected lifecycle carbon, tail instability exposure and confidence-bounded persistence probability. Direct expenditure and model-based reserve terms should also be distinguished wherever possible. A favourable cost estimate should not be allowed to obscure an unfavourable physical-risk trajectory or a materially larger tail exposure. A practical deployment pathway would couple the present decision model to an asset's monitoring and information infrastructure. Building-management-system and IoT streams could provide electricity, indoor-environment and equipment-status data; inspection and maintenance records could inform deterioration and envelope states; and BIM or asset-information models could supply system metadata and intervention histories. These data could support periodic state estimation and

parameter updating, after which the lifecycle optimiser could be rerun at scheduled decision points rather than used as a real-time safety controller. Recent smart-building studies demonstrate the value of integrating monitoring, building automation and deployable predictive decision-support tools in operational energy management (Onweh *et al.* 2025; Razak *et al.* 2026). Such an architecture would move the present simulation-based framework toward a monitoring-informed digital-twin decision service while preserving the distinction between model-based instability screening and formal structural or regulatory assessment.

4.4. Carbon and policy implications

Conventional retrofit assessments often extrapolate commissioning-year energy performance across the remaining service life. The reference-model results indicate that such an assumption can be misleading when structural deterioration, envelope loss and energy demand interact dynamically. In the no-intervention case, degradation-related amplification increased operational demand and progressively moved the coupled system towards a positive local modal regime.

The timing analysis further showed that the earliest tested intervention was not necessarily the least risky and that the lowest-cost representative strategy was not the lowest-risk strategy. Later intervention reduced discounted risk-adjusted lifecycle cost but increased lifecycle carbon, $CVaR_{95}$ and persistent-instability probability. Conversely, earlier intervention reduced operational carbon but did not produce the lowest tail-risk values among the tested timings.

Public retrofit programmes, green-finance assessments and asset-decarbonisation plans may therefore benefit from degradation-adjusted energy projections, explicit embodied-carbon boundaries and a dynamic-risk screen in addition to conventional payback and energy-saving metrics. Such assessments should distinguish operational carbon from intervention embodied carbon and should avoid treating the numerical intervention timing identified in this reference case as a universal policy recommendation.

4.5. Generalisation assumptions and limitations

The construction–climate cases are deterministic coefficient perturbations of the reference office model, not calibrated statistical representations of reinforced-concrete educational buildings, hybrid residential buildings or location-specific climate conditions. Their purpose is to examine whether the direction of benefit persists under structured changes in deterioration, envelope loss, load feedback, HVAC share and climate-related coefficients. Building-stock generalisation would require archetype-specific load profiles, floor areas, occupancy schedules, HVAC systems, deterioration observations, weather files, intervention-performance data and local cost and carbon inventories. The present formulation is primarily an asset-level lifecycle decision model. Extension to a building portfolio would add a second decision layer in which asset-specific intervention vectors compete for shared annual capital, labour and carbon budgets. Candidate–scenario trajectory evaluations are separable, so the scenario engine is amenable to parallel execution; however, direct re-optimisation for every asset would substantially increase computational burden. Portfolio deployment would therefore benefit from parallel computing, archetype-specific parameter libraries, preliminary asset screening and, where required, surrogate or hierarchical optimisation. Transfer to other building typologies would likewise require redefinition or calibration of load profiles, HVAC shares, deterioration parameters, decision bounds and local cost and carbon coefficients rather than direct reuse of the numerical solution reported here.

The 25-year analysis represents one coordinated intervention and does not model subsequent HVAC replacement, refrigerant leakage, repeated envelope maintenance or multiple structural-repair cycles. These exclusions may affect both cost and carbon rankings, particularly for interventions implemented early in the lifecycle.

The uncertainty model represents selected deterioration, envelope, load-feedback and climate parameters but excludes occupant behaviour, rebound effects, implementation delay, cost overrun

and uncertainty in the realised repair and efficiency improvements. The reported confidence intervals are therefore conditional on successful implementation of the specified intervention depths and on the adopted probability distributions.

The state-model parameters were prespecified to define a transparent computational reference case and were not calibrated to asset-specific longitudinal degradation measurements. The selected intervention year and component depths are consequently scenario-conditioned outputs rather than universal retrofit recommendations. For asset-specific deployment, a practical calibration strategy would combine longitudinal energy and indoor-environmental measurements with inspection and maintenance records. The current parameter distributions could be treated as prior distributions and updated as observations accumulate, for example through Bayesian updating or other probabilistic state-estimation procedures. The resulting posterior parameter and state uncertainty could then be propagated through the same lifecycle and chance-constraint workflow. Monitoring-informed state estimation and prospective post-intervention validation would remain necessary to identify model discrepancy and verify that the assumed repair and efficiency benefits persist in service.

The cost and carbon coefficients are disclosed scenario assumptions rather than contractor quotations, supplier prices or product-specific environmental product declarations. Accordingly, the term “cost reduction” refers throughout the paper to a reduction in the adopted risk-adjusted lifecycle-cost metric. Investment-grade application would require asset-specific intervention prices, disruption estimates, energy tariffs, grid-emission trajectories and product-level environmental declarations.

The selected values of r_D, r_L and r_{COP} lie close to the upper bounds of the admissible decision space. The reported intervention depths are therefore conditional on the selected bounds, and no claim of an unconstrained physical optimum is made.

The search-budget audit uses nested prefixes of a single scrambled Sobol sequence and does not quantify variability across independent scrambles. Similarly, the fixed-strategy economic sensitivity varies selected monetary coefficients but does not examine alternative exponents or reference values for the nonlinear corrective-reserve function. Monetary rankings therefore remain conditional on the adopted fourth-power formulation and the deterioration reference value of 0.35.

The fixed-decision component ablation isolates component mechanisms at the selected intervention timing and depths but does not re-optimise the reduced component classes. The construction–climate cases likewise apply the selected package without re-optimisation and should not be interpreted as archetype-specific optimal solutions.

Finally, CSEI is a local frozen-time modal indicator. A negative CSEI does not rule out finite-time transient amplification in a non-normal, time-varying system. The deterioration-screening score and thermal-load-memory threshold are also decision-screening surrogates rather than code-based safety or comfort criteria. CSEI and the associated constraints must therefore not be used as substitutes for structural safety, serviceability, occupant-comfort or regulatory compliance assessments.

5. Conclusions

Lifecycle retrofit planning changes when dynamic fragility is treated as an explicit decision dimension rather than an unobserved consequence of cost–carbon optimisation. The framework developed here combines risk-adjusted lifecycle cost, lifecycle carbon and tail instability exposure with confidence-bounded persistence constraints while separating candidate screening from independent out-of-sample evaluation. In the reference case, the selected instability-aware strategy improved the principal decision metrics relative to no intervention: expected risk-adjusted lifecycle cost decreased by 23.41%, lifecycle carbon by 6.28% and mean instability exposure by 99.799%, while persistent instability fell from 96.38% to 1.60% of scenarios.

The decision-level implication is not that a particular intervention year or package should be transferred directly to other buildings. Rather, the results show that cost-oriented and carbon-oriented strategies can retain materially different tail and persistence-related dynamic risks even

when they appear attractive under conventional objectives. Introducing instability metrics therefore changes what constitutes an acceptable lifecycle compromise and provides a means of screening solutions whose dynamic risk would remain hidden in cost–carbon space. The timing, component-ablation and parameter-perturbation analyses further indicate that the mechanism and magnitude of benefit depend on intervention timing, coupled deterioration pathways and the adopted model assumptions.

These conclusions remain conditional on the prespecified reduced-order state model, uncertainty distributions, decision bounds and scenario-based cost and carbon coefficients. The framework is not a structural-safety model, a code-compliance tool or a deployed operational digital twin, and asset-specific use requires monitored-data calibration, inspection-informed updating, local economic and carbon inputs and prospective validation. Within these limits, the principal scientific contribution is a lifecycle decision framework in which dynamic instability risk is made visible and testable alongside cost and carbon, enabling more defensible long-horizon retrofit screening.

Appendix A. Prespecified model and search settings

The reference dynamics were k_D=0.01382 year^{-1}, q_D=0.03911 year^{-1}, c_H=0.00400 year^{-1}, k_L=0.00986 year^{-1}, c_D=0.02046 year^{-1}, c_{LH}=0.00139 year^{-1}, u_1=1.3919, u_2=4.3504 and τ=0.4585 year. Initial states were $D(0)$=0.02, $L(0)$=0.02 and $H(0)$=1.00. Positive parameters followed independent lognormal distributions.

Let $j \in \{0, \ldots, 299\}$ denote lifecycle month and $q(j)=j \bmod 12$ denote calendar month. Climate forcing was

$$C_j = 1 + 0.15 \sin\left(\frac{2\pi q(j)}{12}\right) + \varepsilon_j$$

with the stationary process

$$\varepsilon_0 \sim \mathcal{N}\left(0,\ 0.04^2\right), \qquad \varepsilon_j = 0.65\varepsilon_{j-1} + 0.04\sqrt{1-0.65^2}\, z_j, \qquad z_j \sim \mathcal{N}(0,\ 1)$$

The disturbance evolved continuously across the full lifecycle and was not reset at year boundaries. For CSEI evaluation, the climate term was fixed at C_J=1.0. Defining $\sigma_8(H-1) = [1 + \exp\{-8(H-1)\}]^{-1}$, the analytical Jacobian was

$$J_{11} = -k_D + 3q_D D^2 - c_H s(H-1), \qquad J_{12} = 0, \qquad J_{13} = c_H \sigma_8(H-1)(1-D)$$

$$J_{21} = 2c_D D(1-L), \qquad J_{22} = -k_L - c_D D^2 - c_{LH} s(H-1), \qquad J_{23} = c_{LH} \sigma_8(H-1)(1-L)$$

$$J_{31} = 0, \qquad J_{32} = \frac{C_J(u_1 + 2u_2 L)}{COP_{rel}\tau}, \qquad J_{33} = -\frac{1}{\tau}$$

Table A1. Reference state-model parameter distributions.

Parameter	Mean	CV	Distribution
k_D	0.01382 year^{-1}	15%	Lognormal
q_D	0.03911 year^{-1}	20%	Lognormal
c_H	0.00400 year^{-1}	20%	Lognormal
k_L	0.00986 year^{-1}	15%	Lognormal
c_D	0.02046 year^{-1}	20%	Lognormal
c_{LH}	0.00139 year^{-1}	20%	Lognormal

Table A1. (continued)

Parameter	Mean	CV	Distribution
u_1	1.3919	12%	Lognormal
u_2	4.3504	15%	Lognormal
τ	0.4585 year	10%	Lognormal

Intervention time was restricted to the monthly grid from 5 to 22 years. Decision ranges were r_D=0.15–0.55, r_L=0.15–0.50 and r_{COP}=0.10–0.35. The candidate design contained 1,024 scrambled Sobol vectors generated with seed 20260729. Training used 256 common-random scenarios generated with seed 20260730. Final evaluation used 5,000 independent holdout scenarios generated with seed 20260728. Final feasibility was assessed using one-sided 95% exact upper confidence bounds, and $CVaR_{95}$ confidence intervals were obtained from 2,000 bootstrap resamples. All state simulations used monthly fourth-order Runge–Kutta integration. Expected-value confidence intervals used the student-t formula; $CVaR_{95}$ intervals used 2,000 percentile bootstrap resamples with replacement and seed 20260802. The nested search-stability audit extended the same scrambled Sobol sequence to 2,048 candidates.

Table A2. Optimisation, holdout and statistical evaluation settings.

Setting	Value
Candidate design	1,024 scrambled Sobol vectors
Candidate seed	20260729
Training scenarios	256 common-random scenarios
Training seed	20260730
Holdout scenarios	5,000 independent scenarios
Holdout seed	20260728
CVaR level	0.95
CVaR bootstrap resamples	2,000
Integration	Monthly fixed-step RK4
CVaR bootstrap seed	20260802
Training reliability/load limit	Maximum 7 violations of 256
Training persistence limit	Maximum 12 violations of 256

Search-set coverage was assessed using normalised hypervolume and inverted generational distance (Zitzler and Thiele 1999; Ishibuchi *et al.* 2015). For a non-dominated set S and reference point r, hypervolume was

$$\mathrm{HV}(S;\, r) = \lambda_d\left[\bigcup_{x\in S}\prod_{k=1}^{d}\left[f_k(x),\, r_k\right]\right]$$

With the 2,048-candidate non-dominated set R as the post-hoc reference set, inverted generational distance was

$$\mathrm{IGD}(S,\, R) = \frac{1}{|R|}\sum_{y\in R}\min_{x\in S}\|y - x\|_2$$

All three objectives were normalised before calculation. Hypervolume used reference point (1.1,1.1,1.1). The 2,048-candidate normalisation ranges and reference set were used only for post-hoc

search diagnostics and did not alter simulation, feasibility screening or selection of the reported 1,024-candidate strategy. The intervention candidates were generated with a scrambled Sobol sequence (Owen 1998).

Table A3. Nested Sobol search-stability diagnostics on the training scenario set.

Candidates	Feasible	Non-dominated	Selected year	Candidate ID	Hypervolume	IGD
256	5	4	12.4167	60	0.5730	0.1849
512	9	6	12.9167	492	0.7454	0.1558
1,024	20	9	12.9167	492	0.8242	0.0941
2,048	44	17	12.9167	492	0.8796	0.0000

Objectives were normalised using the feasible 2,048-candidate set. Hypervolume used reference point (1.1,1.1,1.1), and IGD was calculated against the 2,048-candidate non-dominated set. Candidate 492 is t_I=12.9167 years, r_D=0.543905, r_L=0.467926 and r_{COP}=0.339490.

Table A4. Preference-weight sensitivity of the representative compromise selection.

Cost weight	Carbon weight	Instability weight	Selected ID	Selected year
0.3333	0.3333	0.3333	492	12.9167
0.4000	0.3000	0.3000	492	12.9167
0.3000	0.4000	0.3000	492	12.9167
0.2800	0.3200	0.4000	492	12.9167
0.2000	0.3000	0.5000	492	12.9167

Declarations

Author contributions. Seyed Reza Samaei: Conceptualisation, methodology, software, formal analysis, investigation, data curation, validation, visualisation, writing – original draft, writing – review and editing, and project administration. Assed N. Haddad: Methodology, validation, supervision, and writing – review and editing. James Riffat: Conceptualisation, methodology, validation, supervision, project administration, writing – original draft, and writing – review and editing. All authors contributed to the interpretation of the results, critically reviewed the manuscript, approved the final submitted version, and agreed to be accountable for the integrity of the work.

Funding. This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors. The work was completed using the authors' existing academic and computational resources.

Conflicts of interest. The authors declare that they have no competing financial or non-financial interests that could have influenced the research, analysis, interpretation, or reporting of the findings presented in this manuscript.

Data availability. The public energy data used to define the reference load profile are available through the Building Data Genome Project 2 associated with Miller et al. (2020), DOI: 10.1038/s41597-020-00712-x. Appendix A reports the state-model parameters, uncertainty distributions, decision bounds, random seeds, statistical evaluation settings and search diagnostics required to interpret the optimisation. The processed inputs, candidate matrix, training and holdout outputs, numerical implementation and figure-generation scripts are available from the corresponding author upon reasonable request.

Ethics. This study is based exclusively on synthetic operational data, mathematical modelling, and computational simulation. It did not involve human participants, animals, biological materials,

personal data, or identifiable operational records. Ethical approval and informed consent were therefore not required.

Editorial independence. One or more authors of this paper serve as editors of this journal. To ensure editorial independence, these authors were fully recused from the handling of this manuscript. They took no part in the selection of reviewers, the peer-review process, or the editorial decision, which were managed independently by another editor with no involvement in the work.

References

Boje, Calin; Guerriero, Annie; Kubicki, Sylvain; Rezgui, Yacine (2020) 'Towards a semantic Construction Digital Twin: Directions for future research', *Automation in Construction*, 114, pp. 103179. Available at: https://doi.org/10.1016/j.autcon.2020.103179

Clopper, C. J.; Pearson, E. S. (1934) 'The use of confidence or fiducial limits illustrated in the case of the binomial', *Biometrika*, 26(4), pp. 404-413. Available at: https://doi.org/10.1093/biomet/26.4.404

Deng, Min; Menassa, Carol C.; Kamat, Vineet R. (2021) 'From BIM to digital twins: a systematic review of the evolution of intelligent building representations in the AEC-FM industry', *Journal of Information Technology in Construction*, 26, pp. 58-83. Available at: https://doi.org/10.36680/j.itcon.2021.005

European Committee for Standardization (2011) 'EN 15978:2011 Sustainability of construction works—Assessment of environmental performance of buildings—Calculation method', *Brussels: CEN*.

Evins, Ralph (2013) 'A review of computational optimisation methods applied to sustainable building design', *Renewable and Sustainable Energy Reviews*, 22, pp. 230-245. Available at: https://doi.org/10.1016/j.rser.2013.02.004

Hamdy, Mohamed; Hasan, Ala; Siren, Kai (2013) 'A multi-stage optimization method for cost-optimal and nearly-zero-energy building solutions in line with the EPBD-recast 2010', *Energy and Buildings*, 56, pp. 189-203. Available at: https://doi.org/10.1016/j.enbuild.2012.08.023

International Organization for Standardization (2017) 'ISO 15686-5:2017 Buildings and constructed assets—Service life planning—Part 5: Life-cycle costing', *Geneva: ISO*.

Ishibuchi, Hisao; Masuda, Hiroyuki; Tanigaki, Yuki; Nojima, Yusuke (2015) 'Modified Distance Calculation in Generational Distance and Inverted Generational Distance', *Lecture Notes in Computer Science*. Springer International Publishing, pp. 110-125. Available at: https://doi.org/10.1007/978-3-319-15892-1_8

Kapteyn, Michael G.; Pretorius, Jacob V. R.; Willcox, Karen E. (2021) 'A probabilistic graphical model foundation for enabling predictive digital twins at scale', *Nature Computational Science*, 1(5), pp. 337-347. Available at: https://doi.org/10.1038/s43588-021-00069-0

Luo, Linxi; Wei, Hailu; Lin, Ziqi; Wu, Jiyuan; Wang, Wei; Sun, Yongjun (2025) 'Multi-objective optimal energy-efficient retrofit determination using hybrid urban building energy model: Considering uncertainties between models', *Building Simulation*, 18(1), pp. 183-206. Available at: https://doi.org/10.1007/s12273-024-1206-6

Onweh, Chinedu C.; Al-Habaibeh, Amin; Manu, Emmanuel (2025) 'A Review of Energy Efficiency Strategies in Smart Buildings: Integrating Occupant Comfort, HVAC Optimisation, and Building Automation', *Research and Reviews in Sustainability*, 1, pp. 48-60. Available at: https://doi.org/10.5334/rss.9

Owen, Art B. (1998) 'Scrambling Sobol' and Niederreiter–Xing Points', *Journal of Complexity*, 14(4), pp. 466-489. Available at: https://doi.org/10.1006/jcom.1998.0487

Rasheed, Adil; San, Omer; Kvamsdal, Trond (2020) 'Digital Twin: Values, Challenges and Enablers From a Modeling Perspective', *IEEE Access*, 8, pp. 21980-22012. Available at: https://doi.org/10.1109/ACCESS.2020.2970143

Razak, Tajul Rosli; Kutlu, Cagri; Zheng, Tianhong; Jarimi, Hasila; Su, Yuehong; Riffat, Saffa; Jayakumar, Preethi; Shah, Dhruv (2026) 'From Experiment to Deployment: An Integrated Framework for Predictive Modelling and Visualisation of Thermoelectric HVAC Systems', *Artificial Intelligence for Sustainable Cities*, 1, pp. 137-156. Available at: https://doi.org/10.65582/aifsc.2026.008

Rockafellar, R. Tyrrell; Uryasev, Stanislav (2000) 'Optimization of conditional value-at-risk', *The Journal of Risk*, 2(3), pp. 21-41. Available at: https://doi.org/10.21314/JOR.2000.038

Rockafellar, R.Tyrrell; Uryasev, Stanislav (2002) 'Conditional value-at-risk for general loss distributions', *Journal of Banking & Finance*, 26(7), pp. 1443-1471. Available at: https://doi.org/10.1016/s0378-4266(02)00271-6

Samaei, Seyed Reza; Riffat, James (2026) 'From Optimization to Stability Preservation: A Self-Healing, Action-Bearing Digital Twin for Climate-Stressed Building Energy Systems', *Green Technology & Innovation*, 2, pp. 150-177. Available at: https://doi.org/10.65582/gti.2026.008

Shu, Lei; Hong, Tianzhen; Sun, Kaiyu; Zhao, Dong (2025) 'Framework to select robust energy retrofit measures for residential communities', *Energy and Buildings*, 327, pp. 115077. Available at: https://doi.org/10.1016/j.enbuild.2024.115077

Taileb, Ali; Sherzad, Mohammed Fareed (2026) 'Comprehensive Energy and Decarbonisation Audit of a Multi-Unit Residential Building in Toronto: Methods, Findings, and Recommendations', *Future Cities and Environment*, 12, pp. 18. Available at: https://doi.org/10.70917/fce-2026-014

Zhang, Haonan; Hewage, Kasun; Bakhtavar, Ezzeddin; Sun, Qingqing; Sadiq, Rehan (2025) 'Robust building energy retrofit evaluation under uncertainty: An interpretable machine learning approach', *Energy Conversion and Management*, 345, pp. 120362. Available at: https://doi.org/10.1016/j.enconman.2025.120362

Zitzler, E.; Thiele, L. (1999) 'Multiobjective evolutionary algorithms: a comparative case study and the strength Pareto approach', *IEEE Transactions on Evolutionary Computation*, 3(4), pp. 257-271. Available at: https://doi.org/10.1109/4235.797969

HOW TO CITE THIS ARTICLE

Samaei, S. R., Naked Haddad, A. and Riffat, J. (2026). Instability-Constrained Lifecycle Optimisation of Low-Carbon Buildings under Structural–Envelope Degradation and Energy Feedback. Green Technology & Innovation, 2, 387-409. https://doi.org/10.65582/gti.2026.017

COPYRIGHT AND LICENCE

© 2026 The Author(s). This is an open access article distributed under the terms of the Creative Commons Attribution (CC BY 4.0) licence, which permits unrestricted use, distribution and reproduction in any medium, provided the original work is properly cited.