A Nested Monte Carlo Discrete-Event Framework for Multi-Module SMR Availability and Project Economics

Lachezar Dimitrov

ETH Zurich, Department of Mechanical and Process Engineering, Zurich, Switzerland,
ldimitrov@ethz.ch

Massachusetts Institute of Technology, Department of Nuclear Science and Engineering, Cambridge, MA, USA

Abstract: Plant availability is treated as a scalar input in most nuclear techno-economic analyses. This paper develops it instead as an outcome of failures, support cascades, common-cause events, technical-specification (TS) timers, finite-crew repair with explicit procurement and overhead delays, and operating policy on a multi-module small modular reactor (SMR) plant. A discrete-event simulator with these mechanisms is embedded in a nested Monte Carlo design: the outer loop samples epistemic uncertainty in component failure rates and common-cause beta factors, the inner loop samples aleatory variability in event timing conditional on those parameters. The ensemble returns distributions of lifetime capacity factor, forced-outage hours, present-value corrective-repair spend, and revenue-side summaries including revenue-at-risk (tail shortfalls in discounted revenue relative to the median). The framework is exercised on a NuScale-inspired twelve-module light-water configuration over a 60-year horizon. On the baseline ensemble, lifetime site capacity factor has median 94.0 % with a P5–P95 spread of 4.3 percentage points. Lifetime forced-outage hours are strongly right-skewed (median 6 594 h, P95 21 530 h). The scatter among inner histories at one reliability draw is larger than the shift among outer-draw means, so lifetime availability appears aleatory-dominated within this ensemble under the present priors. Revenue-at-risk at the 5 % level is about \$754 M. A $\lambda/5$ reliability-lever study compares three component families: per-module high-rate equipment, site-wide control with fleet leverage, and low-rate shared support. These return distinct response patterns across median throughput, lower-tail capacity factor, and repair spend, interpreted as a simulation-based analogue, in purpose, to classical PRA importance measures such as RAW, RRW and Fussell–Vesely. The numerical values are illustrative of the modelled scenario. The contribution of the paper is the framework itself: an integrated path from component-level reliability and staged repair logistics to distributional generation-economics outcomes for multi-module SMR plants.

Keywords: probabilistic risk assessment, multi-module SMR, discrete-event simulation, nested Monte Carlo, availability, revenue-at-risk, importance measures.

---

1. INTRODUCTION

Plant availability is central to the economic case for new nuclear build, especially for multi-module small modular reactors (SMRs): shared support systems, staggered refueling, and pooled repair resources couple module-level reliability to site-level output in ways a single-unit reading does not reproduce [1]. Investment decisions therefore depend on the *range* of plant performance a design and operating posture can deliver, not on a single deterministic capacity factor. The same point applies in service: reliability upgrades, repair-crew capacity, spare-parts inventory and operating-policy strictness all reshape that range, but their value is invisible when availability is itself a scalar input.

Three traditions address parts of this chain. Classical PRA provides mature representations of failure logic, dependencies, common-cause failure (CCF) and uncertainty treatment [2, 3, 4], but its endpoint is safety risk rather than generation economics. Techno-economic assessment evaluates levelized cost and net present value, yet typically takes availability as an exogenous scalar [5]. Reliability and maintenance studies model failures and repairs [6, 7], and generation risk assessment connects component reliability to plant-level availability [10]. The resulting operating histories, however, are rarely propagated into distributional discounted cash flows on a multi-module plant. The missing link is a way to propagate component reliability and resource-constrained repair into plant-level economic distributions.

The framework reported here couples a discrete-event plant simulator with support-dependency propagation, TS timers, derates, CCF, standby fail-to-start, and staged finite-crew repair (procurement, overhead, hands-on execution) to a nested Monte Carlo design separating epistemic reliability-parameter draws from aleatory event histories. It returns distributions of capacity factor, forced-outage hours, present-value corrective-repair spend, discounted revenue, and revenue-at-risk. A simulation-based $\lambda/5$ study is used, in the spirit of classical RAW and RRW importance measures [4, 6], to compare lever effects on operational and economic distributions. The framework retains PRA vocabulary for components, dependencies, CCF and importance-style ranking, but applies it to generation and economic outcomes rather than accident-sequence frequencies [8, 9, 10].

2. SIMULATION FRAMEWORK

2.1. Plant abstraction

The plant is a directed graph of modelled elements carrying explicit operating states. Dominant components are represented individually; less consequential subsystems are grouped into a *supercomponent* when internal detail would not change the modelled derate, trip, or repair behaviour. Each element references one *function group*, the unit on which operating policy acts. Three labels couple the inventory to that policy: `CRIT` forces an immediate trip when unavailable; `TS` starts a technical-specification timer; and `DER` contributes to a derate envelope. Support dependencies are encoded as a forward graph with explicit AND/OR aggregation. A change at any element propagates through this graph in clock time, rather than through repeated static top-event evaluations.

2.2. Dynamic mechanisms

Failures. Each element samples its next time-to-failure from a constant-rate hazard. Standby components carry a fail-to-start probability on demand, and latent standby failures remain hidden until surveillance, demand, or operating-cycle inspection. CCF within a defined group is represented by a single-parameter β-factor model [11, 12], a standard public-data-compatible PRA parameterisation.

Operating policy. After each event, unavailable `CRIT` elements force an immediate site or module trip. Function-group counts then trigger derates or start `TS` timers; timer expiry before restoration forces a controlled shutdown. The policy table is a scenario input, so the same plant can be evaluated under different operating tolerances.

Repair pipeline. Restoration is not collapsed into a single mean-time-to-repair scalar. Each repair passes through three sequential stages: spare-parts *procurement* (an in-stock probability and a sampled lead time if the part is not on the shelf), administrative *overhead* (work-order release, hold-points, mobilisation), and *hands-on execution* by a finite crew pool that competes across modules. Online and outage-surge crew counts are explicit parameters, so persistent clustering of failures shows up as waiting time rather than as instantaneous repair. The same plant can therefore be evaluated under different repair-resource postures (crew sizing together with spare-parts and procurement logistics) combined into bundles, without changing the underlying failure model.

2.3. Nested Monte Carlo

The simulator is wrapped in a nested Monte Carlo design [8, 9, 7]. The outer (*epistemic*) loop draws uncertain reliability parameters (failure-rate priors and CCF β factors); the inner (*aleatory*) loop runs independent plant histories conditional on that draw. Percentiles over $N_{\mathrm{epi}} \times N_{\mathrm{alea}}$ histories therefore combine parameter uncertainty with event-process variability, and the nested structure permits the two contributions to be compared.

2.4. Performance metrics and economic outputs

Every run produces an hour-resolved site capacity trace $p_{\max}(t) \in [0, 1]$, with lifetime site capacity factor

$$\mathrm{CF} = \frac{1}{T} \int_0^T p_{\max}(t)\,\mathrm{d}t, \tag{1}$$

where T is the horizon. Forced and planned outages are recorded as separate envelope durations, distinct from the capacity-loss integral that also credits partial derates. Revenue is computed from the same trace under a must-run convention with an exogenous price path; no dispatch coupling, balancing penalties, or price-responsive curtailment is modelled. Discounted lifetime revenue is therefore nearly linear in lifetime CF; the ensemble slope

$$\eta_{\mathrm{av}} = \frac{\mathrm{d}\,\mathrm{Rev}^{\mathrm{PV}}}{\mathrm{d}\,\mathrm{CF}} \quad [\$ \text{ per pp CF}] \tag{2}$$

is reported only to put CF shifts in the same dollar numéraire as repair spend and revenue-at-risk; it is an accounting-implied mapping, not an independent economic response. Downside exposure is reported as revenue-at-risk [14] at the 5 % and 10 % lower-tail percentiles, $\mathrm{RaR}_q = R_{50} - R_q$, where R_q is the q-th percentile of lifetime discounted revenue.

3. CASE STUDY: TWELVE-MODULE NUSCALE-INSPIRED PLANT

The framework is exercised on an illustrative NuScale-inspired twelve-module light-water configuration with site capacity $12 \times 77 = 924$ MWe, informed by the public NuScale Design Certification Application material [13]. It follows the originally proposed twelve-module layout while applying the later 77 MWe per-module rating; it is not a vendor-certified model of the licensed six-module US460 design or of any operating plant.

The inventory contains 375 elements: 30 per-module elements replicated across 12 modules and 15 shared site-level elements (electrical, cooling/support, control/chemistry and virtual aggregators for compound AND/OR logic). Shared support assets are the main route through which one element failure can derate or trip all modules.

Reliability and repair inputs are anchored in public reliability data [12] and generic-LWR templates, not proprietary plant records. The absolute values should therefore not be read as plant forecasts; the useful information is how the distributions move when repair logistics or selected reliability assumptions are changed. Three studies are reported:

- a 60-year nested-MC *baseline* of $N_{\mathrm{epi}} \times N_{\mathrm{alea}} = 10 \times 20 = 200$ runs under operating Policy M and the Balanced repair-resource bundle;
- a 10-year *bundle sensitivity* across Lean, Balanced and ASAP repair-resource postures, each at $5 \times 10 = 50$ runs, varying crew counts and spare-parts/procurement logistics at fixed plant, schedule and policy; and

- a 10-year *reliability-lever study* dividing one component-family failure-rate central value by five at a time (turbine–generator train, TGT; plant control system, PCS; reactor component cooling water, RCCW), each at $5 \times 10 = 50$ runs.

Policy M denotes the medium operating-policy setting, with intermediate TS-timer durations and derate thresholds between the conservative and flexible variants in Table 1. The 60-year ensemble supplies lifetime economic summaries; the 10-year ensembles are used for relative screening of repair-resource and reliability-lever shifts at fixed plant, schedule and policy, and are not intended to be directly comparable to the 60-year baseline in absolute terms.

Table 1: Modeled plant, scenario, and ensemble inputs.

Site	12 reactor modules, $12 \times 77 = 924$ MWe; staggered refueling; 60-year economic life.
Inventory	375 modeled elements (360 per-module + 15 shared, including 5 virtual aggregators).
Support graph	Forward-only AND/OR dependencies (AC, DC, cooling, steam path, control).
Failure model	Constant-rate hazard with optional linear aging; standby failures and fail-to-start; latent failures; β-factor CCF.
Repair pipeline	Procurement (in-stock prob., lognormal lead time) + overhead + hands-on execution; finite online and surge crews.
Operating policy	`CRIT` trip / `TS` timer / `DER` derate logic on function groups; policies C/M/F differ in TS durations and derate thresholds.
Bundles	Lean (3 online + 1 surge), Balanced (5+2), ASAP (8+3); spare-parts and procurement logistics tighten from Lean to ASAP jointly with crew size.
Economics	Must-run revenue under exogenous price path; corrective and PM cost at event; fixed and variable O&M; discounted NPV, LCOE, RaR; optional CF–revenue slope η_{av} (reporting numéraire).
Ensemble sizes	60-y baseline: $10\times20 = 200$ runs (Policy M, Balanced); 10-y sensitivities: $5\times10 = 50$ runs per scenario.

4. RESULTS

4.1. Single-run mechanism illustration

Figure 1 shows one 10-year Balanced, Policy M site-power realisation. Each module refuels on a two-year cycle, staggered evenly across the twelve modules, so the regular short dips from $p_{\max} = 1$ to about $11/12$ correspond to one module being offline for refueling. Deeper drops are forced trips or TS-timer expiries, and intermediate plateaus are derates from function-group counts. Each step is a sampled event propagated through the support graph, policy table and finite repair-crew pool; this hourly trace is the object aggregated in the following results.

4.2. Lifetime distribution and uncertainty decomposition

Across the 200 sixty-year baseline runs, lifetime site capacity factor has median 94.0 %, $P_5 = 91.0$ %, and $P_{95} = 95.3$ % (a 4.3 pp spread; Figure 2(a)). Lifetime forced-outage hours are strongly right-skewed (median 6 594 h, $P_{95} = 21\,530$ h). Every run records at least one forced trip, so the relevant observable is outage accumulation, not trip probability. Planned-outage envelopes are nearly deterministic at about 1 296 h, so forced-outage variability is the dominant source of run-to-run dispersion in revenue across the ensemble.

Figure 2(b) plots the same ensemble by epistemic index. Within this ensemble the scatter among inner histories at one reliability draw is larger than the shift among outer-draw means, so lifetime availability appears aleatory-dominated under the present priors. Sharpening priors alone would not close the band;

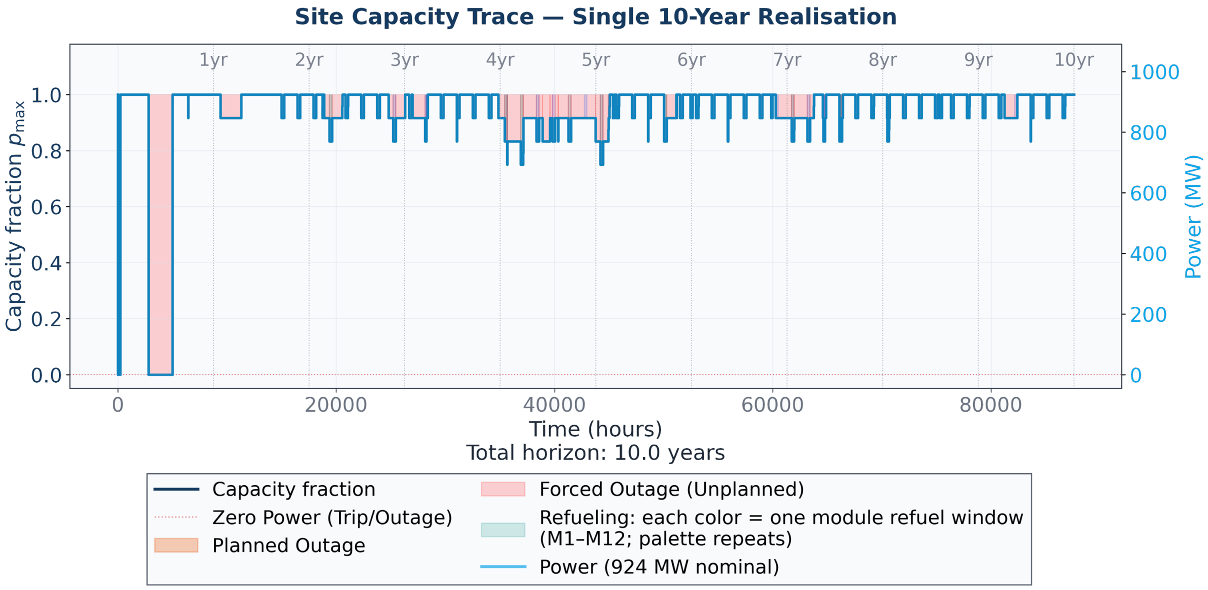

Figure 1: Site power for one Balanced, Policy M, ten-year realisation of the twelve-module configuration. Short steps: staggered refueling. Deep drops: forced trips or TS-timer expiries. Intermediate plateaus: derates from function-group counts.

operational and structural levers are examined next.

4.3. Repair-logistics tail-shaping

With plant, schedule and Policy M fixed, the Lean, Balanced and ASAP bundles vary crew counts (Lean 3+1, Balanced 5+2, ASAP 8+3) jointly with spare-parts and procurement logistics. They therefore encode joint repair-resource postures rather than a single scalar lever.

The main effect appears in the forced-outage tail (Figure 3): P_{95} equals 6 691 h under Lean, 4 258 h under Balanced and 1 469 h under ASAP, so the Balanced→ASAP tail compresses by roughly a factor of three. Median CF barely separates Lean and Balanced (94.2 % vs 94.4 %) and then rises to 96.9 % under ASAP. Under the present inputs, procurement lead time is the longest of the three pipeline stages and therefore dominates the bundle ranking: tightening it from Lean to ASAP shortens failure-to-restoration windows and truncates the forced-outage tail. Crew counts act as a secondary lever: more crews dissolve queueing when failures cluster, but their typical-year effect on CF is modest in this parameterisation. Repair-resource bundles therefore act mainly as a *tail lever*.

4.4. A distributional reliability-importance study

For the reliability-lever study, one component-family failure-rate central value is divided by five while all other inputs remain at the 10-year Balanced, Policy M baseline. The three reported families span distinct roles: per-module high-rate equipment (TGT), site-wide control with fleet leverage (PCS), and low-rate shared support (RCCW). Shifts in 10-year CF, forced-outage hours and present-value corrective-repair spend are reported in Table 2 and Figure 4.

Throughout, H_q^{f} denotes the q-th percentile of lifetime forced-outage hours per run and $C_{\mathrm{rep},q}^{\mathrm{PV}}$ the q-th percentile of discounted corrective-repair spend. Lever shifts Δ compare the study ensemble to the 10-year Balanced, Policy M baseline: $\Delta\mathrm{CF}_q$ entries are absolute percentage-point differences (pp), whereas ΔH_q^{f} and $\Delta C_{\mathrm{rep},q}^{\mathrm{PV}}$ are relative percent changes against the corresponding baseline percentile.

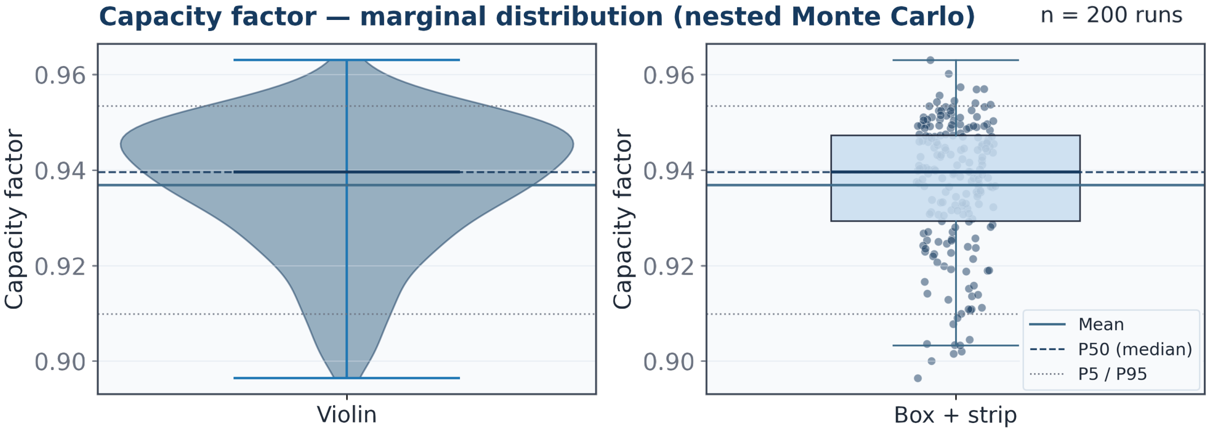

(a) Lifetime site capacity factor.

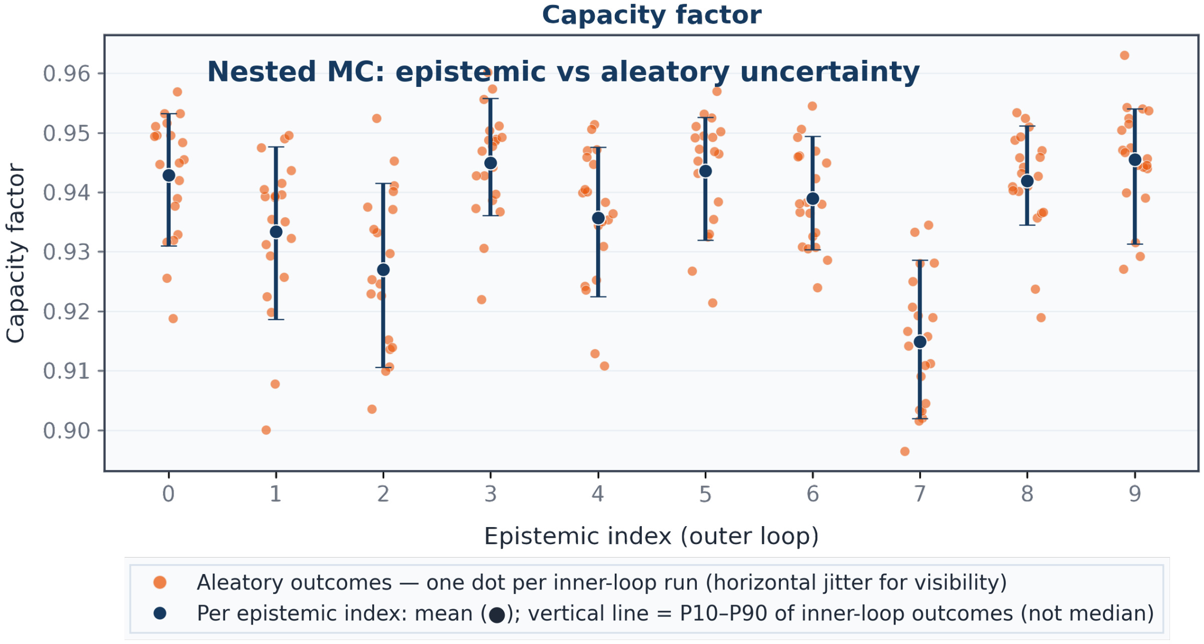

(b) Lifetime CF vs epistemic index (nested MC).

Figure 2: Lifetime baseline ensemble (200 runs, 60 y, Policy M, Balanced bundle). (a) Marginal lifetime CF distribution. (b) The same CF outcomes by epistemic index; orange dots are inner replicas and blue markers show the outer-draw mean with P10–P90 inner spread.

Table 2: Ensemble shifts relative to the 10-year Balanced, Policy M baseline ($N = 50$ per study). CF shifts are percentage points; hour and cost shifts are percent. Positive CF and negative hours/spend are beneficial. Baseline anchors: $\mathrm{CF}_{50} = 94.38\,\%$, $\mathrm{CF}_5 = 91.04\,\%$, $H^{\mathrm{f}}_{50} = 1{,}401$ h, $H^{\mathrm{f}}_{95} = 4{,}258$ h, $C^{\mathrm{PV}}_{\mathrm{rep},50} = 6.8$ MUSD.

Lever	$\Delta\mathrm{CF}_{50}$ (pp)	$\Delta\mathrm{CF}_5$ (pp)	$\Delta H^{\mathrm{f}}_{50}$ (%)	$\Delta H^{\mathrm{f}}_{95}$ (%)	$\Delta C^{\mathrm{PV}}_{\mathrm{rep},50}$ (%)
TGT/5	+2.05	+1.03	−54.8	+3.9	−38.1
PCS/5	+1.12	+1.30	−99.4	−34.7	+4.7
RCCW/5	+1.12	−1.65	−80.5	+35.1	+5.9

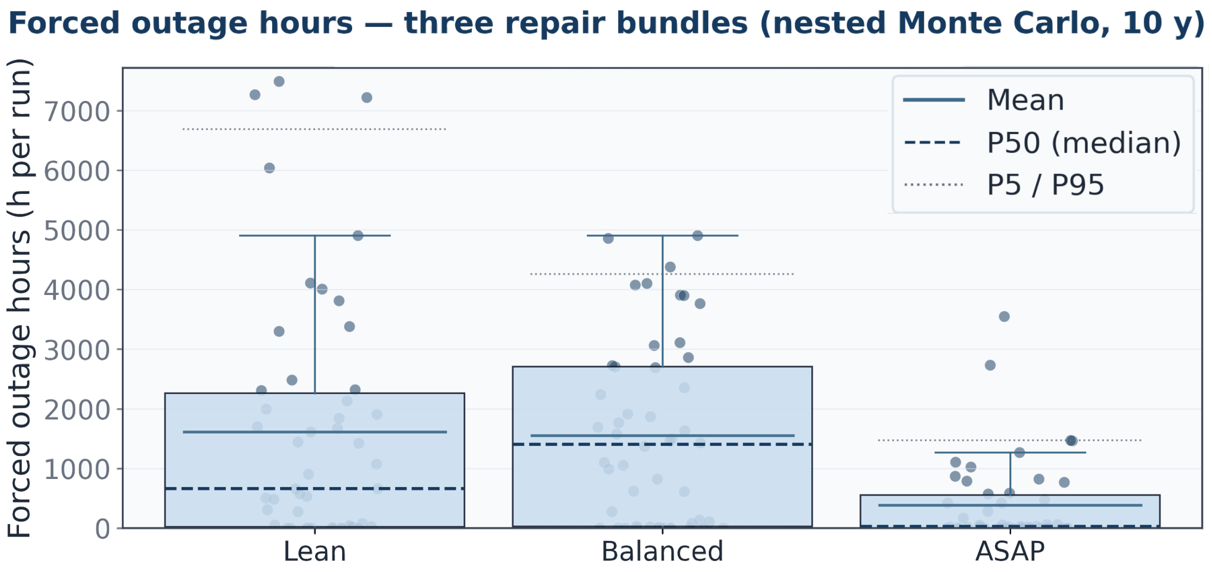

Figure 3: Ten-year forced-outage hours under three repair-resource bundles on one axis with a shared vertical scale (box plot with jittered run outcomes per bundle; 50 runs each, Policy M). Bundles vary online + surge crew counts (Lean 3+1, Balanced 5+2, ASAP 8+3) jointly with spare-parts and procurement logistics. The P95 tail compresses strongly from Lean to ASAP while medians shift less.

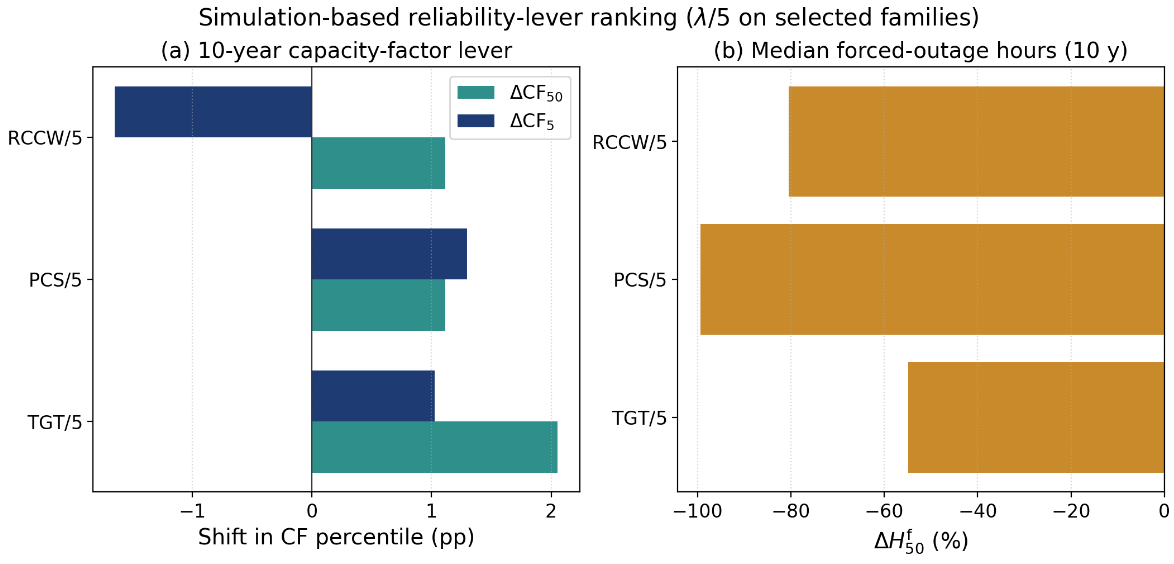

Figure 4: Reliability-lever ranking ($\lambda/5$ on TGT, PCS, RCCW; 10 y, Policy M, Balanced bundle, 5×10 runs per study). (a) $\Delta\mathrm{CF}_{50}$ and $\Delta\mathrm{CF}_{5}$. (b) percent change in median forced-outage hours. Repair-spend shifts are in Table 2. RCCW tail-direction changes at this sample size are not interpreted as robust.

Three response patterns emerge. TGT/5 is a throughput lever: it gives the largest median CF gain (+2.05 pp) and reduces median present-value corrective spend by about 38 %, because per-module TGT failures dominate baseline forced-outage hours and crew demand. PCS/5 is a tail-risk lever: it gives the strongest left-tail CF improvement (+1.30 pp on CF_5) and nearly eliminates median forced-outage hours (−99.4 %). PCS failures are rare but fleet-wide, so reducing their rate removes the deepest years without strongly changing the typical year. RCCW/5 is a low-rate shared-support lever: its median CF gain is comparable to PCS, but at $N = 50$ its tail sign is mixed and should be read as sample noise rather than as a robust adverse tail effect. The remaining small $\lesssim$ 5 % non-target sign changes in Table 2, namely TGT/5's slight rise in $\Delta H^{\mathrm{f}}_{95}$ and PCS/5's slight rise in $\Delta C^{\mathrm{PV}}_{\mathrm{rep},50}$, sit within sampling noise at this ensemble size and are not interpreted as substantive lever effects.

The ranking therefore depends on the output of interest: a median-throughput or spend-oriented stakeholder would choose TGT, while a tail-risk stakeholder would choose PCS. A single scalar ranking is therefore insufficient for this investment-relevant question; the distributional ranking exposes a multi-objective trade-off that Boolean RAW or RRW reports do not.

4.5. Revenue-risk results

The economic layer maps availability to revenue-side summaries. On the 60-year baseline, discounted lifetime actual revenue has median \$23.5 B and a P_5–P_{95} spread of about \$1.1 B. Under the must-run revenue model used here the CF–revenue map is close to linear, with ensemble slope $\eta_{\mathrm{av}} \approx$ \$249 M per +1 pp CF; this is the accounting-implied scale at which CF shifts translate into dollars in the same numéraire as repair spend and RaR. With dispatch coupling, balancing penalties or price-responsive curtailment the mapping would weaken in low-price hours.

Downside exposure is reported at two thresholds: $\mathrm{RaR}_5 \approx$ \$754 M captures the severe lower-tail revenue shortfall relative to the median, while $\mathrm{RaR}_{10} \approx$ \$603 M captures the moderate downside. The ratio $\mathrm{RaR}_5/\mathrm{RaR}_{10} \approx 1.25$ is a compact indicator of right-tail thickness in lifetime revenue losses, which are themselves driven almost entirely by forced outage (Section 4.2). Because corrective-repair spend $C^{\mathrm{PV}}_{\mathrm{rep}}$ is reported in the same numéraire as RaR, a stakeholder can compare the order of magnitude of logistics or reliability investments against avoided revenue loss by reading Table 2 alongside these numbers; it is a screening comparison, not a decision-grade investment analysis.

Absolute NPV and LCOE are not reported because they depend on illustrative CAPEX and price inputs that do not derive from a vendor calibration. What the framework supports is the distributional shape of these revenue-side outputs and how they shift across scenarios.

5. DISCUSSION

The case study points to two non-substitutable management levers. Repair logistics (crew capacity plus spare-parts/procurement posture) mainly changes the tail of forced-outage hours. Reliability improvement changes distribution shape in component-specific ways: TGT shifts median CF and corrective spend, while PCS mainly improves the lower tail without strongly changing the typical year. A site that invests in median capacity through reliability upgrades can still be exposed in the tail through weak logistics, and vice versa.

The reliability-lever study is in the spirit of classical PRA importance measures, but defined by simulation-based shifts in empirical distributions of CF, forced-outage hours and corrective-repair spend rather than by changes in Boolean top-event probability [4, 6]. It is complementary to licensing-grade PRA, not a replacement for it.

Five limitations should be kept in view. First, reliability and repair inputs are stylised public/generic data, not proprietary operating records; absolute outcomes are scenario-dependent. Second, each element currently carries one aggregated failure mode, so mode-specific failure and repair distributions are the main path to sharper attribution. Third, revenue is must-run against an exogenous price path, with no dispatch, balancing penalties or curtailment. Fourth, sample sizes are finite: central percentiles stabilise more readily than deep forced-outage and CF tails, especially for low-rate shared-support levers such as RCCW. Fifth, the model is not validated against operating data for a multi-module SMR plant, since no such plant is yet in commercial operation; results should accordingly be read as a structured exploration of the case study, not as an empirically calibrated forecast.

6. CONCLUSION

A discrete-event simulation framework embedded in a nested Monte Carlo design was demonstrated on a NuScale-inspired twelve-module light-water configuration. The 60-year baseline has a P_5–P_{95} capacity-factor spread of 4.3 pp and appears aleatory-dominated within this nested ensemble under the present priors. Closing this band would require operational or structural interventions, not only sharper parameter priors. Repair-resource bundles mainly act on the forced-outage tail: the 10-year P_{95} falls by roughly a factor of three from Balanced to ASAP at fixed plant and policy. A $\lambda/5$ reliability-lever study returns throughput, tail-risk and low-rate shared-support response patterns, an importance-style ranking expressed through shifts in operational and economic distributions rather than through a Boolean top event. Revenue-at-risk at the 5 % level is about $754 M. The case study is used to demonstrate the framework rather than to forecast the performance of a specific deployed plant; the main contribution is the framework linking component reliability and staged repair logistics to distributional generation-economics outcomes. The plant model, operating-policy table, CCF groups, and repair-resource bundles are configuration inputs, so the framework can be re-parameterised for other multi-module designs, operating tolerances and logistics postures. The main extensions are plant-specific repair calibration and mode-specific failure/repair distributions.

ACKNOWLEDGEMENTS

The author thanks the MIT Department of Nuclear Science and Engineering for hosting the MIT portion of this research and is particularly grateful to Professor Curtis Smith for mentorship and technical guidance during that stay. The author also thanks Professor Annalisa Manera (ETH Zurich) for supervision of the master's thesis on which this paper draws.

REFERENCES

[1] International Atomic Energy Agency. *Advances in Small Modular Reactor Technology Developments*, 2020 Edition, Vienna, 2020.

[2] International Atomic Energy Agency. *Development and Application of Level 1 Probabilistic Safety Assessment for Nuclear Power Plants*, IAEA Safety Standards Series SSG-3 (Rev. 1), Vienna, 2024.

[3] U.S. Nuclear Regulatory Commission. *PRA Procedures Guide: A Guide to the Performance of Probabilistic Risk Assessments for Nuclear Power Plants*, NUREG/CR–2300, Washington, DC, 1983.

[4] U.S. Nuclear Regulatory Commission. *Fault Tree Handbook*, NUREG–0492, Washington, DC, 1981.

[5] OECD Nuclear Energy Agency / International Energy Agency. *Projected Costs of Generating Electricity, 2020 Edition*, OECD Publishing, Paris, 2020.

[6] M. Rausand and A. Høyland. *System Reliability Theory: Models, Statistical Methods, and Applications*, 2nd ed., Wiley, Hoboken, NJ, 2004.

[7] A. M. Law. *Simulation Modeling and Analysis*, 5th ed., McGraw-Hill, New York, 2014.

[8] U.S. Nuclear Regulatory Commission. *Guidance on the Treatment of Uncertainties Associated with PRAs in Risk-Informed Decision Making*, NUREG–1855, Vol. 1, Washington, DC, 2009.

[9] NASA. *Probabilistic Risk Assessment Procedures Guide for NASA Managers and Practitioners*, NASA/SP–2011–3421, Washington, DC, 2011.

[10] Electric Power Research Institute. *Generation Risk Assessment (GRA) Plant Implementation Guide*, EPRI Report 1008121, Palo Alto, CA, 2004.

[11] U.S. Nuclear Regulatory Commission. *Guidelines on Modeling Common-Cause Failures in Probabilistic Risk Assessment*, NUREG/CR–5485, Washington, DC, 1998.

[12] U.S. Nuclear Regulatory Commission. *Industry-Average Performance for Components and Initiating Events at U.S. Commercial Nuclear Power Plants*, NUREG/CR–6928, Washington, DC, 2007.

[13] NuScale Power, LLC. *NuScale Standard Plant Design Certification Application, Final Safety Analysis Report, Part 2 (Tier 2)*, twelve-module configuration, U.S. Nuclear Regulatory Commission Docket No. 52-048; public version available via the NRC ADAMS document system. Used here only as a public source for the twelve-module layout assumed in the case study; this paper does not represent or evaluate the licensed design.

[14] P. Jorion. *Value at Risk: The New Benchmark for Managing Financial Risk*, 3rd ed., McGraw-Hill, New York, 2007.