Survivability-Adjusted Loss Modeling (SALM)

A Loss-Centric Framework for Decision Risk, Capital Survivability, and Opportunity Cost

Adrian Wood
DealStrike.app
adrian@dealstrike.app
January 9, 2026

Abstract

Traditional investment risk frameworks focus primarily on volatility, expected return, or static downside metrics. While these approaches are useful for portfolio-level comparisons, they fail to capture how and why individual investment decisions fail in practice. In parallel, loss-based frameworks such as FAIR (Factor Analysis of Information Risk) provide a rigorous foundation for decomposing risk into frequency and magnitude but are rarely extended to capital-constrained investment decisions.

This paper introduces Survivability-Adjusted Loss Modeling (SALM), a loss-centric risk framework inspired by FAIR-style risk analysis by incorporating time-dependent loss processes, capital survivability, and counterfactual opportunity cost modeling. SALM treats risk as the probability-weighted set of loss trajectories that degrade or terminate an asset's ability to operate, rather than as a single scalar metric. The framework is particularly applicable to illiquid, leveraged investments such as real estate, where recovery, optionality, and capital exhaustion are dominant drivers of real-world outcomes.

Keywords: Risk modeling, survivability analysis, loss modeling, FAIR, Monte Carlo simulation, opportunity cost, real estate investment, decision risk.

---

1 Introduction

Investment analysis has historically prioritized point estimates: internal rate of return (IRR), net present value (NPV), cap rate, or volatility-adjusted metrics such as Sharpe ratios. While useful, these measures obscure the mechanisms by which investments fail. Two investments with identical IRR profiles may differ substantially in their ability to withstand prolonged stress, absorb losses, or preserve future optionality.

In operational risk domains, the FAIR framework reframed risk as a function of loss event frequency and loss magnitude, enabling structured analysis of how adverse events occur and what they cost.

However, FAIR was designed for organizational assets and static loss events, not for capital-bound investment decisions with path-dependent failure modes.

Survivability-Adjusted Loss Modeling (SALM) bridges this gap by reframing investment risk as a dynamic loss process operating over time under capital constraints. Rather than asking "What is the expected return?" SALM asks:

> What loss processes threaten this decision, how do they interact over time, how long can the system survive them, and what was forfeited by choosing this path?

2 Limitations of Existing Investment Risk Models

2.1 Volatility-Centric Risk

Metrics such as standard deviation or drawdown treat deviations symmetrically and fail to distinguish between recoverable impairment and irreversible loss. Volatility does not model capital exhaustion, refinancing failure, or forced liquidation.

2.2 Scenario and Stress Testing

Traditional stress tests examine isolated shocks but rarely model:

- duration of stress
- compounding effects
- capital depletion
- secondary failures triggered by prior stress
- unique property, market, or tenant-type stresses

Traditional scenario analysis typically assumes:

- instantaneous recovery
- independence between shocks
- infinite capital availability

In practice, stress persists across months or years, and the duration of stress is often more damaging than its initial magnitude.

2.3 Expected Value Fallacy

Expected value collapses multi-modal, asymmetric outcomes into a single number, masking tail risks and survivability constraints that dominate leveraged investments. It aggregates outcomes into a mean, hiding tail-driven failure modes. Two investments with identical expected returns may differ radically in survivability due to differences in leverage, liquidity, and recovery dynamics.

3 Design Principles of SALM

SALM is founded on five principles:

1. **Loss-first modeling:** All risk metrics derive from explicit loss definitions.
2. **Time dependency:** Losses evolve, persist, and compound over time.
3. **Capital survivability:** Failure occurs when capital buffers are exhausted.
4. **Comparative decision context:** Risk must be evaluated relative to forgone alternatives.
5. **Conditional failure pathways:** Secondary losses are triggered by prior losses and cannot occur independently.

4 Core Definitions

4.1 Asset Exposure Unit (AEU)

An Asset Exposure Unit is a capital-bound investment characterized by:

- income generation
- operating expenses
- financing structure
- external dependencies
- finite liquidity buffers

Examples include leveraged real estate assets, development projects, or alternative capital allocations.

In the DealStrike implementation, each simulation path tracks the evolving state of an AEU across monthly time steps, including cash balances, debt service coverage, and cumulative stress exposure.

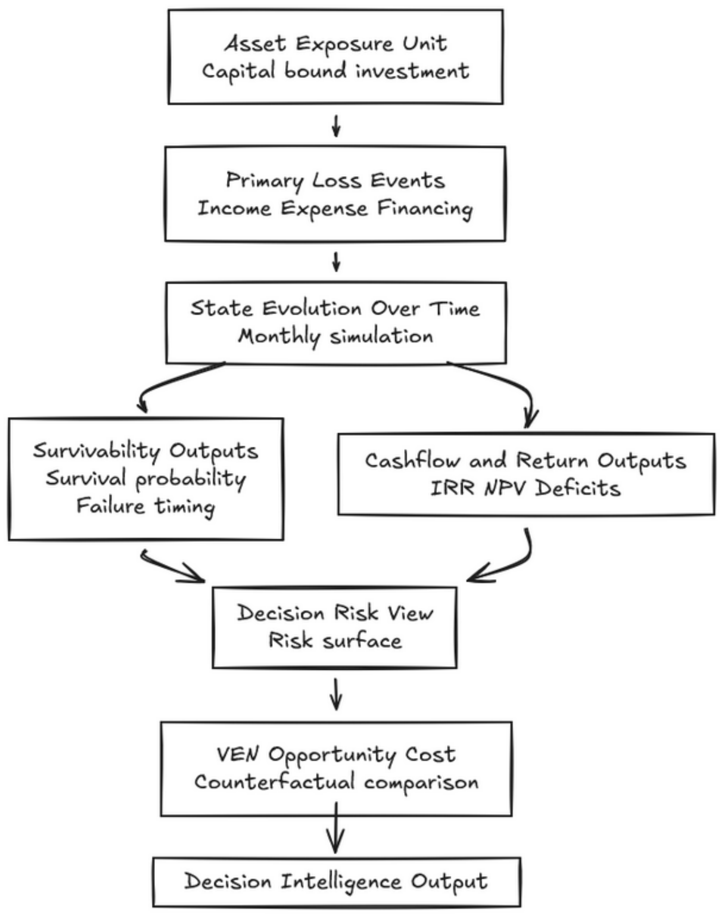

Figure 1: SALM conceptual overview.

4.2 Loss Event

A loss event is any occurrence that degrades the AEU's ability to:

- service obligations
- preserve equity
- maintain optionality
- continue operation

Loss events are continuous and partial, not binary.

5 Loss Event Taxonomy

5.1 Primary Loss Events

Primary losses occur directly to the asset:

- income degradation (vacancy, rent decline)
- expense escalation (tax, insurance, maintenance)
- financing impairment (rate shocks, covenant breaches)
- liquidity drawdown (operating deficits)

These losses are instantiated through monthly hazard sampling. At each time step, the simulation evaluates whether a loss event occurs based on predefined hazard rates. When triggered, the loss event modifies the relevant cash flow components for that month. Importantly, these changes persist across time steps until recovery conditions are met or the simulation terminates.

This approach avoids the unrealistic assumption that losses are immediately reversed and instead allows stress to persist long enough to matter.

5.2 Secondary (Conditional) Loss Events as Emergent Outcomes

Secondary losses occur conditionally, triggered by persistent primary losses, examples include:

- refinancing failure following NOI decline
- forced sale following capital exhaustion

- equity wipeout following prolonged negative cash flow

These losses are path-dependent and cannot occur independently. Critically, these losses are state-dependent. In the current implementation, survivability tracking records the dominant failure cause, but future extensions can explicitly encode conditional triggers.

5.3 Adversarially-Driven Loss Event Parameterization

In SALM, loss events are defined at the conceptual level (income degradation, expense shocks, financing impairment), but their parameterization (hazard rates, severity distributions, persistence) materially determines failure pathways. SALM therefore supports an adversarial layer that proposes targeted perturbations to these parameters based on scenario context. This adversarial parameter control is treated as a risk discovery method that surfaces hidden fragility, rather than as a forecasting mechanism.

6 Loss Event Frequency as Hazard Rates

Loss event frequency is modeled via time-bounded hazard processes at the simulation timestep. In addition to baseline hazard configurations, SALM permits hazard-rate modification by adversarial agents to explore regimes in which hazards cluster, persist, or compound beyond naïve assumptions. This explicitly addresses the under-sampling of rare but plausible hazard combinations that frequently dominate terminal outcomes in leveraged assets.

This approach reflects empirical realities:

- losses cluster temporally
- recovery is slow and capital-dependent
- independent sampling assumptions break down under stress

In DealStrike's Monte Carlo engine, hazard rates govern whether vacancy or expense shocks occur within each month of a stress window, producing realistic loss clustering.

6.1 Implications for Loss Clustering and Tail Risk

Because hazard rates are evaluated repeatedly across time steps, loss clustering emerges naturally. A path that experiences early losses is more likely to experience compounding stress, not because the hazards are explicitly correlated, but because the system state becomes increasingly fragile.

This mechanism captures a key empirical feature of financial failure: losses make future losses more damaging, even if their probability remains unchanged.

6.2 Adversarial Scenario Generation and Parameter Control

6.2.1 Motivation: Why Random Sampling Is Not Enough

Monte Carlo simulation is effective when the analyst has correctly specified the distribution of adverse conditions. In practice, the most damaging failures arise from mis-specification: interactions between parameters, hidden sensitivities, and rare combinations that are under-sampled by naïve stochastic draws. This problem is amplified in illiquid, leveraged assets where small deviations can push the system across terminal thresholds.

To address this, SALM incorporates an adversarial parameter control layer that systematically searches for failure-inducing configurations of model parameters, rather than relying solely on passive sampling.

6.2.2 Adversarial Agents as Stress Search, Not Forecasting

The adversarial layer does not attempt to predict the future. Instead, it is framed as a stress search mechanism: a bounded optimization over parameter space intended to reveal plausible failure modes and sensitivity cliffs.

Each adversarial agent proposes parameter perturbations consistent with:

- the scenario context (asset type, leverage, market regime)
- configured bounds and constraints
- an explicit objective function tied to survivability and loss

This yields a structured set of "hard cases" for SALM evaluation.

6.2.3 Agent Committee Structure and Minimum-Success Gate

In implementation, the adversarial layer uses a multi-agent committee pattern: multiple expert agents operate in parallel to propose stress parameterizations. A minimum-success gate (quorum requirement) ensures the system only proceeds to synthesis if enough agent outputs are valid and parseable under a structured contract.

This design has two methodological benefits:

- **Reliability under model failure:** if one or more agents fail, the quorum gate prevents low-quality stress suggestions from driving conclusions.
- **Diversity of stress proposals:** multiple agent viewpoints reduce the chance of missing failure modes due to a single agent's bias.

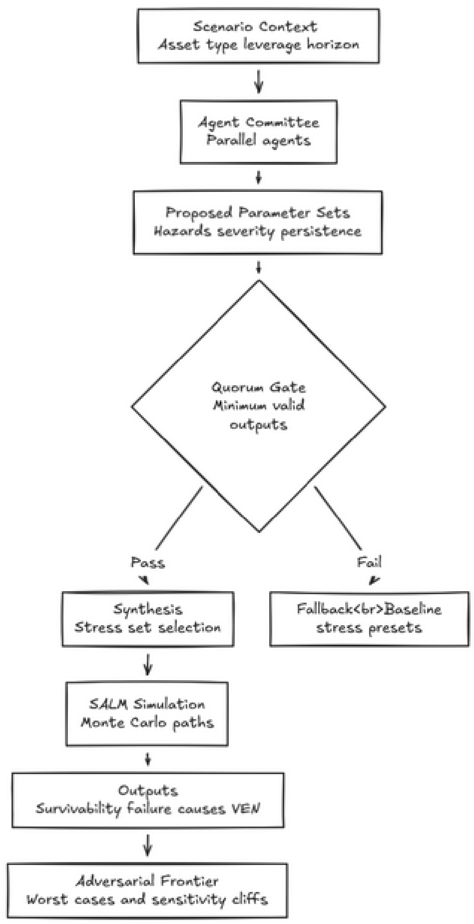

Figure 2: committee agents proposing bounded parameter perturbations; validation + quorum gate; simulation runs; outputs are survivability/VEN distributions, not agent assertions

6.2.4 Evidence-Bound Parameter Contracts and Reproducibility

Agent outputs are constrained to a structured format (e.g., JSON schema) that includes:

- proposed parameter changes
- rationale
- referenced evidence or assumptions (where available)
- confidence / uncertainty tags
- constraint checks (bounds, feasibility)

The evidence-bound contract is essential for academic credibility: it allows auditability and reduces the risk that adversarial parameterizations are arbitrary or hallucinated.

6.2.5 Objective Functions for Adversarial Search

Agents may optimize for one or more of the following SALM-aligned objectives:

- survivability degradation: maximize probability of terminal failure within horizon
- time-to-failure minimization: minimize median survival duration
- capital exhaustion intensity: maximize reserve drawdown rate
- VEN underperformance: maximize probability and magnitude of VEN < 0
- failure mode targeting: identify parameterizations that trigger specific failure pathways (e.g., refi failure vs. liquidity exhaustion)

This turns "stress testing" into a goal-directed discovery process that is directly interpretable in SALM terms.

7 Loss Magnitude: A Multi-Dimensional Concept

Loss magnitude in SALM includes direct, indirect, and terminal components. Adversarial agents may target magnitude pathways by proposing parameterizations that increase (i) severity of cashflow impairment, (ii) duration of impairment (persistence), or (iii) structural loss such as refinancing infeasibility. Importantly, these proposals remain bounded by feasibility constraints and are evaluated through survivability outcomes rather than accepted as true ex ante.

SALM decomposes loss magnitude into three layers:

7.1 Direct Loss

Immediate financial impact:

- cash deficits
- reserve drawdowns
- equity erosion

Each Monte Carlo path tracks these values explicitly, allowing cumulative loss magnitude to be measured over time rather than inferred from terminal outcomes.

7.2 Indirect Loss

Indirect losses arise not from immediate cash impact, but from structural degradation of the AEU's future options.

Examples include:

- reduced refinancing probability following NOI decline
- higher cost of capital due to deteriorating cash flow stability
- loss of growth optionality due to capital depletion

While these losses are not always represented as explicit line items in the current simulation, they are implicitly reflected in survivability outcomes. Paths that experience prolonged stress are more likely to fail at refinancing checkpoints or to terminate due to insufficient capital.

SALM treats these indirect losses as first-class components of loss magnitude, even when they are operationalized through downstream failure conditions.

7.3 Terminal Loss

Irreversible outcomes:

- forced liquidation
- permanent capital loss
- adverse tax consequences

The simulation tracks these effects implicitly through capital state transitions and terminal conditions.

By distinguishing terminal loss from temporary impairment, SALM avoids conflating recoverable stress with true failure. This distinction is critical for understanding which loss trajectories are survivable and which are not.

8 Survivability Modeling as Core Risk Lens

8.1 Survivability as a Time-Indexed Probability

Survivability in SALM is defined as the probability that an AEU remains operational—i.e., avoids terminal loss—over a given time horizon.

Rather than asking whether failure occurs at all, survivability asks:

- when does failure occur?
- how long does the system resist stress?
- which loss processes dominate failure outcomes?

In DealStrike, survivability outputs include probability of failure, median time to failure, and the most common failure causes across Monte Carlo paths.

Formally, survivability is expressed as:

- survival probability at time t
- median survival duration
- dominant failure pathways
- minimum capital required to avoid failure

8.2 Survivability vs Risk

Scalar risk scores compress diverse failure pathways into a single number, obscuring how failure occurs. Survivability curves, by contrast, preserve temporal structure and expose fragility that may not be apparent from aggregate metrics.

> Risk describes what can go wrong.
> Survivability describes how long the system resists failure.

Two AEUs with similar loss frequencies may differ dramatically in survivability due to differences in leverage, reserves, and recovery dynamics.

DealStrike's survivability output captures:

- probability of failure
- time-to-failure distribution
- primary vs secondary failure causes
- what you need to change to survive the scenario

8.3 Adversarial Survivability Stress Search

Survivability is an emergent property of loss processes interacting with finite capital buffers. Adversarial agents are used to identify stress parameterizations that most rapidly degrade survivability—either by accelerating capital exhaustion, increasing the probability of terminal thresholds being crossed, or steering failures into particular modes (e.g., liquidity vs. refinancing). The resulting stress cases form a structured frontier of survivability risk, revealing sensitivity cliffs that may be missed by passive Monte Carlo sampling.

8.4 Capital as the Limiting Factor

A central insight of SALM is that capital, not volatility, is the limiting factor for survivability. The simulation explicitly tracks capital depletion over time, allowing survivability to emerge from the interaction between loss processes and finite reserves.

This framing aligns more closely with investor experience, where failure occurs not because returns fluctuate, but because obligations cannot be met for long enough.

9 Opportunity Cost as Loss: VEN

Traditional investment analysis treats opportunity cost as an external comparison rather than an internal risk factor. SALM integrates opportunity cost directly into the loss framework via Value Equivalent Net (VEN).

VEN measures the difference in terminal value between the chosen AEU and a counterfactual alternative evaluated under the same time horizon and capital constraints.

9.1 Value Equivalent Net (VEN)

SALM introduces Value Equivalent Net (VEN) to model opportunity cost as a counterfactual loss distribution.

VEN measures the difference in terminal wealth between:

- the selected AEU
- the best available alternative allocation under identical constraints

9.2 VEN as Risk Signal

VEN reframes opportunity cost from a theoretical concept into an observable loss:

- negative VEN indicates hidden underperformance
- high VEN variance indicates fragile decisions
- VEN integrates liquidity, tax treatment, and volatility

In implementation, DealStrike compares real-estate AEUs against market alternatives such as index investing using tax-aware consistent time horizons and capital assumptions.

VEN is not computed as a single deterministic number. Instead, it emerges as a distribution derived from comparative simulations.

This allows SALM to answer questions such as:

- how often does this decision underperform a passive alternative?
- what is the magnitude of underperformance in adverse scenarios?
- how volatile is the relative advantage of this decision?

Negative VEN is interpreted as opportunity loss, making forgone alternatives visible as a quantifiable downside.

9.3 VEN as a Risk Multiplier

High variance in VEN signals fragility: decisions that outperform alternatives in some scenarios but underperform catastrophically in others. SALM treats such variability as a risk signal, even if expected returns appear favorable.

9.4 Adversarial Opportunity Loss Discovery

Opportunity cost is formalized via VEN as a counterfactual loss distribution. Adversarial agents can be configured to search for parameterizations where the chosen AEU underperforms the alternative allocation—maximizing the probability mass of VEN < 0 or the severity of left-tail VEN outcomes. This reframes opportunity cost as an adversarially testable property of the decision, rather than a purely retrospective comparison.

10 Integrated SALM Risk Representation

10.1 Why SALM Rejects a Single Risk Score

SALM outputs are intentionally multi-view (survivability curves, failure causes, VEN distributions). The adversarial layer complements this by generating a set of stress-discovered parameterizations that define a practical risk envelope around the decision. Decision risk is therefore represented not only by stochastic sampling under assumed distributions, but also by adversarially discovered regimes where the decision exhibits sharp fragility or unfavorable counterfactual performance.

Instead, it evaluates:

- loss event chains
- survivability curves
- capital exhaustion probabilities
- VEN distributions

Collectively, these define decision risk as the probability-weighted set of failure trajectories over time.

The system also uses "Evidence-Bound, Quorum-Validated Inference Using Probabilistic Agents" (patent pending) to draw conclusions based upon the SALM scores to guide the user.

10.2 Decision Risk as a Set of Failure Trajectories

Under SALM, decision risk is defined as the probability-weighted set of ways an AEU can fail or underperform over time.

This framing allows decision-makers to reason about:

- which failures are likely

- which are catastrophic
- which are survivable with additional capital
- which represent irreversible opportunity loss

10.3 Alignment with Real Decision-Making

By preserving temporal structure, conditional loss pathways, and comparative context, SALM aligns risk modeling with how investors actually experience outcomes. It replaces abstract notions of "riskiness" with concrete questions of survivability, capital adequacy, and relative advantage.

11 Threat Model and Validation of Adversarial Agents

In SALM, adversarial agents are treated as untrusted stress proposers, not authoritative predictors. The threat model assumes that agents may:

- propose implausible or internally inconsistent parameterizations,
- exaggerate stress beyond feasible bounds,
- introduce bias or noise through model failure or hallucination,
- or fail outright to produce usable outputs.

Accordingly, SALM is designed such that no agent output is ever applied directly to the simulation without validation, bounding, and consensus filtering. Adversarial agents are constrained to operate within a tightly controlled parameter surface and are evaluated as inputs to a stress discovery process, not as sources of ground truth.

11.1 Threat Assumptions

The adversarial agent layer is evaluated under the following threat assumptions:

- **Unreliable generation:** malformed, incomplete, or incoherent parameter suggestions.
- **Over-stress bias:** extreme but unrealistic parameter values.
- **Context misalignment:** ignoring scenario context (asset type, leverage, horizon).
- **Non-determinism:** outputs vary across runs due to stochastic model behavior.
- **Correlated failure:** multiple agents converge on the same flawed proposal.

The system explicitly does not assume agents are benign, correct, or calibrated.

11.2 Bounding and Constraint Enforcement

To prevent nonsensical or infeasible stress proposals, all agent outputs are subject to hard and soft constraints prior to acceptance.

11.2.1 Hard Bounds

Each tunable parameter (e.g., hazard rates, severity multipliers, persistence durations) is constrained to predefined numeric ranges derived from:

- historical observations where available,
- conservative domain-specific assumptions,
- explicit analyst-defined limits.

Any proposal exceeding hard bounds is rejected automatically.

11.2.2 Structural Constraints

Agents are restricted to modifying only parameters exposed by the SALM loss model, such as:

- loss event frequency modifiers,
- loss persistence durations,
- magnitude scaling factors.

They cannot introduce new loss types, alter simulation logic, or modify termination criteria.

11.2.3 Scenario Consistency Validation

Agent proposals are validated against the scenario context before acceptance. This includes checks such as:

- asset-type compatibility (e.g., STR vs NNN retail),
- leverage and capital structure alignment,
- horizon consistency (e.g., refi assumptions within loan term).

Proposals that violate contextual assumptions are discarded, even if numerically valid.

11.2.4 Evidence-Bound Parameter Contracts

All agent outputs must conform to a structured parameter contract that includes:

- explicit parameter deltas,
- stated rationale for each modification,
- optional evidence or assumptions,
- confidence or uncertainty indicators.

This contract serves two purposes:

- **Machine validation:** ensuring parseability and schema compliance.
- **Human auditability:** allowing analysts to inspect why a stress case exists.

Outputs that fail schema validation or omit required fields are rejected.

11.3 Committee Structure and Quorum Gate

To mitigate single-agent failure modes, SALM employs a multi-agent committee architecture.

Multiple agents independently propose stress parameterizations. A minimum-success quorum is required before proceeding to synthesis. If fewer than the required number of valid proposals are produced:

- the adversarial step is aborted, and
- the system falls back to baseline stress presets.

This prevents low-quality or adversarially malformed outputs from contaminating the simulation.

11.4 Consensus and De-Duplication

Accepted agent proposals are clustered to identify:

- duplicate stress patterns,
- trivially equivalent parameterizations,
- minor variations that do not materially change outcomes.

Only representative stress cases are retained, reducing overfitting to a single pathological configuration and improving interpretability of results.

11.5 Separation of Stress Discovery and Risk Estimation

A critical design safeguard is the strict separation between stress discovery and risk estimation.

- Agents identify fragile regimes.
- Final risk metrics (survivability, failure timing, VEN distributions) are computed only from simulation outputs, not from agent claims.
- Agents cannot assert that a scenario is "high risk"; they can only propose conditions under which the simulation exhibits failure.

11.6 Reproducibility and Audit Controls

To address non-determinism, SALM enforces:

- deterministic simulation seeds for accepted stress cases,
- full logging of agent proposals and validation outcomes,
- replayable simulation runs for any published result.

This ensures that adversarially discovered failure modes are inspectable, repeatable, and falsifiable.

11.7 Threats Explicitly Out of Scope

SALM does not attempt to:

- defend against malicious human manipulation of parameter bounds,
- guarantee coverage of all possible stress configurations,
- or claim predictive accuracy of adversarial proposals.

The goal is structured discovery of plausible failure modes, not exhaustive enumeration or forecasting.

12 Discussion

SALM aligns more closely with how investors experience failure: gradual erosion, compounding stress, and irreversible exits rather than isolated shocks. By elevating survivability and opportunity cost to first-class modeling components, SALM provides a richer, more realistic representation of decision risk.

13 Conclusion

Survivability-Adjusted Loss Modeling extends traditional loss-based risk analysis into the domain of capital-constrained investment decisions. By integrating time-dependent loss processes, survivability dynamics, and opportunity cost modeling, SALM offers a robust framework for understanding not just expected outcomes, but the pathways by which decisions succeed or fail.