Deterministic Human Structural-Health Monitoring Under Bounded-Telemetry Release

A simulated multi-cohort pathway evaluation of retained time, burden reallocation, and evidence-governed safety inference

MarTaize K. Fails, LAKANA Sovereign Systems, Norman, Oklahoma, USA

2026
Public Research Suite v2.0 – Journal Candidate J2

Abstract

Background. Human safety systems in sport and load-bearing work face a paradox. The more a system can observe bodily state, the more it may also expose the person being protected. Current evaluation habits often emphasize external load, internal load, readiness scores, or post-event reporting, while leaving three scientific questions under-tested: whether a system changes the pathway of structural-load accumulation, whether it makes managed burden visible rather than hidden, and whether public review can occur without raw trajectory exposure.

Objective. This study evaluates LAKANA Sovereign Structural Intelligence (SSI) as a deterministic structural-load inference framework. The aim is not to claim clinical validation or field injury prevention. The aim is to test, inside a declared simulator, whether an SSI policy changes retained time, burden movement, terminal-proxy behavior, uncertainty, and public evidence-governance posture relative to a comparator policy.

Methods. A two-arm Monte Carlo simulation compared an SSI policy with a non-LAKANA comparator policy across five sport profiles and thirteen sport-position strata. The design preserves basketball_female and basketball_male as source-labeled strata where present; the headline basketball retained-time profile is therefore not interpreted as a sex-collapsed validation claim. The primary outcome was a restricted-mean-survival-time-style retained-time delta over fixed horizons. Secondary analyses examined burden-delta behavior, terminal-event proxy movement, sensitivity screening, convergence, tail bookkeeping, and evidence-release governance. The analytic stack included RMST retained-time summaries, Aalen-Johansen cumulative-incidence summaries, Fine-Gray subdistribution models, cause-specific Cox models, Markov pathway logic, Morris screening, Sobol/GSA sensitivity, bootstrap convergence, tail/GPD bookkeeping, and the eight-layer LAKANA Sovereign Structural Evidence Stack (SSES), a post-run evidence-governance layer that audits source linkage, claim boundaries, and public-release discipline without generating new simulation outputs. Negative controls and claim boundaries were built into the design: the comparator was not treated as a named commercial system, terminal proxies were not treated as clinical events, burden was not allowed to disappear behind favorable retained-time summaries, and public evidence-release metrics were not treated as privacy proofs.

Results. Retained-time deltas were positive in all five sport profiles at both 90 and 365 simulator days. At 90 days, deltas were basketball +0.247574, football +0.214367, female soccer +0.210225, male soccer +0.201622, and softball +0.134529 simulated career-years. At 365 days, the corresponding deltas were +0.255039, +0.224717, +0.218515, +0.210103, and +0.145092. The 90-to-365-day shift remained small (+0.007465 to +0.010563). Bootstrap convergence auditing showed low relative drift (maximum 0.009864; median 0.002119). The block-level burden-delta precision interval was [953,444.21, 1,205,922.71] across 390 effective block units, with a stable favorable sign. Public evidence-governance metrics showed strong audit posture (public evidence sufficiency 0.7886; lineage completeness 0.9185), while remaining explicitly non-clinical and non-privacy-proof metrics.

Interpretation. The study supports a bounded pre-field claim: in this simulation, SSI produced coherent retained-time, burden, event-history, sensitivity, tail-bookkeeping, and evidence-governance behavior across multiple independent checks. The whole-frame paper preserves the female/male basketball split as a design requirement while leaving female-lane interpretation to the companion manuscript. In the present public results table, basketball is displayed only as the retained-time horizon profile because the public manuscript does not yet present a complete lane-specific basketball_female / basketball_male table spanning retained-time, burden, terminal-proxy, event-history, and sensitivity summaries under one provenance standard. This is a reporting boundary, not a sex-collapsing assumption. The contribution is a research framework for evaluating structural-load safety inference without converting safety evidence into unrestricted surveillance. The findings do not establish clinical validity, field effectiveness, injury-prevention efficacy, return-to-play authority, regulatory readiness, or formal privacy guarantees.

Keywords: structural-load intelligence; restricted mean survival time; Monte Carlo simulation; structural health monitoring; competing risks; athlete safety; burden visibility; evidence governance; anti-surveillance design.

Reader navigation and package dependencies

This public research suite is organized as a four-document stack. This main manuscript contains the whole-frame simulation argument, the retained-time results, the burden-visibility interpretation, the comparator-policy framing, and the public claim boundaries. The companion manuscript contains the female-lane and managed-burden equity interpretation; it should be used when reading source-labeled women's basketball, female soccer, and softball results. The engineering technical supplement contains the state-space mechanics, TSARO/NICOLE governance abstractions, saturation-bounded control formulation, formal public-safe results, complete model-family stack, SSES layer definitions, and reproducibility boundary. The concise evidence summary is a front-page reading guide for reviewers and partners.

The dependency order is: main manuscript for the whole-frame claim, companion manuscript for female-lane interpretation, technical supplement for mathematical and engineering definitions, and the concise summary for rapid orientation. No document converts simulator-internal evidence into clinical validation, field validation, return-to-play authority, regulatory readiness, device validation, or a formal privacy proof.

1. Introduction: the monitoring paradox

Wearable sensors, external-load metrics, internal-load scores, and readiness dashboards have become common in high-performance sport and are increasingly relevant to physically demanding work. These tools can measure motion, strain, fatigue, recovery, and exposure with increasing detail. Yet the central safety problem is not merely measurement. A system may collect more data while still failing to answer whether structural-load accumulation is recognized earlier, whether managed burden is created or reduced, and whether the person being monitored retains protection against institutional misuse.

This creates the monitoring paradox: the same bodily signals that may support safety can also reveal vulnerability. Fatigue, gait disturbance, thermal stress, recovery debt, prior load, and medical-adjacent patterns can be useful for safety inference, but they can also become leverage for coaching, selection, scholarship, employment, insurance, or disciplinary decisions. A credible safety architecture therefore has to be evaluated as both a structural-load system and a custody-bound system.

This paper proposes that athlete and load-bearing safety systems should be evaluated through **structural-load intelligence without surveillance**. The scientific object is not a diagnosis, a medical clearance, or a commercial dashboard score. The object is a pathway: how modeled load, recovery, recognition, restriction, burden, and terminal-proxy behavior evolve under a stated policy. The governance object is equally explicit: public review should be possible through aggregate, claim-bounded, integrity-tracked artifacts rather than unrestricted raw subject trajectories.

The paper reports a whole-frame simulation across five sport profiles and thirteen sport-position strata. Basketball is handled carefully: the retained-time horizon table reports a basketball profile, while the design layer preserves basketball_female and basketball_male as source-labeled strata where the evidence distinguishes them. A companion manuscript separately analyzes female-lane and minority-protective governance implications. This separation prevents pooled results from being overextended into female-specific or demographic claims and prevents the companion paper from claiming global whole-frame evidence.

2. Conceptual contribution

The conceptual leap in this work is to move the evaluation target away from "more monitoring" and toward **burden-visible structural-load pathway evaluation**. In ordinary monitoring, the system may report whether load is high, recovery is low, or an event occurred. In structural-load intelligence, the question is different: what policy changes the trajectory of retained time, managed burden, terminal-proxy movement, and evidence exposure?

The framework is inspired by structural health monitoring, where cumulative exposure, fatigue, degradation signals, and precursor states are evaluated over time rather than after failure alone (Farrar & Worden, 2007). The analogy is limited. A human body is not a bridge, aircraft component, or machine. The useful scientific parallel is that load-bearing systems require pathway-level reasoning: exposure, degradation, constraint, and uncertainty matter before terminal failure.

The public-evidence contribution is also part of the method. If a paper about safety requires unrestricted raw subject trajectories to be credible, then the review process itself may reproduce the extraction problem. This paper instead treats aggregate publication, figure-level traceability, claim boundaries, and checksum-indexed artifact release as part of the scientific design. That posture supports reviewability but does not claim differential privacy, anonymity, de-identification, no re-identification risk, telemetry validation, or device-level security.

2.1 Public-safe mathematical framing

The manuscript uses a deliberately public-safe mathematical statement of the architecture. A simulated subject at time t is represented by a bounded state vector $x_{i,t}$. The simulator can be described at the public level as a bounded state transition,

$$x_{i,t+\Delta t} = F_\theta(x_{i,t}, u_{i,t}, e_t) + \varepsilon_{i,t}, \qquad \|\varepsilon_{i,t}\| \le \sigma,$$

where $u_{i,t}$ is the policy action, e_t is environmental or contextual input, F_θ is a public abstraction over non-public parameters, and σ denotes a bounded uncertainty envelope. The modeled admissible region is written as

$$\mathcal{C} = \{x : g(x) \le 0\},$$

where g is a public-safe constraint functional. When a state approaches or violates the modeled boundary, the SSI policy is represented as a minimum-intervention projection:

$$u^*_{i,t} = \arg\min_{u \in \mathcal{U}} \|u\|^2_W \quad \text{subject to} \quad F_\theta(x_{i,t}, u, e_t) \in \mathcal{C}.$$

The load-pressure term used for interpretation can be stated without disclosing private thresholds or coefficients:

$$\rho_i(t) = f\{L_i(t) - R_i(t),\ T_i(t),\ S_i(t)\},$$

where $L_i(t)$ denotes cumulative load, $R_i(t)$ recovery, $T_i(t)$ thermal stress, and $S_i(t)$ strain accumulation. The recalibration principle is monotone envelope non-expansion:

$$\mathcal{E}_{k+1} \subseteq \mathcal{E}_k.$$

These expressions define the public mathematical object evaluated here. They are not clinical risk equations, tissue-failure thresholds, device specifications, or a disclosure of proprietary coefficients.

3. Study design and falsification strategy

3.1 Simulation design

The study used a deterministic two-arm Monte Carlo design. The SSI arm and comparator arm operated over the same simulated state space, sport profiles, horizons, and outcome definitions. The comparator was a non-LAKANA policy abstraction and was not treated as a reproduction of any named product or institution.

The whole-frame analysis included five retained-time sport profiles: basketball, football, female soccer, male soccer, and softball. The stratum-rich design layer preserved thirteen sport-position strata, including basketball_female and basketball_male where source-labeled. Soccer is displayed as female soccer and male soccer in the headline retained-time table because those sex-labeled profiles are public-packaged at that level. Basketball is displayed as a single retained-time profile in the headline horizon table while basketball_female and basketball_male remain explicit design strata for lane-integrity accounting. Thus, the basketball row should be read as a public reporting level, not as a claim that female and male basketball are physiologically equivalent or analytically interchangeable. The effective precision unit for the burden-delta analysis was trial by sport by stratum, yielding 390 block-level units.

3.1A Comparator-policy families represented in the simulation

The comparator arm is not a named product and is not a caricatured null. It is a family of policy abstractions used to test what SSI changes under identical simulation assumptions. Five comparator-policy classes are relevant for interpretation.

Comparator-policy family	Public abstraction	Why it matters
External-load policy	Responds primarily to motion, exposure, acceleration, or player-load style variables.	Tests whether external demand alone is enough to reproduce SSI pathway behavior.
Internal-recovery policy	Responds primarily to recovery, readiness, HRV-like, wellness-like, or reserve-like variables.	Tests whether internal state alone is enough to govern structural-load pathways.
ACWR-style workload policy	Responds to acute/chronic workload contrast or workload-ratio logic.	Anchors the comparison to a widely discussed training-load paradigm without claiming ACWR field validation here.
Hybrid-SOTA policy	Combines external load and internal recovery into a richer monitoring abstraction.	Prevents SSI from being compared only against a weak single-signal baseline.
Industry-ablation policy	Represents a centrally governed monitoring platform with weaker release-boundary and claim-governance constraints.	Separates structural-load inference from evidence-governance and surveillance-conversion risk.

These comparator classes are simulation-policy abstractions. They do not reproduce any named commercial system, do not assert vendor superiority, and do not claim that current clinical, coaching, or occupational practices are equivalent to the modeled comparator. Their role is to make the contrast more concrete: SSI differs not only by the signals it reads, but by its saturation-bounded projection, burden-visibility requirement, TSARO/NICOLE governance boundary, and SSES post-run evidence discipline.

3.1B Comparator specification and anti-straw-man guard

The comparator arm is intentionally defined as a non-LAKANA policy-abstraction family, but the comparison is not interpreted as superiority over all real monitoring, coaching, clinical, wearable, occupational, or vendor systems. For journal review, the comparator must be read under three constraints. First, the comparator and SSI arms share the same simulator plant, sport/profile definitions, horizons, state classes, and public estimands. Second, comparator families differ by policy rule and evidence-governance boundary, not by access to a different population or a different outcome definition. Third, no comparator family is tuned to represent a specific commercial product or clinical protocol.

Comparator concern	Required symmetry	Current public mitigation	Remaining boundary
Information asymmetry	Arms must operate on the same simulator state classes and horizon definitions.	Comparator and SSI share public state, sport/profile, horizon, and terminal-proxy definitions.	Non-public implementation parameters are not disclosed.
Weak baseline risk	Comparator must not be a trivial null.	External-load, internal-recovery, ACWR-style, hybrid-SOTA, and industry-ablation families are named as distinct policy abstractions.	No named external system is reproduced.
Outcome mismatch	Arms must use the same estimands.	RMST-style retained time, burden, terminal-proxy, sensitivity, and governance summaries are interpreted under one claim boundary.	Field endpoints are not measured.
Governance confounding	Structural-load inference and evidence-governance effects must not be silently merged.	The manuscript distinguishes policy pathway behavior from SSES release-governance metrics.	Formal decomposition of policy-only versus governance-only ablations is future work.

3.1C Public comparator-policy specification register

To reduce the risk that the comparator is read as a straw baseline, the journal-candidate package treats comparator families as explicit public-safe policy objects. The table below defines the information available to each comparator family, the public decision-rule class, and the symmetry guard used in this manuscript. Proprietary coefficients and thresholds remain withheld, but the evidentiary comparison is not allowed to change horizons, endpoint definitions, exposure grammar, or terminal-proxy coding between arms.

Comparator family	Information available	Public policy rule class	Symmetry / anti-straw-man guard
External-load	Motion, exposure, acceleration, workload-like state variables.	Responds to external demand or accumulated exposure.	Same simulator population, horizons, terminal-proxy grammar, and censoring/event definitions as SSI.
Internal-recovery	Recovery, readiness, reserve, fatigue-like internal-state variables.	Responds to modeled recovery or reserve deterioration.	Not compared against a null; it receives an internal-state information channel.
ACWR-style	Acute and chronic workload windows or ratio-like workload contrast.	Responds to acute/chronic imbalance in modeled load history.	Used as a workload-paradigm abstraction, not an endorsement or field validation of ACWR.

Hybrid-SOTA	External load plus internal recovery/readiness variables.	Combines multiple information families under one monitoring abstraction.	Prevents SSI from being compared only to single-signal policies.
Industry-ablation	Monitoring under weaker release-governance and claim-boundary constraints.	Separates pathway monitoring from public evidence discipline.	Not a named-vendor reproduction and not a vendor-superiority claim.

This table does not make the comparator externally validated. It makes the comparator auditable enough for the current simulation claim: the comparison is between declared public policy abstractions under shared simulator assumptions, not between SSI and a real deployed product.

3.2 Primary and secondary outcomes

The primary outcome was a restricted-mean-survival-time-style retained-time delta: the modeled time difference before terminal structural-load proxy behavior over a fixed horizon. For mode m and horizon τ, the public estimand can be written as $\mathrm{RMST}_m(\tau) = \int_0^\tau S_m(t)\,dt$, with $\Delta_{\mathrm{RMST}}(\tau) = \mathrm{RMST}_{SSI}(\tau) - \mathrm{RMST}_{Comparator}(\tau)$. RMST-style reporting was selected because it is expressed in time units and remains interpretable when proportional-hazards assumptions are not the primary claim surface (Royston & Parmar, 2013). Competing-risk and multi-state ideas informed interpretation of terminal pathways, because one terminal or restriction pathway can alter later observation of another pathway (Fine & Gray, 1999; Putter et al., 2007).

Secondary outcomes included recognition and latency proxies, block-level burden-delta behavior, terminal-event proxy movement, sensitivity screening, convergence auditing, tail-risk bookkeeping, and public evidence-governance metrics. These outputs are simulator-internal. They are not clinical outcomes.

3.3 Negative controls and alternative-explanation guards

The study was constructed to make several alternative explanations visible.

First, **the comparator control** prevents the simulation from being interpreted as named-vendor superiority. The comparison is against a documented policy abstraction, not against a commercial product.

Second, **the burden control** prevents retained time from being reported as costless benefit. A favorable retained-time delta must be read together with managed burden and terminal-proxy movement.

Third, **the horizon control** checks whether the primary pattern reverses when the horizon expands from 90 to 365 simulator days.

Fourth, **the convergence control** checks whether the reported mean behavior is unstable under replication.

Fifth, **the tail-control boundary** prevents rare-event bookkeeping from being read as real-world catastrophic-event prediction.

Sixth, **the public-evidence control** prevents governance metrics from being misread as privacy guarantees. Public evidence-release discipline supports reviewability; it does not establish formal privacy.

3.4 Reproducibility boundary

The public package is designed to reproduce and audit the reported public claims, figures, captions, references, and checksums. It is not designed to reproduce proprietary implementation details. Thresholds, scoring weights, trigger logic, raw trace schemas, custody internals, key-management internals, sensor-fusion coefficients, and raw subject trajectories are outside the public release boundary.

3.5 Analysis provenance and specification status

This section is an analysis-provenance statement, not a claim of formal prospective registration. The public package separates four kinds of material: fixed simulator outputs, post-run source reconciliation, public evidence-governance layers, and interpretive or future-validation discussion. This distinction is included to prevent readers from treating a post-run evidence audit as if it were a field trial registry or a prospective clinical protocol.

The key rule is that SSES is post-run and non-generative. It evaluates completed artifacts, tables, figures, manifests, and claim statements. It may admit, exclude, qualify, classify, or bound a public claim, but it does not rerun the simulator, generate new subject trajectories, create new Monte Carlo evidence, or fill missing source tables.

Analysis component	Status in this public package	Role	Claim boundary
Whole-frame retained-time horizon spine	Fixed public reporting output	Primary simulator-internal retained-time comparison across five profiles	Not field survival or clinical outcome evidence
B=390 MEBVB burden-delta precision layer	Source-reconciled post-run precision layer	Tests whether block-level burden-delta sign remains stable under the declared empirical-Bernstein rule	Not real-world burden reduction
Aalen-Johansen, Fine-Gray, cause-specific Cox, Markov/Semi-Markov	Event-history model families represented in the public model stack	Checks pathway and competing-risk coherence inside the simulator	Not clinical incidence, causality, or field prediction
Morris, Sobol/GSA, bootstrap/convergence	Sensitivity and robustness families	Tests numerical stability and parameter leverage under declared simulation assumptions	Not parameter truth or external calibration
Tail/GPD and rare-event side-harness	Tail-bookkeeping and denominator-governance outputs	Prevents rare-event statements from being denominator-free or hidden behind averages	Not real-world catastrophic-event prediction
SSES post-run evidence-governance layers	Source-reconciled governance stack	Audits source linkage, claim boundaries, lineage completeness, sensitivity-to-claim leverage, tail pairing, and release sufficiency	Not a privacy proof or field-validation substitute
Female-stack companion B=180 analysis	Separate companion scope	Interprets women's basketball, female soccer, and softball lane evidence without owning whole-frame claims	Not demographic subgroup validation
Basketball female/male public lane completion	Lane-integrity reporting boundary	Preserves `basketball_female` and `basketball_male` as design-layer strata while avoiding incomplete public sex-comparison tables	Not sex-neutral biological equivalence or sex-comparison proof
Field validation and empirical calibration	Future validation stage	Requires prospective data, adjudicated outcomes, governed access, and independent review	Not established by this public simulation package

3.6 Estimand hierarchy, multiplicity posture, and inference strength

The public package contains many model-family outputs, but they do not carry equal inferential status. The retained-time horizon spine is the primary simulator-internal estimand. Burden-delta precision and retained-time/burden tradeoff surfaces are key secondary evidence because they test whether retained time is being purchased by hidden burden. Event-history, sensitivity, tail, and governance layers are supportive or audit-oriented unless explicitly elevated by a future registered protocol.

This paper does not claim family-wise error-rate or false-discovery-rate control across all reported model families, profiles, lanes, terminal-proxy families, and governance scores. The inferential posture is hierarchical and bounded: primary and key secondary outputs define the main simulator-internal claim; additional model families evaluate coherence, fragility, denominator discipline, and public-release eligibility. Nominal intervals in supportive layers should not be read as a single confirmatory multiplicity-controlled trial.

Tier	Components	Purpose	Inference boundary
Primary	RMST-style retained-time delta at declared horizons	Tests whether the modeled SSI policy changes retained time relative to comparator across the whole-frame profiles.	Simulator-internal retained-time contrast only.
Key secondary	B=390 MEBVB burden-delta interval; RET-BURD retained-time/burden surface; terminal-proxy movement	Tests whether favorable retained time coexists with burden visibility and source-reconciled precision.	Not field burden reduction or clinical effect.

Supportive event-history	Aalen-Johansen, Fine-Gray, cause-specific Cox, Markov, Semi-Markov	Checks pathway and competing-risk coherence under the simulator grammar.	Not clinical incidence or causal hazard evidence.
Supportive robustness	Bootstrap/convergence, Morris, Sobol/GSA, multiverse/specification stability	Tests numerical stability and parameter leverage under declared simulator assumptions.	Not external calibration or parameter truth.
Governance-only	SAFE-N, ELCI, SCLG, TAF, MCAC, SPCS	Audits public evidence release, lineage, denominator pairing, and claim eligibility.	Not privacy proof, safety proof, or regulatory evidence.
Exploratory/future	Field validation, hardware-in-the-loop testing, full tau-grid RMST sensitivity, external comparators	Defines next validation stages.	Not established by this package.

3.7 Claim registry summary

The journal-candidate reading of the paper depends on explicit claim-to-evidence traceability. The table below is a compact registry for the main public claims; a full machine-readable registry should be included in any journal submission package.

Claim ID	Public claim	Evidence source in suite	Allowed interpretation	Prohibited interpretation
C1	SSI shows positive retained-time deltas across five public sport profiles.	Main retained-time horizon table and Figure 2.	Simulator-internal retained-time contrast.	Field survival, career extension, or injury prevention.
C2	Whole-frame burden-delta sign remains favorable under MEBVB.	Main burden interval and SSES table; technical supplement MEBVB section.	Block-level precision governance over completed simulation outputs.	Real-world burden reduction.
C3	Female-stack results preserve women's basketball, female soccer, and softball lanes.	Companion female-stack tables and concise summary lane table.	Source-reconciled female-lane simulation interpretation.	Demographic subgroup validation or sex-comparison superiority.
C4	Event-history models support pathway coherence.	Main event-history table; technical supplement model-family definitions.	Internal pathway and competing-risk consistency.	Clinical incidence or causal protection.
C5	SSES is post-run and non-generative.	Technical supplement SSES proof sketches and theorem boxes.	Evidence-governance admissibility over existing outputs.	New simulation evidence or external truth.
C6	Bounded-telemetry release limits public/role-facing outputs.	Technical supplement bounded-release section.	Custody and release-boundary architecture.	Differential privacy, anonymity, zero knowledge, or device security proof.

3.8 RMST horizon sensitivity posture

The current public package reports a two-horizon RMST-style sensitivity spine at 90 and 365 simulator days. This is not a continuous tau-grid and is not presented as a full horizon-optimization analysis. It is, however, sufficient to test the immediate reviewer concern that the retained-time result depends on a single arbitrary horizon. Across all five public sport profiles, the retained-time delta remains positive at both horizons, and the 90-to-365 shift remains small relative to the 90-day profile delta.

Sport profile	90-day delta	365-day delta	Absolute shift	Relative shift vs. 90-day delta
Basketball retained-time profile	+0.247574	+0.255039	+0.007465	3.02%
Football	+0.214367	+0.224717	+0.010350	4.83%
Female soccer	+0.210225	+0.218515	+0.008290	3.94%
Male soccer	+0.201622	+0.210103	+0.008481	4.21%
Softball	+0.134529	+0.145092	+0.010563	7.85%

The journal-candidate interpretation is therefore bounded: the primary retained-time finding is directionally stable across the two declared public horizons, but the package does not claim that 90 or 365 simulator days is an optimal tau. A future registered field or external-calibration protocol should add a full tau-grid with predeclared horizon-selection rules. In the present manuscript, all non-primary event-history and sensitivity outputs remain supportive or exploratory unless explicitly tied to the primary retained-time spine.

4. Convergent analytic stack

The paper does not rely on a single favorable table. The public interpretation is built from converging model families that ask different questions about the same structural-load pathway. RMST asks whether modeled retained time changes over a fixed horizon. Aalen-Johansen estimates cumulative incidence under competing terminal pathways. Fine-Gray and cause-specific Cox ask whether the subdistribution and cause-specific event processes move in the same general direction. Markov state-transition logic preserves pathway movement rather than collapsing the simulation into one terminal endpoint. Morris and Sobol/GSA ask whether conclusions are fragile to parameter leverage. The Sovereign Structural Evidence Stack (SSES) asks whether the public evidence itself remains auditable, bounded, and non-extractive. SSES is LAKANA's post-run evidence-governance layer: it links reported simulation values to source artifacts and claim boundaries, checks whether public evidence can be reviewed without raw trajectory disclosure, and prevents simulator-internal findings from being converted into clinical, field, regulatory, or privacy-proof claims. It does not rerun the simulator, generate new estimates, impute missing artifacts, or disclose non-public coefficients, thresholds, weights, sensor-fusion details, or trigger logic.

Model family	Public status	Role in inference
RMST retained-time	Completed/source-reconciled	Primary retained-time horizon spine across five profiles.
Aalen-Johansen CIF	Public-packaged model-output summary	Competing-risk cumulative incidence pattern; not clinical event incidence.
Fine-Gray subdistribution	Packaged in final model outputs and women-equity branch	Subdistribution hazard pattern under simulation; not causal prevention.
Cause-specific Cox	Packaged in final model outputs	Cause-specific hazard pattern; not field prediction.
Markov state-transition	Source-reconciled pathway family with transition/state-space summaries	State/pathway interpretation; not real-world transition probabilities.
Semi-Markov dwell/transition extension	Evidence-lineage present/public packaging pending	Methods-status only unless public-safe table is displayed.

Model family	Public status	Role in inference
Morris screening	Completed/source-reconciled	Local leverage screening over parameter knobs.
Sobol/GSA	N=2048 ladder present; public table summarized here by top parameter per metric	Variance-attribution under declared sampling; not causal validation.
Bootstrap/convergence	Completed/source-reconciled	Stability of reported means under resampling.
Tail/GPD and rare-event side-harness	Completed/source-reconciled bookkeeping	Internal tail adjudication; not external rare-event prediction.

The LAKANA-specific evidence-governance stack is included because the scientific object is not only the event model. A safety inference architecture can appear favorable while hiding burden, losing claim provenance, or making evidence release inseparable from raw surveillance. The eight SSES layers are therefore treated as public research instruments for evidence governance, not as branding or as a substitute for field validation.

LAKANA layer	Expansion	Public result / status	Allowed interpretation
MEBVB	Monotone Empirical Bernstein Variance Bound	Whole-frame: B=390, interval [953,444.21, 1,205,922.71]; female stack: B=180, interval [2,280,391.93, 3,111,841.08]	Precision and uncertainty governance; not field burden reduction.
SPCS	Source-path / structural persistence concordance	Whole-frame coverage 1.0 with 0 contradictions under neutral sign rule.	Source-lineage coherence; not directional proof by itself.
RET-BURD	Retained-time / burden reallocation surface	Whole-frame: 390 eligible blocks; female stack: 180 eligible blocks; female stack 162 dominant favorable and 18 managed-burden blocks.	Reads retained time jointly with burden; not simple benefit scoring.
MCAC	Model-family concordance and adjudication cube	19 model-family rows; mean readiness 0.868; 14 source-reconciled and 5 public-packaging-pending.	Prevents treating packaged and pending model families as the same.
SAFE-N	Source-and-aggregate fidelity evidence, normalized	Adjusted public evidence sufficiency score 0.7886.	Release-governance posture; not privacy proof.
ELCI	Evidence-lineage completeness index	Global lineage completeness 0.9185 across audited claims/sections.	Audit completeness; not evidence strength.
SCLG	Sensitivity-to-claim leverage gate	Morris leverage linked to claim-boundary checks where source-linked.	Sensitivity-to-claim stability; not causality.
TAF	Tail adjudication frontier	Tail/GPD side-harness bookkeeping with explicit internal-prior warning.	Tail governance; not real-world catastrophic-event prediction.

4.1 SSES as a portable evidence-governance method

SSES is introduced here through LAKANA SSI, but it is not logically limited to LAKANA. The stack can be read as a general post-run evidence-governance method for simulation studies that produce source artifacts, model outputs, claim registers, sensitivity summaries, tail denominators, figures, tables, and release-boundary decisions. In a non-LAKANA simulation, SSES would take those objects as inputs and return claim eligibility, lineage completeness, sensitivity-to-claim leverage, denominator eligibility, retained-time/burden classification, and release sufficiency.

The present limitation is novelty: SSES is a proposed evidence-governance stack and has not yet been independently benchmarked across external simulation programs. That limitation should not be hidden. It defines the next methods question. The current paper demonstrates SSES on SSI; future work should test whether the same source-linkage, tail-denominator, sensitivity-to-claim, and release-boundary rules improve reviewability in other simulation domains.

As an illustrative example, consider a peer-reviewed study that reports an RMST difference or a competing-risk cumulative-incidence result. SSES would not re-prove the clinical result and would not decide whether the study is true. It would audit whether the public claim is tied to the correct source table, whether the horizon or event definition matches the figure caption, whether sensitivity or competing-risk assumptions are disclosed, whether denominators are present for tail or rare-event statements, and whether the public statement overreaches the reported estimand. Thus SSES proves claim admissibility and reviewability, not biological truth or external validity.

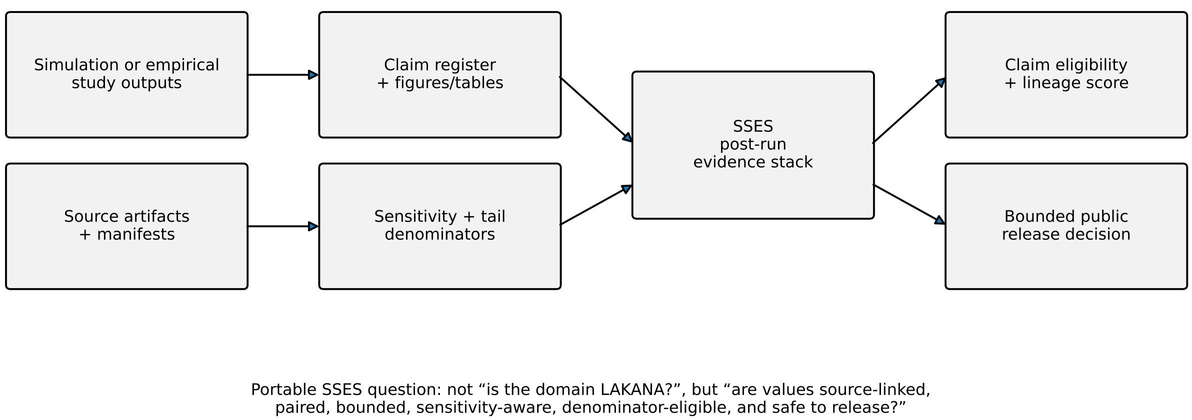

Figure 1: SSES portability beyond LAKANA SSI. The figure shows SSES as a post-run evidence-governance pattern that can sit above any simulation or quantitative study with source artifacts, model outputs, claim registers, sensitivity summaries, denominators, and release rules. It establishes claim eligibility and bounded public release, not external validity or field truth.

The event-history summaries below are included to keep Aalen-Johansen, Fine-Gray, and cause-specific Cox visible in the public paper without turning them into clinical claims. The ACL-family row is displayed because it is the clearest cross-model terminal-proxy check. Degenerative-attrition rows are discussed as pathway reallocation rather than as benefit or harm by themselves.

Profile	Fine-Gray HR	AJ CIF delta	Cox HR
Basketball profile	0.1215 [0.1209, 0.1221]	-0.6866	0.0988 [0.0984, 0.0992]

Profile	Fine-Gray HR	AJ CIF delta	Cox HR
Football	0.1196 [0.1190, 0.1201]	-0.6135	0.1100 [0.1098, 0.1102]
Female soccer	0.1491 [0.1485, 0.1497]	-0.6172	0.1321 [0.1319, 0.1324]
Male soccer	0.0901 [0.0896, 0.0906]	-0.6299	0.0814 [0.0812, 0.0816]
Softball	0.0604 [0.0599, 0.0609]	-0.4460	0.0545 [0.0543, 0.0547]

The Sobol/GSA ladder is reported as sensitivity evidence, not as a causal mechanism. The table shows the top total-order parameter for each public outcome metric at the N=2048 ladder.

Metric	Top parameter	ST
Leakage proxy	freeze hazard threshold	0.9767
Restriction burden delta	freeze hazard threshold	0.4718
RMST delta years	freeze hazard threshold	0.6599
Terminal-event delta proxy	freeze hazard threshold	0.7118

5. Results

Figure 1 summarizes the evaluation object: structural-load inference, public evidence release, and the boundary between reviewable aggregate evidence and non-public raw trajectories.

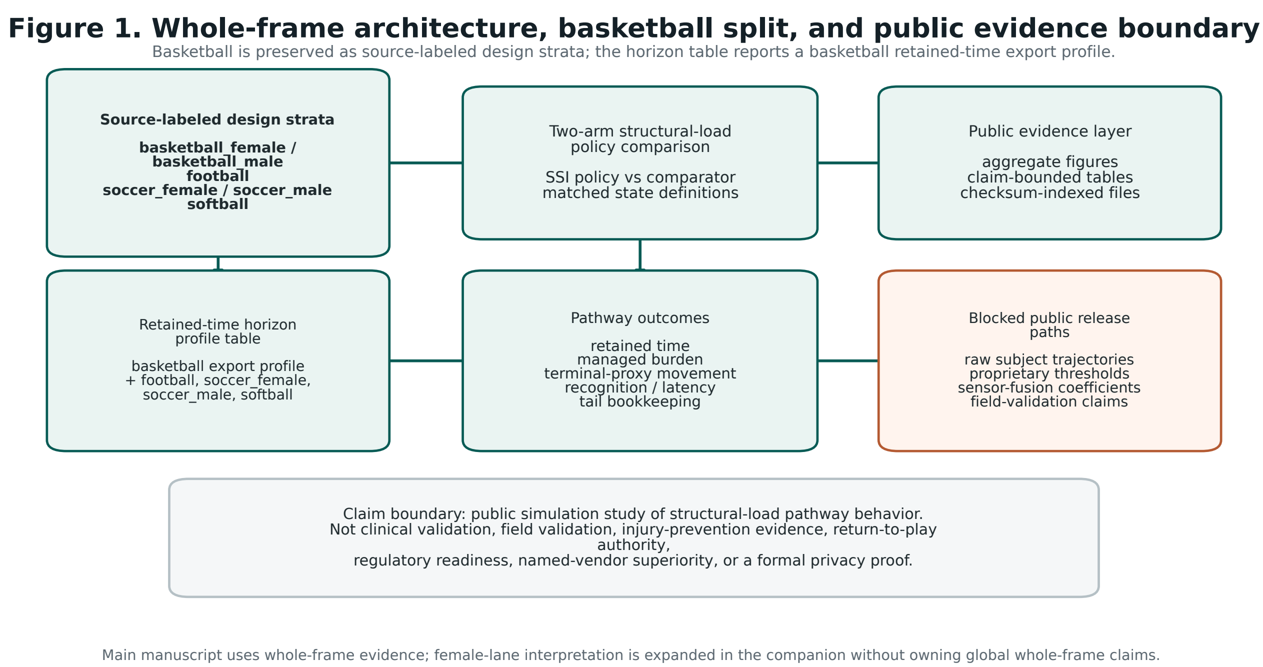

Figure 2: Whole-frame structural-load intelligence, basketball split, and public evidence-release boundary. The figure summarizes the evaluation architecture, including source-labeled basketball__female and basketball__male design strata, deterministic structural-load inference, model-supported interpretation, evidence-governed public release, and blocked raw-data pathways. It is an architecture figure only and does not disclose proprietary thresholds, trigger logic, sensor-fusion coefficients, raw telemetry, or field-validation claims.

5.1 Retained-time horizon stability

Retained-time deltas were positive for all five sport profiles at both horizons. Basketball showed the largest delta, and softball showed the smallest, but no profile reversed direction. The 90-to-365-day shift remained small for every profile, ranging from +0.007465 to +0.010563 simulated career-years.

Sport profile	90-day retained-time delta	365-day retained-time delta	90-to-365 shift
Basketball retained-time profile	+0.247574	+0.255039	+0.007465
Football	+0.214367	+0.224717	+0.010350
Female soccer	+0.210225	+0.218515	+0.008290
Male soccer	+0.201622	+0.210103	+0.008481
Softball	+0.134529	+0.145092	+0.010563

The important result is not that one sport is ranked above another. The important result is that the directional pattern persisted across both horizons under the declared simulation. For basketball, the horizon row is a retained-time export profile; basketball_female and basketball_male remain preserved as source-labeled design strata and must be validated separately in future work.

5.1.1 Basketball lane-integrity status

Basketball is handled as a lane-integrity problem, not as sex-comparison rhetoric. The simulation design preserves basketball_female and basketball_male as separate sex-labeled strata where source labels distinguish them, and the companion manuscript preserves the source-labeled women's basketball lane within the female-stack analysis. The main paper displays the basketball retained-time horizon result only at the public profile level because the present manuscript does not yet display a complete, source-reconciled basketball_female / basketball_male table covering the same public fields used for other lane summaries: retained time, burden, terminal-proxy behavior, event-history behavior, and sensitivity behavior. That choice prevents an incomplete or uneven table from being read as a sex-comparison result. It also prevents the main manuscript from silently collapsing basketball into a sex-neutral biological claim. The correct interpretation is narrower: basketball remains included in the whole-frame simulation and in the female stack where source-labeled, while complete side-by-side basketball sex-lane public reporting is a lane-packaging requirement for the next public evidence release.

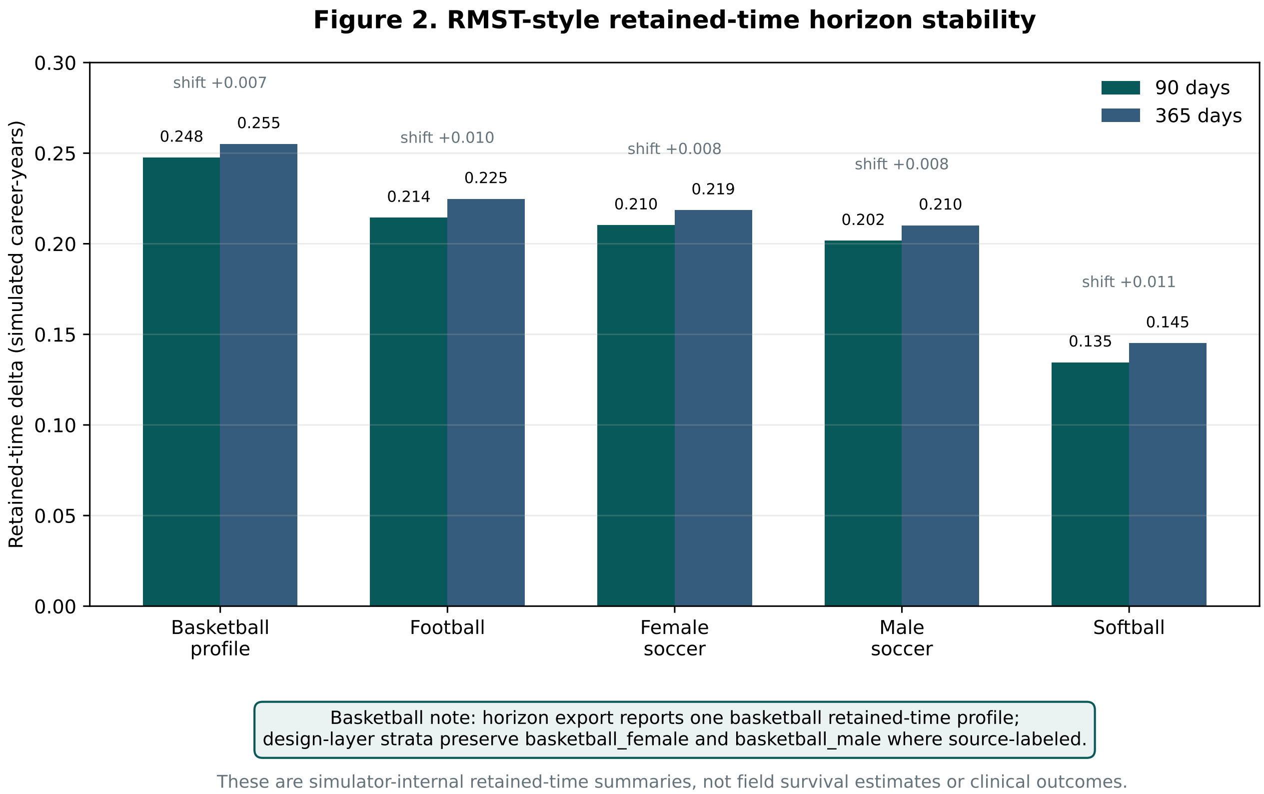

Figure 3: RMST-style retained-time horizon stability across five sport profiles. The figure displays 90-day and 365-day simulator-internal retained-time deltas and the 90-to-365 shift for the basketball retained-time export profile, football, female soccer, male soccer, and softball, while explicitly preserving basketball_female and basketball_male as design-layer strata. These are simulator-internal retained-time summaries, not field survival estimates, clinical outcomes, or real-world career-longevity claims.

5.2 Burden-delta precision

The block-level burden-delta analysis yielded a mean burden delta of 1,079,683.46 and a precision interval of [953,444.21, 1,205,922.71] across 390 block-level units. The interval is computed as $[\bar{\mu} - \varepsilon_B, \bar{\mu} + \varepsilon_B]$, where $\bar{\mu}$ is the block-level mean and ε_B is the variance-aware empirical-Bernstein radius for $B = 390$ effective block units. The lower bound remained positive. This does not mean real-world burden has been reduced. It means the simulator's block-level burden-delta summary retained a stable favorable sign under the declared precision calculation.

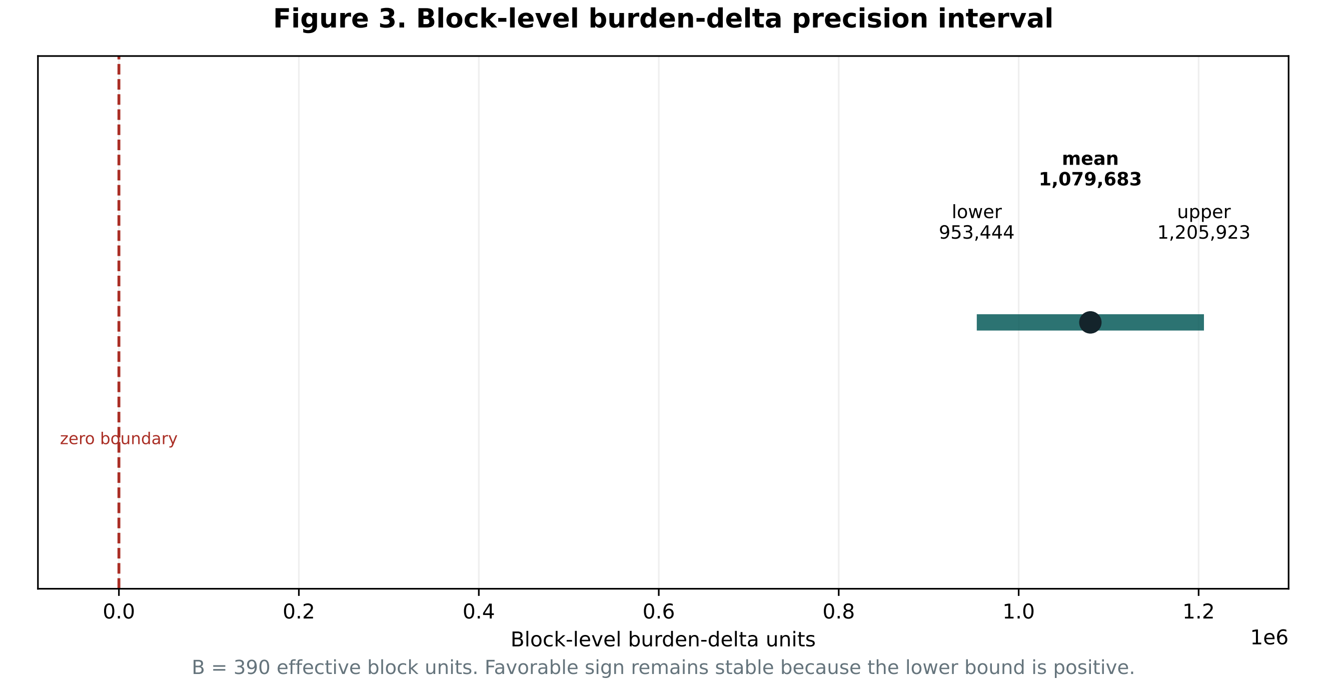

Figure 4: Block-level burden-delta precision interval. The figure displays the variance-aware precision interval around the simulator-internal burden-delta summary across 390 effective block units. The interval supports uncertainty governance inside the simulation; it is not a field confidence interval, clinical validation, or injury-prevention estimate.

5.3 Convergence and sensitivity checks

Bootstrap convergence auditing reached 10,000 primary replicates and 20,000 convergence replicates. Across 80 audited metrics, the maximum relative mean shift was 0.009864 and the median relative shift was 0.002119. This supports numerical stability of the reported public summaries under the chosen replication procedure.

Sensitivity screening identified interpretable movement on retained-time, restriction-burden, terminal-proxy, and leakage-facing outputs. These screening results were used as a robustness check, not as causal proof and not as field parameter validation.

5.4 Public evidence-governance results

The public evidence-governance review returned a public evidence sufficiency score of 0.7886 and a lineage completeness score of 0.9185. These scores summarize whether the public package is organized for review through aggregate evidence, documented limitations, and claim boundaries. They do not validate field effectiveness and do not prove privacy.

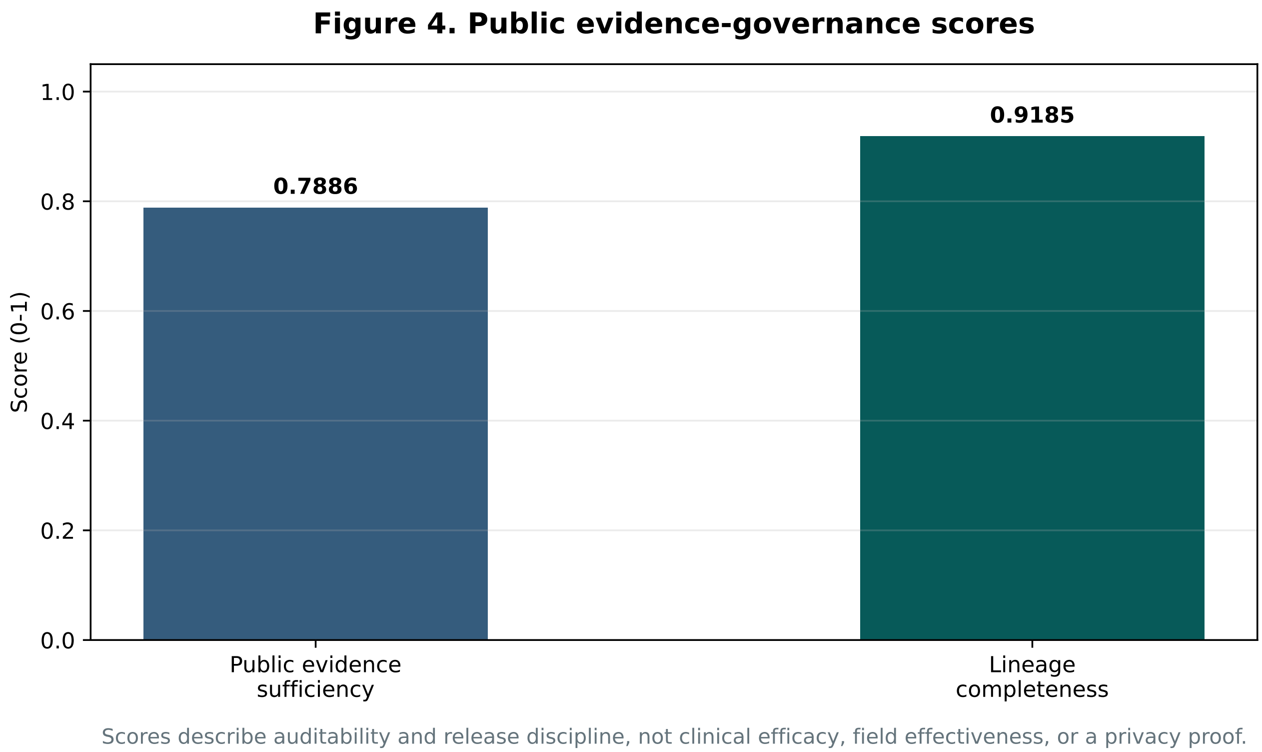

Figure 5: Public evidence-governance scores. The figure summarizes exposure-adjusted public evidence sufficiency and evidence-lineage completeness. These are auditability and release-discipline metrics, not clinical, field, regulatory, privacy-proof, or real-world safety-effectiveness metrics.

5.5 Tail-risk bookkeeping

Tail-risk bookkeeping covered 120 profile-by-metric rows, all with status OK. Among rows with tail-fraction values, the mean numeric tail fraction was approximately 0.058; 23 rows carried numeric generalized-Pareto shape estimates. These outputs are retained as internal tail-adjudication bookkeeping only. They are not empirical estimates of real-world catastrophic-event rates (Coles, 2001).

6. Mechanistic interpretation

The mechanism suggested by the simulation is pathway reallocation. Under the SSI policy, terminal-proxy behavior can be delayed or redirected while more simulated subjects remain in the risk set long enough to accumulate managed-state time. In the public mathematical framing, this is the difference between allowing $\rho_i(t)$ to accumulate toward terminal-proxy behavior and applying a minimum-intervention projection toward $\mathcal{C}$ while keeping managed burden visible. That is why retained time and burden must be read together. A safety policy that only reports fewer terminal proxies can hide the burden it creates; a policy that only reports burden can miss retained-time context.

This joint mechanism explains why the paper does not use a single triumph metric. The simulated policy is evaluated by whether retained time, terminal-proxy movement, burden-delta behavior, horizon stability, convergence, and public evidence governance remain coherent together.

The same logic gives the framework broader relevance. The method is not limited to one sport. The evaluation grammar can apply to any load-bearing human system where exposure, fatigue, recovery, managed restriction, and governance all matter: sport, tactical work, first response, warehouse labor, construction, and other physically demanding domains. That broader applicability is a research-program claim, not proof that the present numerical results transfer to those domains.

7. Negative controls: what the results do not show

The results do not show that SSI prevents injuries. Terminal-event families are structural-load proxies, not clinical events. ACL-like, heat-like, commotio-like, ECAST-like, and soft-event constructs remain simulation variables.

The results do not show that SSI outperforms commercial systems. The comparator was a policy abstraction. Named-vendor superiority would require vendor-specific implementation evidence or an independently defined external comparator protocol.

The results do not show that burden disappears. The analysis explicitly requires burden visibility, and some future validation questions should focus on whether restriction burden is acceptable, contestable, and non-punitive.

The results do not show formal privacy. Public evidence-release governance limits what is disclosed in the paper and figure package. It does not prove differential privacy, de-identification, anonymity, no re-identification risk, telemetry security, or device-level privacy.

The results do not show demographic outcome improvement. The main paper preserves source-labeled basketball_female and basketball_male design strata, but it does not claim female-specific, male-specific, minority-specific, or institution-specific effectiveness. The companion paper addresses female-lane and equity-governance interpretation without owning global whole-frame evidence.

8. Limitations

First, all findings are simulation-bounded. External validity requires prospective studies with independent review, adjudicated outcomes, and prospectively specified endpoints.

Second, the comparator is a modeled abstraction. It is useful for internal contrast but cannot support claims against any named commercial system.

Third, terminal-event families are proxies. They cannot be interpreted as diagnosed injuries or medical outcomes.

Fourth, the modeled structural-load boundary is not an empirically calibrated tissue-failure surface. Calibration against biomechanics, physiology, exposure records, clinical adjudication, and longitudinal outcomes remains required.

Fifth, some event-history and global-sensitivity analyses require additional public-safe packaging before quantitative display. They support the methodological frame but are not used here as headline quantitative figures.

Sixth, public evidence governance is not privacy proof. The technical supplement formalizes bounded-telemetry release and role-scoped custody as a public-safe architecture model, but those equations are not presented as differential privacy, anonymity, de-identification, device-level security validation, or deployed cryptographic proof. Security, device-level custody, formal privacy analysis, and field anti-surveillance validation remain future work.

Seventh, the public package supports review of reported public claims but does not disclose clone-enabling implementation details.

8A. Transportability and external validity

Transportability is not assumed. The present results describe simulator-internal pathway behavior under declared model, policy, and evidence-governance assumptions. Movement from this evidence surface to real athletes or load-bearing workers requires future studies that test how measured signals, exposure patterns, adjudicated outcomes, institutional workflows, and governance rules map onto the simulator's variables and decisions.

SSI is therefore designed to be testable in pilots, not presumed field-valid. A future pilot would need prospectively specified endpoints, an independently described comparator policy, bounded access rules, adverse-use controls, and a plan for contestability, medical adjudication, and participant representation.

Transportability factor	Why it matters	Required future evidence
Sensor fidelity	The simulator uses public-safe state abstractions, while field signals depend on device placement, sampling, calibration, missingness, and latency.	Hardware-in-the-loop and wearable-signal studies with known error properties.
Sport/position exposure rhythm	Load pathways differ by sport, position, season structure, practice/game density, and recovery rhythm.	Prospective sport- and role-specific exposure capture.
Female/male and position-specific lane structure	Lane integrity can be lost if sex-labeled or position-labeled strata are collapsed into generic profiles.	Source-labeled female/male and role-specific datasets with common estimands and provenance.
Medical adjudication	Terminal proxies in the simulator are not diagnoses.	Independently adjudicated outcomes and medical review protocols.
Institutional workflow	Real decisions occur through coaches, clinicians, administrators, employers, and athletes or workers with different incentives.	Pilot workflows that define authority, appeal, override, and non-punitive use.
Privacy/custody governance	Bounded-telemetry release must survive real access pressure and misuse incentives.	Governed-access audits, consent/revocation tests, and adverse-use review.
Comparator policy	Simulation-policy comparators do not automatically represent field practice.	Prospectively specified external-load, internal-recovery, ACWR-style, hybrid, and clinician-reviewed comparator protocols.

9. Ten-year research roadmap

The next decade of work opened by this paper should move in five stages.

1. **External comparator expansion:** evaluate SSI against prospectively specified external-load, internal-load, ACWR, hybrid, and clinician-reviewed policy classes.
2. **Prospective field protocols:** test structural-load endpoints under independent ethics review without granting coaches or employers unrestricted raw-data access.
3. **Biomechanical and physiological calibration:** connect modeled structural-load boundaries to instrumented biomechanics, physiology, recovery, exposure, and adjudicated outcomes.
4. **Basketball lane completion:** publish a dedicated basketball_female / basketball_male lane table only when retained-time, burden, terminal-proxy, event-history, and sensitivity summaries can be displayed for both lanes under the same public provenance standard. The first goal is lane integrity, not sex-comparison rhetoric.
5. **Advanced model packaging:** convert the Aalen-Johansen, Fine-Gray, cause-specific Cox, Markov/Semi-Markov, Morris/Sobol, and SSES results into a stable public table-and-figure set for future preprints and supplements without exposing clone-enabling internals.
6. **Governance validation:** test consent, revocation, appeal, role separation, misuse prevention, and public evidence-release limits under real institutional conditions.
7. **Local-first deployment integration:** evaluate future coupling of SSI with local-first, fail-closed safety operating-system substrates such as LAKANA SOS/CivOS under separate privacy, survivability, hardware, and field-validation protocols. This is an integration roadmap item, not evidence used to upgrade the present SSI simulation into a privacy or deployment claim.
8. **Cross-domain translation:** extend the evaluation grammar to load-bearing occupations such as construction, warehouse work, emergency response, tactical work, and other physically demanding roles.

10. Data, code, and public evidence availability

The public deliverable contains the manuscripts, figures, references, figure registry, and checksum manifests. These files are sufficient to inspect the reported public claims and verify that the figure package is organized and bounded. The public release does not include raw subject trajectories, private implementation logic, proprietary thresholds, scoring weights, trigger logic, custody internals, key-management internals, or sensor-fusion coefficients.

This is not a claim of full independent reimplementation. It is a claim of public claim traceability. Full simulator replication would require governed access to additional materials under purpose limitation, consent, and trade-secret protection.

11. Conclusion

Sovereign Structural Intelligence (SSI) reframes safety evaluation for athletes and other load-bearing populations as a pathway-level problem rather than a single-score prediction task. In this simulation study, the central object is not a diagnosis, a clearance decision, or a commercial readiness score. It is the modeled pathway through retained time, burden accumulation, terminal-proxy movement, recognition, restriction, event-history structure, sensitivity, tail behavior, and public evidence governance.

Across the reported simulation outputs, SSI produced positive retained-time deltas across the five retained-time sport profiles while preserving visibility into the burden tradeoffs that make safety evaluation scientifically honest. The favorable retained-time pattern is not presented as clinical efficacy or field validation. Its value is methodological: it demonstrates a way to evaluate structural-load safety architectures without collapsing burden into hidden attrition, without erasing source-labeled lane structure, and without making unrestricted subject surveillance the price of public evidence.

The convergent analytic stack strengthens that interpretation by asking different questions of the same bounded evidence object. RMST-style summaries describe retained time. Aalen-Johansen, Fine-Gray, cause-specific Cox, Markov, and Semi-Markov layers describe event-history and pathway structure. Morris and Sobol/GSA sensitivity layers test whether interpretation is fragile to influential assumptions. Tail/GPD bookkeeping and rare-event governance prevent low-frequency behavior from disappearing behind averages. The eight SSES layers then audit whether the reported public evidence is source-linked, claim-bounded, and governed rather than merely asserted.

The basketball reporting boundary illustrates the paper's broader discipline. When source-labeled sex-specific strata exist, the manuscript should preserve them as design-layer strata; when a side-by-side public table is not presented at a common provenance standard, the correct response is not to collapse sex-labeled structure into a generic cohort or to force a premature comparison. The current paper therefore preserves basketball_female and basketball_male as distinct design-layer strata while treating sex-comparison claims as outside the displayed public evidence surface.

The contribution is a research grammar for structural-load intelligence without surveillance: define the pathway, preserve lane integrity, report retained time and managed burden together, expose negative controls, audit sensitivity and tail behavior, and release only aggregate, claim-bounded evidence. Future work should extend this grammar through fully packaged sex-lane basketball summaries, externally governed validation cohorts, multi-comparator simulation, measured biomechanics and recovery data, and institutionally auditable protocols that test whether the simulator-internal pathway findings remain meaningful under real measurement, consent, contestability, and oversight.

Appendix A. Journal analysis-plan and provenance dossier

This appendix is a retrospective journal-submission provenance dossier, not formal prospective pre-registration. Its purpose is to make the public inference hierarchy explicit enough for peer review and to prevent model-family breadth from being mistaken for uncontrolled confirmatory testing.

Item	Frozen/public status	Primary artifact or location	Inferential role	Multiplicity handling
Primary estimand	Fixed in public reporting spine	Main retained-time table; Figure 2	Primary simulator-internal retained-time contrast	Interpreted as the primary family
Primary horizons	Declared public horizons: 90 and 365 simulator days	Main Table 5.1 and RMST horizon-sensitivity table	Horizon-stability check	No continuous tau-grid claim
Primary population	Whole-frame public profile set	Basketball export profile, football, female soccer, male soccer, softball	Whole-frame scope	No demographic outcome claim
Key secondary burden	Source-reconciled B=390 MEBVB interval	Main burden table; supplement MEBVB section	Burden-visibility and precision governance	Secondary, not separate confirmatory proof
Female-stack analysis	Separate B=180 companion scope	Companion manuscript and female-stack tables	Lane-integrity interpretation	Not whole-frame ownership or subgroup validation
Event-history families	Source-reconciled supportive models	Event-history model register in supplement	Pathway coherence and competing-risk grammar	Exploratory/supportive; no FWER/FDR claim
Sensitivity families	Source-reconciled robustness models	Supplement sensitivity sections; top-parameter tables	Parameter leverage and numerical stability	Exploratory/supportive
Tail/rare-event models	Denominator bookkeeping layer	Supplement TAF/GPD/POT sections	Prevents denominator-free tail claims	Governance/supportive only
SSES governance metrics	Post-run, non-generative evidence-governance stack	SAFE-N, ELCI, SCLG, TAF, MCAC, SPCS sections	Claim eligibility, lineage, release discipline	Not privacy proof or evidence-strength score

Future external validation	Not performed in this package	Roadmap and limitations	Field, hardware, clinical, and governance validation	Outside present inference set

For journal submission, the machine-readable companion file `JOURNAL_ANALYSIS_PLAN_PROVENANCE_J2.csv` carries the same structure with artifact identifiers and checksum references. The presence of this dossier does not convert the study into a pre-registered trial; it makes the simulation and post-run evidence-governance posture auditable.

Appendix B. Reviewer reproducibility boundary

The public reviewer package supports verification of reported public outputs against manuscripts, tables, manifests, and claim registers. It does not support independent cloning of the simulator because proprietary coefficients, thresholds, trigger rules, sensor-fusion internals, raw trajectories, and protected policy parameters are not public. This is a deliberate reproducibility boundary. The appropriate peer-review test for the current package is source-to-claim verification, not independent production reimplementation.

A governed reviewer archive should contain: (i) manuscripts and supplements, (ii) the claim registry, (iii) artifact inventory, (iv) source-reconciled public tables, (v) SHA256 manifest, (vi) a manifest-verification script, and (vii) a reviewer README describing which claims are machine-verifiable and which require future external validation.

References

The bibliography below is shared by the main manuscript, companion manuscript, and engineering supplement where applicable. LAKANA DOI records are cited as public record links, not as validation evidence.

Aalen, O. O., & Johansen, S. (1978). An empirical transition matrix for non-homogeneous Markov chains based on censored observations. *Scandinavian Journal of Statistics, 5*(3), 141-150. https://www.jstor.org/stable/4615704

Coles, S. (2001). *An Introduction to Statistical Modeling of Extreme Values*. Springer. https://doi.org/10.1007/978-1-4471-3675-0

Cowley, E. S., Olenick, A. A., McNulty, K. L., & Ross, E. Z. (2021). Invisible sportswomen: The sex data gap in sport and exercise science research. *Women in Sport and Physical Activity Journal, 29*(2), 146-151. https://doi.org/10.1123/wspaj.2021-0028

Cummins, C., Orr, R., O'Connor, H., & West, C. (2013). Global positioning systems and microtechnology sensors in team sports: A systematic review. *Sports Medicine, 43*(10), 1025-1042. https://doi.org/10.1007/s40279-013-0069-2

Dietrick, T. A. (2021). Biometric monitoring devices: Modern solutions to protecting athletes' data privacy. *Pittsburgh Journal of Technology Law & Policy, 21*(1). https://doi.org/10.5195/tlp.2021.245

Fails, M. K. (2026a). *LAKANA Sovereign Structural Intelligence white paper*. Zenodo. https://doi.org/10.5281/zenodo.20117305

Fails, M. K. (2026b). *In Silico Comparative Evaluation of a Deterministic Structural-Load Monitoring Framework*. Zenodo. https://doi.org/10.5281/zenodo.19488201

Fails, M. K. (2026c). *LAKANA SOS*. Zenodo. https://doi.org/10.5281/zenodo.19956213

Farrar, C. R., & Worden, K. (2007). An introduction to structural health monitoring. *Philosophical Transactions of the Royal Society A, 365*(1851), 303-315. https://doi.org/10.1098/rsta.2006.1928

Fine, J. P., & Gray, R. J. (1999). A proportional hazards model for the subdistribution of a competing risk. *Journal of the American Statistical Association, 94*(446), 496-509. https://doi.org/10.1080/01621459.1999.10474144

Gabbett, T. J. (2016). The training-injury prevention paradox: Should athletes be training smarter and harder? *British Journal of Sports Medicine, 50*(5), 273-280. https://doi.org/10.1136/bjsports-2015-095788

Hewett, T. E., Myer, G. D., Ford, K. R., Heidt, R. S., Colosimo, A. J., McLean, S. G., van den Bogert, A. J., Paterno, M. V., & Succop, P. (2005a). Biomechanical measures of neuromuscular control and valgus loading of the knee predict anterior cruciate ligament injury risk in female athletes: A prospective study. *The American Journal of Sports Medicine, 33*(4), 492-501. https://doi.org/10.1177/0363546504269591

Hewett, T. E., Zazulak, B. T., Myer, G. D., & Ford, K. R. (2005b). A review of electromyographic activation levels, timing differences, and increased anterior cruciate ligament injury incidence in female athletes. *British Journal of Sports Medicine, 39*(6), 347-350. https://doi.org/10.1136/bjsm.2005.018572

Maurer, A., & Pontil, M. (2009). Empirical Bernstein bounds and sample variance penalization. *COLT 2009*. https://doi.org/10.48550/arXiv.0907.3740

Midoglu, C., Winther, A. K., Boeker, M., Pettersen, S. D., Pedersen, S., Ragab, N., Kupka, T., Hicks, S. A., Randers, M. B., Jain, R., Dagenborg, H. J., Pettersen, S. A., Johansen, D., Riegler, M. A., & Halvorsen, P. (2024). A large-scale multivariate soccer athlete health, performance, and position monitoring dataset. *Scientific Data, 11*, 553. https://doi.org/10.1038/s41597-024-03386-x

Myer, G. D., Ford, K. R., Palumbo, J. P., & Hewett, T. E. (2005). Neuromuscular training improves performance and lower-extremity biomechanics in female athletes. *Journal of Strength and Conditioning Research, 19*(1), 51-60. https://doi.org/10.1519/13643.1

Plews, D. J., Laursen, P. B., Stanley, J., Kilding, A. E., & Buchheit, M. (2013). Training adaptation and heart rate variability in elite endurance athletes: Opening the door to effective monitoring. *International Journal of Sports Physiology and Performance, 8*(6), 688-691. https://doi.org/10.1123/ijspp.8.6.688

Pryor, R. R., Casa, D. J., Vandermark, L. W., Stearns, R. L., Attanasio, S. M., Fontaine, G. J., & Wafer, A. M. (2015). Athletic training services in public secondary schools: A benchmark study. *Journal of Athletic Training, 50*(2), 156-162. https://doi.org/10.4085/1062-6050-50.2.03

Putter, H., Fiocco, M., & Geskus, R. B. (2007). Tutorial in biostatistics: Competing risks and multi-state models. *Statistics in Medicine, 26*(11), 2389-2430. https://doi.org/10.1002/sim.2712

Royston, P., & Parmar, M. K. B. (2013). Restricted mean survival time: An alternative to the hazard ratio for the design and analysis of randomized trials with a time-to-event outcome. *BMC Medical Research Methodology, 13*, 152. https://doi.org/10.1186/1471-2288-13-152

Saltelli, A., Ratto, M., Andres, T., Campolongo, F., Cariboni, J., Gatelli, D., Saisana, M., & Tarantola, S. (2008). *Global Sensitivity Analysis: The Primer*. Wiley. https://doi.org/10.1002/9780470725184

Sandve, G. K., Nekrutenko, A., Taylor, J., & Hovig, E. (2013). Ten simple rules for reproducible computational research. *PLoS Computational Biology, 9*(10), e1003285. https://doi.org/10.1371/journal.pcbi.1003285

Taylor, L. M. D. (2014). The times they are a-changin': Shifting norms and employee privacy in the technological era. *Minnesota Journal of Law, Science & Technology, 15*(2). https://doi.org/10.24926/15529541.3389

Wallerstein, N., & Duran, B. (2010). Community-based participatory research contributions to intervention research: The intersection of science and practice to improve health equity. *American Journal of Public Health, 100*(S1), S40-S46. https://doi.org/10.2105/AJPH.2009.184036

Werner, S. L., Jones, D. G., Guido, J. A., Jr., & Brunet, M. E. (2006). Kinematics and kinetics of elite windmill softball pitching. *The American Journal of Sports Medicine, 34*(4), 597-603. https://doi.org/10.1177/0363546505281796