Technical Supplement to Deterministic Human Structural-Health Monitoring

State-space mechanics, TSARO/NICOLE governance, saturation-bounded control, full model-family stack, SSES proof layers, sensitivity gates, and reproducibility protocols

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

2026
Public Research Suite v2.0 – Journal Candidate J2

Contents

1 Purpose and technical reading order

This technical manuscript is the engineering and mathematical companion to the main public research paper and the female-lane companion paper. Its function is not to repeat the narrative argument of those papers. Its function is to expose the engineering structure that makes the public claims reviewable: the subject-time state model, the deterministic safe-set control abstraction, the physics-pressure variables, the TSARO physics-authority layer, the NICOLE cryptographic governance layer, the post-run Sovereign Structural Evidence Stack (SSES), the event-history model families, the sensitivity and robustness model families, the tail and rare-event model families, and the trade-secret-safe evidence-release boundary.

This technical supplement includes five safeguards that address systems, biomechanics, and privacy review: (i) compact actuator-saturation bounds on the TSARO/SSI projection variable, (ii) dimensional-consistency guards for the public load-accumulation equations, (iii) explicit phenomenological limits on the ACL-family creep proxy, (iv) female-stack parameter-regime language that avoids clinical biological profiling, and (v) a sharper distinction between release-governance scores and device-level cryptographic or statistical privacy proofs. These additions strengthen reviewability without exposing implementation coefficients, thresholds, weights, sensor-fusion details, trigger logic, or raw subject trajectories.

The document is written for reviewers who need to know not only *what* was reported, but *why each model exists*, *how each model is mathematically represented*, *what output it produces*, and *what public claim it can and cannot support*. A model family is included only if it has a defined technical role in the SSI pipeline. Numerical values are retained where they are part of the public reconciled package; proprietary coefficients, thresholds, trigger rules, sensor-fusion weights, raw subject trajectories, and device-level implementation details are not disclosed.

The document should be read in six passes:

1. Sections 2–4 define the public-safe architecture and notation.
2. Sections 5–7 define SSI as a deterministic state-space and physics-pressure system.
3. Sections 8–9 explain how TSARO and NICOLE operate without turning the paper into a hardware implementation disclosure.
4. Sections 10–12 define the model-family stack: event-history, sensitivity, tail, robustness, and lane-integrity models.
5. Sections 13–22 define SSES as eight separate post-run evidence-governance models.
6. Sections 23–29 define reproducibility, claim boundaries, attack surfaces, limitations, and validation roadmap.

2 Suite navigation and dependency map

This technical supplement supports the public research suite rather than replacing the main or companion manuscripts. The main manuscript defines the whole-frame simulation claim and summarizes retained time, burden movement, event-history, sensitivity, tail bookkeeping, and SSES evidence-governance posture. The companion manuscript defines the female-lane interpretation for source-labeled women's basketball, female soccer, and softball. This supplement defines the mathematical objects that make those papers reviewable: state vectors, safe sets, saturation-bounded projection, TSARO physics authority, NICOLE ledger governance, comparator-policy abstractions, event-history models, sensitivity models, tail models, lane-integrity models, and the eight SSES layers. The concise evidence summary is a front-page orientation document.

The documents have different claim surfaces. The main and companion papers interpret results. This supplement defines mechanisms and proof obligations. None of the four documents claims clinical validation, field validation, regulatory clearance, device validation, formal privacy proof, or independent reimplementation from public materials alone.

3 Claim boundary and engineering stance

The public papers make simulator-internal claims. This technical manuscript preserves that boundary. It does not claim clinical validation, field validation, injury-prevention proof, return-to-play authority, medical-device authorization, deployment readiness, production readiness, regulatory clearance, formal verification, telemetry validation, hardware validation, differential privacy, de-identification, anonymity, a no-reidentification-risk guarantee, or a formal zero-knowledge proof.

The technical contribution is narrower and stronger: it defines a deterministic, reproducible, simulator-internal structural-load pathway model and a post-run evidence-governance stack that prevents unsupported claims from escaping the simulation boundary. That is the engineering claim. The public evidence surface is not "trust me" and not "black-box model says so." It is a chain of state-space definitions, source-reconciled estimands, event-history and sensitivity model families, SHA-tracked artifacts, and explicit claim gates.

4 Definitions and notation

Term	Public-safe definition	Claim boundary
SSI	Sovereign Structural Intelligence; the deterministic structural-load inference framework evaluated in simulation.	Simulator-internal safety-inference framework; not a medical device or return-to-play authority.
TSARO	Threat-Adaptive Safety Response Orchestrator; the physics-authority layer that validates sensor/model assertions, enforces safety envelopes, checks environmental plausibility, and computes deterministic boundary acceptance.	Architecture-level doctrine here; no hardware, firmware, or field validation is claimed.
NICOLE	Non-Interactive Cryptographic Oversight and Ledger Enforcement; the cryptographic governance layer for identity binding, evidence custody, immutable ledger recording, revocation, and decision-flow accountability.	Governance-level doctrine here; not a deployed privacy guarantee or formal cryptographic proof.
SSES	Sovereign Structural Evidence Stack; the post-run evidence-governance stack that audits outputs after simulation and determines whether values are source-eligible for public claims.	Does not create new estimates, rerun simulation, or validate real-world outcomes.
State vector	$x_{i,t} \in \mathbb{R}^d$, the public abstraction of subject i at time t.	Model construct; not a real subject record.
Safe set	$\mathcal{C}_t$, a modeled admissible region defined by public-safe constraint functions.	Not a clinical injury threshold or empirically calibrated tissue-failure surface.
Action	$a_{i,t} \in \mathcal{U}$, the modeled policy action.	Does not imply clinical instruction or return-to-play decision.
Aggregate map	$z = \mathcal{A}(X)$, the many-to-one public-release map from internal trajectory collection X to public aggregate artifact z.	Structural non-invertibility orientation only; not differential privacy.

5 System-level architecture

SSI has four engineering layers. The first is the subject-time simulator. It advances modeled athletes or workers through discrete-time state transitions under shared environmental and exposure inputs. The second is the deterministic policy layer. It contrasts a modeled comparator with SSI's minimum-intervention safe-set projection. The third is the TSARO/NICOLE governance boundary. TSARO supplies deterministic physics authority; NICOLE supplies cryptographic custody, revocation, and decision-causality recording. The fourth is SSES. SSES reads completed outputs and determines which values are source-eligible for public claims.

Figure ES-1. SSI system architecture (conceptual schematic).

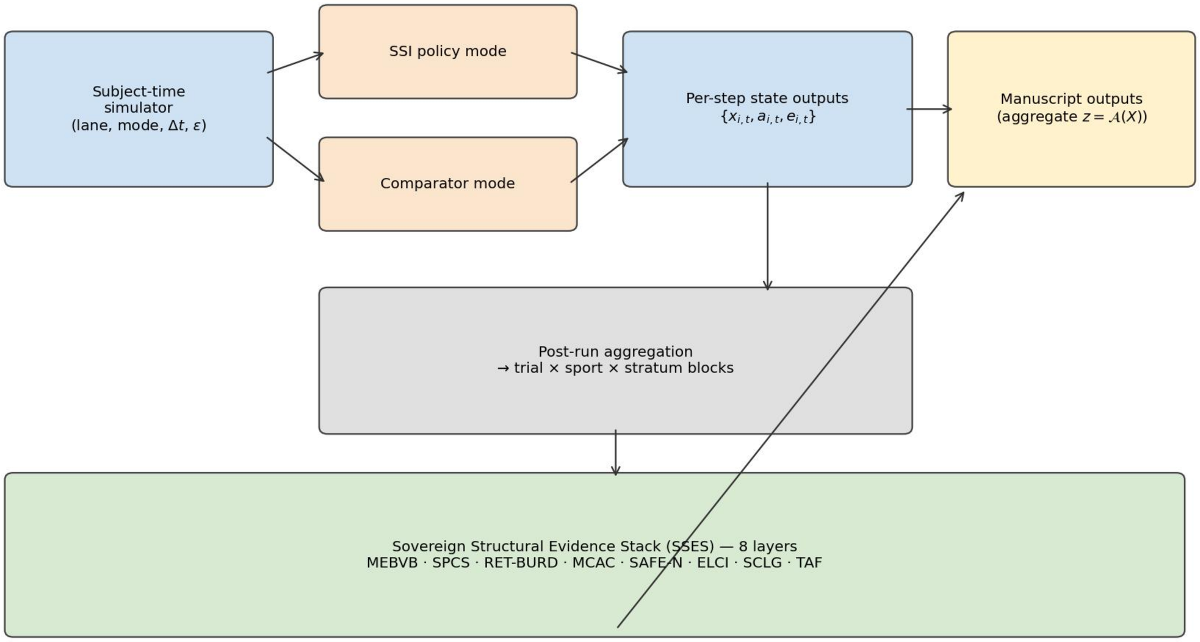

Figure 1: SSI system architecture. The simulator, policy modes, post-run evidence-governance layer, and public manuscript outputs are decoupled. The figure is conceptual and does not disclose coefficients, thresholds, sensor-fusion weights, or trigger logic.

The design separation matters because each layer answers a different reviewer question. The simulator asks: what happens inside the modeled state system? TSARO asks: is a proposed state or load claim admissible under deterministic physics boundaries? NICOLE asks: was access, override, revocation, or decision flow recorded and bounded? SSES asks: is a reported value source-linked, paired correctly, and bounded by a permitted claim surface?

6 Subject-time simulation mechanics

6.1 State vector

Each simulated subject is indexed by i, lane s, mode m, and discrete time t. The public state vector is

$$x_{i,t} = \big(c_{i,t}, q_{i,t}, b_{i,t}, L_{i,t}, \Delta L_{i,t}, F_{i,t}, R_{i,t}, H_{i,t}, G_{i,t}, D_{i,t}, A_{i,t}, M_{i,t}\big),$$

where c is career state, q is escalation or restriction state, b is terminal/active status, L is external load, ΔL is load acceleration, F is fatigue pressure, R is reserve, H is thermal pressure, G is metabolic reserve, D is cumulative structural-load damage memory, A is ACL-family creep proxy, and M is a vector of counters for managed burden, recognition, terminal events, and state occupancy.

The vector is deliberately public at the level of variable classes and private at the level of coefficients, thresholds, weights, and trigger logic. This gives reviewers enough structure to understand the system without exposing the implementation recipe.

6.2 Discrete-time transition

The public transition abstraction is

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

where F_θ is a deterministic transition operator with non-public parameter vector θ, $a_{i,t} \in \mathcal{U}$ is an admissible policy action, $e_{i,t}$ is an environmental input, and $\varepsilon_{i,t}$ is a bounded disturbance term. Monte Carlo replication enters through controlled seeds, lane draws, and scenario instantiation; conditional on the run configuration and seed, the update path is reproducible.

6.3 State machine

The public state machine contains unrestricted, caution/escalation, restricted, chronic-management, and terminal states. Terminal families include ACL-family proxy, degenerative/attritional terminal paths, heat-like terminal paths, commotio-like terminal paths, ECAST-like terminal paths, and administrative horizon exit. The state machine is not a clinical taxonomy. It is a controlled simulation grammar for testing whether a deterministic structural-load policy changes retained time, burden allocation, terminal-proxy distribution, and recognition/latency behavior.

Figure ES-2. Subject-time state machine (conceptual schematic).

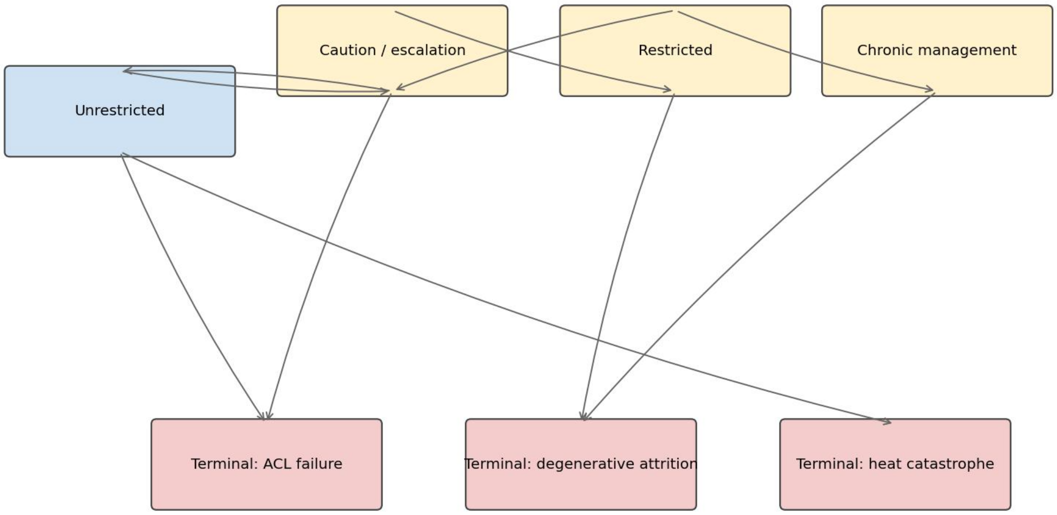

Figure 2: Subject-time state machine. The figure displays state classes and admissible transitions only. Transition coefficients, trigger logic, and thresholds remain non-public.

7 Physics-pressure model

The central physical idea is that structural risk is not only a function of external load. It is a relation between demand, acceleration of demand, internal reserve, thermal pressure, fatigue, and accumulated damage memory. The public biological-boundary pressure index is

$$\rho_{i,t} = \frac{L_{i,t} + \Delta L_{i,t} + H_{i,t} + F_{i,t} + D_{i,t}}{C_{i,t} + \epsilon_C},$$

where $C_{i,t}$ is modeled capacity and $\epsilon_C > 0$ prevents division singularity. The condition $\rho_{i,t} \leq 1$ is a modeled admissibility condition, not a true biological law.

7.1 Load and fatigue update

A public-safe fatigue recurrence can be written as

$$F_{i,t+\Delta t} = \text{clip}\{\eta_F F_{i,t} + \phi_L(L_{i,t}) + \phi_\Delta(\Delta L_{i,t}) - \phi_R(R_{i,t}), 0, F_{\max}\}.$$

The functions ϕ_L, ϕ_Δ, and ϕ_R are public shape classes only. Their coefficients and thresholds are non-public.

Dimensional consistency is enforced by non-dimensionalizing the public pressure variables before aggregation. In a continuous-time reading, if $S_i(t)$ denotes a dimensionless structural-strain pressure proxy, a public form is

$$\frac{dS_i}{dt} = \Phi_L(L_i, \dot{L}_i, H_i, F_i, D_i) - \gamma_R R_i,$$

where Φ_L has dimensions of $[\text{time}]^{-1}$ after non-dimensionalization and $\gamma_R R_i$ carries the same time-normalized units. The dimensional constants that convert sensor-scale quantities into dimensionless pressure terms are implementation-defining and are not disclosed. The public claim is therefore the sign, structure, and boundedness of the update, not a transferable calibration constant.

7.2 Thermal pressure

Thermal pressure is represented as

$$H_{i,t+\Delta t} = \text{clip}\{\eta_H H_{i,t} + \psi_E(e_{i,t}) + \psi_L(L_{i,t}) - \psi_G(G_{i,t}), 0, H_{\max}\}.$$

The term $\psi_E(e_t)$ captures environmental heat/load context; $\psi_L(L_t)$ captures exertional contribution; $\psi_G(G_t)$ captures reserve or recovery offset. This is an engineering pressure variable, not a measured core-temperature claim.

7.3 Damage memory and ACL-family creep proxy

Damage memory is monotone-dominant but recovery-sensitive:

$$D_{i,t+\Delta t} = \text{clip}\{D_{i,t} + \alpha_D \chi_L(L_{i,t}) + \beta_D \chi_F(F_{i,t}) + \gamma_D \chi_H(H_{i,t}) - \delta_D \chi_R(R_{i,t}), 0, D_{\max}\}.$$

ACL-family creep proxy is represented similarly:

$$A_{i,t+\Delta t} = \text{clip}\{A_{i,t} + \alpha_A \chi_{valgus}(x_{i,t}) + \beta_A \chi_{load}(L_{i,t}) - \delta_A \chi_R(R_{i,t}), 0, A_{\max}\}.$$

These equations explain the physics logic: load without reserve, repeated acceleration, thermal stress, and fatigue contribute to structural pressure; recovery can offset but not erase the entire pathway history. The coefficients are not disclosed because they are implementation-defining.

The ACL-family creep proxy is deliberately phenomenological. Biological ligaments and tendons exhibit multi-phase viscoelastic creep, hydration dependence, temperature dependence, collagen-fiber behavior, and remodeling processes that are not represented by the public equation above. The term $A_{i,t}$ is a low-dimensional simulator memory of ACL-family mechanical stress, not a cellular tissue model, not a ligament-failure law, and not a clinically calibrated knee-injury predictor. Its role is to preserve temporal stress memory inside the simulator while keeping microscopic tissue biology outside the public claim surface.

Figure ES-5A. SSI public state-space and physics-pressure map

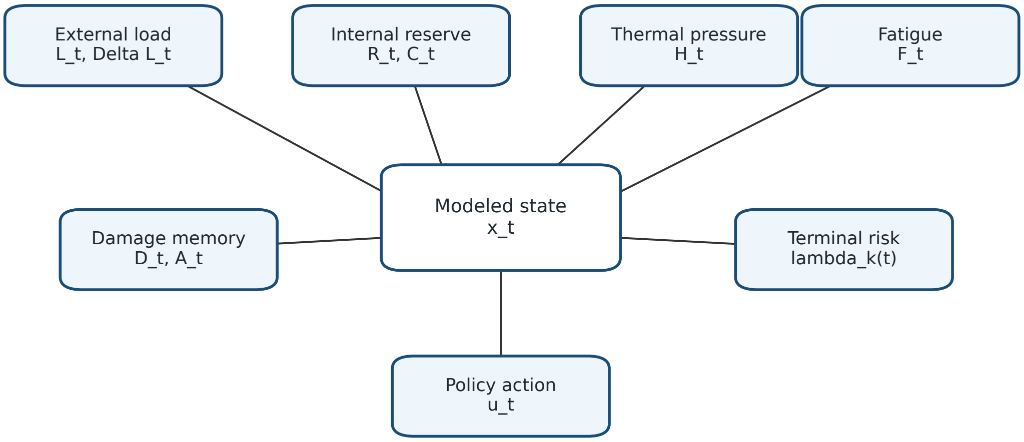

SSI models structural load as a pathway of external demand, internal reserve, thermal pressure, fatigue, and accumulated damage memory.

Figure 3: SSI public state-space and physics-pressure map. The map shows how external load, internal reserve, thermal pressure, fatigue, damage memory, terminal-risk family, and policy action connect inside the public abstraction.

8 Deterministic policy mechanics

8.1 Safe set

The modeled admissible region is

$$\mathcal{C}_t = \{x \in \mathbb{R}^d : h_j(x, e_t) \le 0,\ j = 1, \dots, m\}.$$

The h_j functions represent public-safe constraint classes such as load saturation, thermal pressure, fatigue envelope, damage memory, terminal-proxy proximity, and reserve depletion. The explicit thresholds are non-public.

8.2 Minimum-intervention projection

SSI is represented by a saturation-bounded minimum-intervention projection:

$$u_t^* = \underset{u \in \mathcal{U} \cap \mathcal{U}_{\max}}{\arg\min} \ \|u\|_W^2 \quad \text{subject to} \quad F_\theta(x_t, u, e_t) \in \mathcal{C}_t,$$

where

$$\mathcal{U}_{\max} = \{u \in \mathbb{R}^m : u_{\min} \le u \le u_{\max}\}$$

is a non-empty compact saturation set representing public-safe upper and lower bounds on physiologically executable advisory correction. In engineering terms, $\mathcal{U}_{\max}$ prevents the public control abstraction from requesting impossible joint torque, angular deceleration, neuromuscular response, thermal correction, or load-reduction action. The numerical entries of $u_{\min}$ and $u_{\max}$ are not public because they are implementation-defining and context-dependent.

If $\mathcal{U} \cap \mathcal{U}_{\max}$ contains no action that can return the projected next state to $\mathcal{C}_t$, the public policy does not expand the feasible set. It routes the pathway to the conservative fail-closed state family: restriction, hard stop, escalation, or terminal protective status depending on the modeled state class. This infeasibility rule is essential. It prevents the projection from becoming a mathematically elegant but physically impossible controller.

This explains why SSI can simultaneously improve some terminal-proxy outcomes and increase restriction burden. The policy is not trying to maximize performance or minimize every burden. It is selecting the smallest physically bounded modeled action that keeps the next state inside the admissible set. If the admissible set requires restriction, then restriction is part of the cost structure rather than an accidental side effect.

8.3 Comparator

The comparator uses the same state vector and estimands. It differs in its policy rule. The comparator does not apply SSI's deterministic projection to the safe set. It therefore permits silent accumulation of burden until terminal events occur under the modeled baseline.

Figure ES-14. Comparator vs. SSI policy schematic (conceptual schematic).

Comparator is a modeled baseline policy. It is not a census of every athletic-training, coaching, wearable, clinical, or occupational practice.

Figure 4: Comparator vs. SSI policy schematic. The comparator is a modeled baseline, not a census of all athletic-training, coaching, wearable, clinical, or occupational practices.

9 Comparator-policy abstractions

The comparator side of the simulation is best understood as a set of policy families rather than as a single named external system. This is necessary because the paper does not claim vendor superiority and does not reproduce any proprietary external device. Each comparator family answers a different engineering question.

Comparator family	Abstract policy rule	Engineering question	Boundary
External-load	$a_t = f_L(L_t, \Delta L_t)$	Is external demand alone sufficient?	No named device claim.
Internal-recovery	$a_t = f_R(R_t, F_t, G_t)$	Is recovery/readiness alone sufficient?	Not HRV validation.

Comparator family	Abstract policy rule	Engineering question	Boundary
ACWR-style	$a_t = f_A(Acute_t/Chronic_t)$	Does workload-ratio logic reproduce pathway behavior?	Not ACWR field endorsement.
Hybrid-SOTA	$a_t = f_H(L_t, \Delta L_t, R_t, F_t, H_t)$	Does a richer monitoring baseline close the gap?	Still a policy abstraction.
Industry-ablation	$a_t = f_I$(monitoring signals) with weaker release governance	What changes when evidence governance is removed or weakened?	Not a commercial attack.

These policy families make the comparison more concrete without weakening the claim boundary. SSI differs by combining state-space pressure, saturation-bounded projection, fail-closed TSARO plausibility, NICOLE custody/decision accountability, and SSES post-run evidence governance. The comparator families test parts of that structure but do not constitute clinical practice standards or named systems.

10 TSARO mechanics: physics authority layer

TSARO is the physics-authority layer. It is not a scoring model and not an optimizer of performance. It validates whether a proposed sensor assertion, model assertion, load recommendation, or state transition is plausible and admissible under the deterministic safety envelope.

10.1 Assertion validity

Let y_t be an incoming sensor or model assertion and $\hat{x}_t$ the inferred state. TSARO defines a plausibility predicate

$$\Pi(y_t, \hat{x}_t, e_t) = \mathbb{I}\{d_y(y_t, \hat{x}_t, e_t) \le \tau_y\} \cdot \mathbb{I}\{d_e(e_t, \hat{x}_t) \le \tau_e\} \cdot \mathbb{I}\{\text{integrity}(y_t) = 1\}.$$

The distances d_y, d_e and thresholds τ_y, τ_e are not public. The public rule is simple: a model or sensor assertion must be physically plausible, environmentally plausible, and integrity-valid before it can influence SSI.

10.2 Deterministic acceptance rule

For candidate action a_t and inferred state $\hat{x}_t$, TSARO accepts the transition only if

$$\Pi(y_t, \hat{x}_t, e_t) = 1, \qquad a_t \in \mathcal{U}_{\max}, \qquad F_\theta(\hat{x}_t, a_t, e_t) \in \mathcal{C}_t, \qquad \rho_t \le 1.$$

If any condition fails, the candidate transition is denied or routed to a conservative projection. The saturation condition matters because TSARO is a physics authority layer, not a purely algebraic optimizer: accepted actions must remain inside the public abstraction of feasible human load correction.

10.3 Fail-closed contraction

Let ω_t measure uncertainty from disagreement, missingness, sensor conflict, environmental divergence, or model confidence loss. TSARO maps uncertainty to envelope contraction:

$$\omega_2 \ge \omega_1 \quad \Rightarrow \quad \mathcal{C}_t(\omega_2) \subseteq \mathcal{C}_t(\omega_1).$$

This is the mathematical meaning of fail-closed behavior. Uncertainty cannot expand the safety boundary. It can only preserve or contract it.

Figure ES-2A. TSARO deterministic envelope enforcement

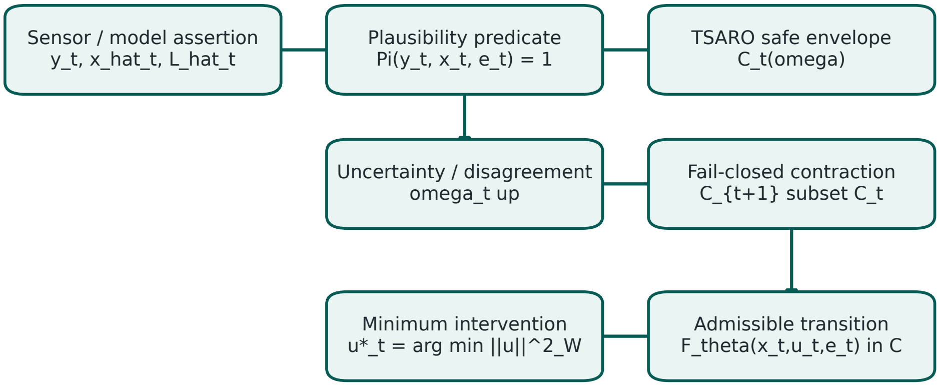

TSARO accepts only physically plausible, safe-envelope-consistent transitions; disagreement narrows the envelope rather than expanding it.

Figure 5: TSARO deterministic envelope enforcement. The diagram shows the chain from model/sensor assertion to plausibility predicate, safe-envelope check, uncertainty contraction, and admissible transition.

11 NICOLE mechanics: cryptographic governance layer

NICOLE governs custody and decision accountability. It does not judge whether the SSI model is accurate. It records what was accessed, by whom, under what scope, with what authorization, and whether the access or override remained valid under the current key epoch and governance state.

11.1 Ledger primitive

A public hash-chain abstraction is

$$\ell_k = H(\ell_{k-1} \parallel r_k \parallel d_k \parallel s_k \parallel \tau_k \parallel c_k \parallel q_k),$$

where r_k is role, d_k is device or credential identifier, s_k is data scope, τ_k is timestamp or monotonic counter context, c_k is content digest, and q_k is key epoch. This captures the ledger logic without exposing implementation.

11.2 Access-validity rule

Access is valid only if

$$\mathrm{AccessValid}_k = \mathbb{I}\{r_k \in \mathcal{R}\}\mathbb{I}\{s_k \subseteq \mathcal{S}(r_k)\}\mathbb{I}\{q_k \in \mathcal{Q}_{active}\}\mathbb{I}\{\mathrm{sig}_k = 1\}\mathbb{I}\{\mathrm{lease}_k = 1\}.$$

This expresses role, scope, key epoch, signature, and lease validity. If any term is zero, access is denied or revoked.

11.3 Decision Causality Ledger

When a safety warning is overridden, NICOLE records a causality event rather than silently discarding the warning:

$$\mathrm{DCL}_k = H(\ell_{k-1} \parallel \mathrm{warning} \parallel \mathrm{override} \parallel \mathrm{role} \parallel \mathrm{ack} \parallel \tau_k).$$

The ledger documents decision flow. It does not make the decision clinically correct and does not replace institutional policy.

11.4 Revocation and key epoch

A revocation event advances the key epoch:

$$q \leftarrow q + 1, \qquad K_q = \mathrm{KDF}(K_{root}, q, \mathrm{context}), \qquad K_{q-1} \to \varnothing.$$

The public claim is key-epoch invalidation at the governance level, not a proof of deployed cryptographic security.

11.5 Transport-survivability boundary and common-cause coupling guard

This SSI supplement does not make a field claim about emergency payload delivery, mesh networking, RF resilience, or CivOS transport survivability. Those are adjacent architecture problems. If a future transport-survivability appendix is added, any independent-channel lower bound must be explicitly qualified. For channel survival probabilities p_i, the idealized independent redundancy term

$$P_{\text{survive}}^{\text{ind}} = 1 - \prod_i (1 - p_i)$$

would be an optimistic bound unless common-cause failure is modeled. A public-safe beta-factor correction can be written as

$$P_{\text{fail}}^{\beta} = \beta P_{\text{ccf}} + (1 - \beta) \prod_i (1 - p_i), \qquad \beta \in [0, 1],$$

where β represents shared RF, spatial, antenna-body-blocking, power, or host-environment coupling. This paper does not report such transport estimates; the equation is included only as a guardrail against misreading NICOLE/TSARO governance as a validated communications survivability result.

Figure ES-3A. NICOLE ledger and revocation mechanics

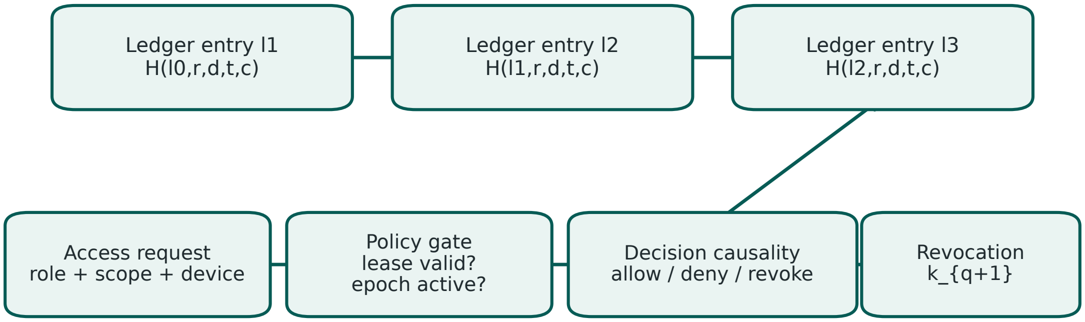

NICOLE governs access, ledger causality, and revocation; it records custody and decision flow without evaluating model accuracy.

Figure 6: NICOLE ledger and revocation mechanics. The diagram shows hash-chained ledger entries, access gating, decision-causality recording, and key-epoch revocation.

12 Deterministic but trade-secret-safe disclosure

The public manuscript must disclose enough structure to make the method reviewable while withholding implementation details that would clone the system. This is achieved by partitioning the parameterization:

$$\theta = (\theta_{pub}, \theta_{sec}),$$

where θ_{pub} contains public equation classes, estimand definitions, claim gates, and reported aggregate values, and θ_{sec} contains coefficients, thresholds, sensor-fusion weights, trigger rules, and implementation internals.

Determinism is preserved because the public computational form is fixed and the reported outputs are source-linked. Trade-secret safety is preserved because the public release map

$$z = \mathcal{A}(X; \theta_{pub})$$

is many-to-one with respect to trajectory collection X and secret parameterization θ_{sec}. Many different internal trajectories and secret parameter choices can produce the same public aggregate z. That fact supports trade-secret-safe reviewability. It is not a privacy theorem.

Across release stages,

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

meaning later public evidence envelopes cannot disclose more raw information than earlier governed envelopes unless explicitly authorized by a new release rule. This monotone non-expansion is a governance invariant, not a formal cryptographic proof.

Figure ES-18A. Deterministic but trade-secret-safe public disclosure

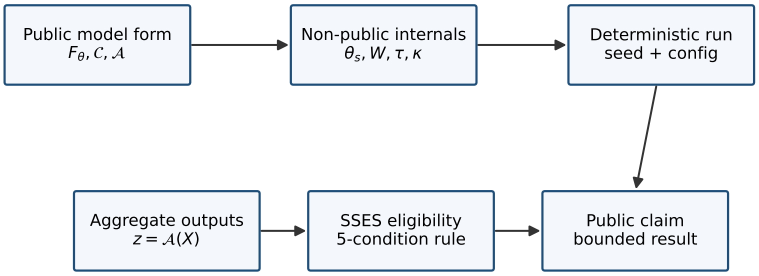

Determinism is preserved by fixed forms and reproducible seeds/configs; cloning is prevented by withholding coefficients, thresholds, fusion weights, and raw trajectories.

Figure 7: Deterministic but trade-secret-safe public disclosure. The figure separates public model form, non-public internals, deterministic runs, aggregate outputs, SSES eligibility, and bounded public claims.

13 Bounded-telemetry release and role-scoped custody model

The SSI research suite uses the phrase bounded-telemetry release in a narrow engineering sense. It does not mean that the simulator or a future deployment has no telemetry. It means that raw subject trajectories are not the ordinary public, coach-facing, vendor-facing, or institutional review object. The public and role-facing outputs are generated through governed maps from a non-public trajectory collection to role-scoped views.

Let

$$X = \{x_{i,t} : i = 1, \ldots, n,\ t = 1, \ldots, T\}$$

denote the non-public subject-time trajectory collection. For role r, SSI exposes a role-bounded view

$$O_r = \mathcal{A}_r(X; \theta_{pub}),$$

where $\mathcal{A}_r$ is a public-safe aggregation, filtering, or release map whose structure is governed by role, purpose, scope, and release envelope. For non-medical roles, the rule is not that raw state is hidden by a label; the rule is that the role output is not an individual raw trajectory:

$$O_{coach} \neq X_i, \qquad O_{vendor} \neq X_i, \qquad O_{public} \neq X_i \quad \text{for subject trajectory } X_i.$$

A coach-facing output may summarize an aggregate safety envelope, an institution-facing output may summarize bounded operational posture, a vendor-facing output may receive pass/conditional/fail assertion feedback, and the public paper may report SSES-adjudicated aggregates. None of those views is equivalent to unrestricted raw biometric, workload, recovery, gait, medical-adjacent, or environmental trajectory custody.

Raw or near-raw clinical telemetry is treated as exceptional access. A public-safe Medical Lease gate can be represented as

$$\begin{aligned}\text{MedicalLease}_{i,r,d,s,t} = {} & \mathbb{I}\{r \in \mathcal{R}_{med}\}\mathbb{I}\{d \in \mathcal{D}_{auth}\}\mathbb{I}\{s \subseteq \mathcal{S}_r\}\mathbb{I}\{t \in [t_0, t_1]\} \\ & \cdot \mathbb{I}\{\text{sig} = 1\}\mathbb{I}\{q \in \mathcal{Q}_{active}\}\mathbb{I}\{\text{revoked} = 0\}.\end{aligned}$$

The terms express authorized clinical role, registered device, scoped data purpose, time-bounded access, valid signature, active key epoch, and absence of revocation. This is a custody model, not a claim that a deployed clinical device or mobile client has been validated.

NICOLE records governed access and decision flow through a hash-chained public abstraction:

$$\ell_k = H(\ell_{k-1} \parallel r_k \parallel d_k \parallel s_k \parallel \tau_k \parallel c_k \parallel q_k \parallel decision_k).$$

The entry binds role, device or credential context, scope, time or monotonic counter, content digest, key epoch, and decision state. The equation states the reviewable governance property; it does not disclose production cryptographic construction, key schedule, hardware boundary, or implementation code.

Proposition: role-bounded non-release under the public model. Assume each non-medical role map $\mathcal{A}_r$ is many-to-one over X, omits raw individual trajectory coordinates, and is constrained by a release envelope $\mathcal{E}_r$. Then the released object $O_r = \mathcal{A}_r(X; \theta_{pub})$ is not an unrestricted raw subject trajectory for any non-medical role $r \notin \mathcal{R}_{med}$.

Proof sketch. By assumption, $\mathcal{A}_r$ maps multiple trajectory collections to the same role output and omits individual raw coordinates. Therefore O_r lacks the one-to-one information required to equal a raw subject trajectory X_i. The release envelope further restricts admissible fields. The conclusion is structural non-release, not privacy in the differential, cryptographic, or adversarial sense.

This bounded-release model establishes five narrow public properties: (i) public review does not require raw trajectory publication; (ii) coach, vendor, public, and general institutional outputs are role-bounded views rather than raw biometric copies; (iii) medical raw access is exceptional, time-bounded, scope-bounded, device-bounded, and logged; (iv) later public release stages cannot silently expand the evidence envelope without a new release rule; and (v) governed access can be tied to a ledger entry and key epoch.

It does not prove differential privacy, anonymity, de-identification, no re-identification risk, side-channel resistance, secure hardware implementation, cryptographic zero knowledge, quantum resistance, legal compliance, or deployed device security. Those are separate validation problems.

14 Complete model-family stack

SSI uses model families because no single statistic can answer all engineering questions. Retained time answers whether modeled pathways persist longer. Event-history models answer where and how pathways terminate or transition. Sensitivity models answer whether claims are fragile to parameter choices. Tail models answer whether rare-event denominators are handled honestly. SSES models answer whether public claims are source-eligible, bounded, and reproducible.

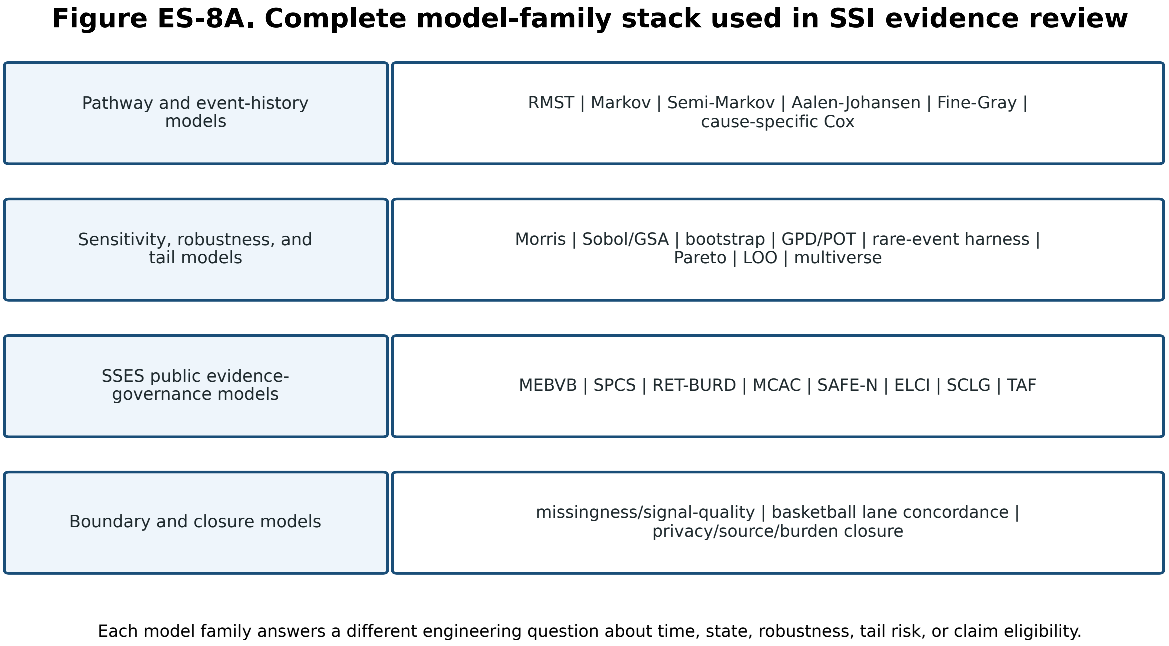

Figure ES-8A. Complete model-family stack used in SSI evidence review

Pathway and event-history models	RMST | Markov | Semi-Markov | Aalen-Johansen | Fine-Gray | cause-specific Cox
Sensitivity, robustness, and tail models	Morris | Sobol/GSA | bootstrap | GPD/POT | rare-event harness | Pareto | LOO | multiverse
SSES public evidence-governance models	MEBVB | SPCS | RET-BURD | MCAC | SAFE-N | ELCI | SCLG | TAF
Boundary and closure models	missingness/signal-quality | basketball lane concordance | privacy/source/burden closure

Each model family answers a different engineering question about time, state, robustness, tail risk, or claim eligibility.

Figure 8: Complete model-family stack used in SSI evidence review. The figure groups the event-history, sensitivity/robustness, SSES, and closure model families.

Model family	What it is	Why SSI uses it	What it establishes
RMST	Time retained before event or horizon.	Converts survival-style comparison into interpretable time.	Simulator-internal retained-time contrast.
Markov	Transition matrix over discrete states.	Shows state occupation and pathway flow.	State-space movement inside the simulator.
Semi-Markov	Transition model with holding-time structure.	Separates state transitions from dwell-time behavior.	Duration-aware pathway dynamics.
Aalen-Johansen	Nonparametric transition/CIF estimator.	Handles multi-state and competing-risk transition probabilities.	Cumulative transition behavior under censoring.
Fine-Gray	Subdistribution hazard model.	Represents target-event cumulative incidence in presence of competing events.	Competing-risk interpretation, not causality.
Cause-specific Cox	Cause-specific process hazard.	Separates process intensity from cumulative incidence.	Cause-specific transition tendency.
Bootstrap / convergence	Resampling and convergence audit.	Checks stability of estimates.	Monte Carlo precision behavior.
Morris	Elementary-effect screening.	Identifies high-leverage parameters efficiently.	Local/global screening signal.
Sobol/GSA	Variance-decomposition sensitivity.	Quantifies contribution of parameters and interactions.	Global uncertainty attribution.
GPD/POT	Extreme-value tail modeling.	Prevents rare-event tail behavior from being ignored.	Tail-exceedance structure under simulator priors.
Rare-event side-harness	Side model for rare terminal families.	Keeps rare outputs visible without making them headline claims.	Rare-event bookkeeping.
Pareto	Trade-off frontier analysis.	Shows retained-time/burden trade spaces.	Non-dominated simulator trade-offs.
Empirical-Bayes shrinkage	Stabilizes noisy lane estimates.	Avoids overreading thin cells.	Regularized lane-level interpretation.
Leave-one-lane-out	Lane-deletion robustness check.	Tests whether one lane drives the finding.	Lane dependence.
Multiverse / specification stability	Alternative specification audit.	Tests fragility to admissible model choices.	Specification robustness.

Model family	What it is	Why SSI uses it	What it establishes
Basketball female/male concordance	Lane-integrity check.	Prevents sex-lane collapse or false comparison.	Proper lane preservation.
Missingness/signal-quality	Missingness and signal audit.	Prevents missing data from masquerading as evidence.	Data-integrity posture.
BLI	Biological-boundary index.	Connects load, fatigue, thermal pressure, damage, and capacity.	Modeled boundary pressure.
Privacy / source / burden closure	Release-boundary closure.	Ensures source and burden claims are consistent with release rules.	Bounded public evidence closure.

15 Specification status of model families

The model-family stack combines fixed simulation outputs, post-run governance models, and interpretive support layers. The table below prevents readers from treating every model family as the same evidentiary object. Some families report primary pathway summaries, some test robustness, some govern release eligibility, and some define future or adjacent integration boundaries.

Model family	Specification status	Primary output	Not used to claim
RMST retained-time	Fixed public retained-time reporting spine	Horizon-specific retained-time delta	Field survival or career longevity
Markov	Source-reconciled pathway family	State-transition pattern	Real-athlete transition probabilities
Semi-Markov	Duration-aware pathway family	Dwell/transition interpretation	Empirical dwell-time validation
Aalen-Johansen	Source-reconciled event-history family	Transition/CIF pattern	Clinical incidence
Fine-Gray	Source-reconciled competing-risk family	Subdistribution pattern	Causal prevention
Cause-specific Cox	Source-reconciled cause-specific process family	Cause-specific hazard pattern	Field prediction
Bootstrap/convergence	Robustness/precision audit	Resampling stability	External validation
Morris	Sensitivity screening	Elementary-effect leverage	Causality
Sobol/GSA	Global-sensitivity family	Variance-attribution ladder	Parameter truth
GPD/POT	Tail-bookkeeping family	Threshold-exceedance structure	Catastrophic-event prediction
Rare-event side-harness	Rare-family visibility layer	Count/denominator bookkeeping	Empirical rare-event incidence
Pareto	Tradeoff interpretation layer	Non-dominated retained-time/burden rows	Universal optimum
Empirical-Bayes shrinkage	Thin-cell stabilization layer	Regularized lane interpretation	Imputation of missing evidence
Leave-one-lane-out	Lane-deletion robustness check	Whole-frame dependence on a lane	Subgroup validation
Multiverse/specification stability	Specification-robustness audit	Direction under admissible alternatives	Mechanistic proof
MEBVB	SSES precision layer	Empirical-Bernstein burden-delta interval	Field burden reduction
SPCS	SSES source-path concordance	Sign/source contradiction status	Outcome superiority
RET-BURD	SSES retained-time/burden surface	Tradeoff quadrant classification	Simple benefit score
MCAC	SSES model-family readiness cube	Model-family execution/readiness taxonomy	Evidence strength
SAFE-N	SSES release-sufficiency metric	Public evidence sufficiency score	Privacy proof
ELCI	SSES lineage-completeness metric	Claim/section completeness score	Evidence truth
SCLG	SSES sensitivity-to-claim gate	Sensitivity-linked claim stability	Causality
TAF	SSES tail adjudication frontier	Tail denominator eligibility	Real-world rare-event rate
Bounded-telemetry release / role-scoped custody	Public-safe custody abstraction	Role-bounded release and ledger model	Differential privacy, anonymity, or device security
SOS/CivOS integration	Future/adjacent integration only; not current SSI evidence	Roadmap boundary	Current SSI privacy, transport, hardware, or survivability validation

16 Event-history model families

16.1 Consolidated event-history model specification register

This register consolidates the public-safe specification for the event-history families used in the manuscript suite. It is not a full disclosure of proprietary simulator coefficients or protected state-transition rules. Its function is to make the estimand, event grammar, censoring/competing-risk treatment, and artifact pathway visible enough for peer review.

Model	Public estimand	Time scale	Event/censoring grammar	Public covariate / stratification abstraction	Public artifact path
RMST	Retained modeled time difference through declared horizon τ.	Simulator day / career-time horizon.	Terminal structural-load proxy or retained-state horizon.	Arm, sport/profile, lane where source-labeled, horizon.	Main retained-time table; Figure 2.
Aalen-Johansen	Transition probability / cumulative incidence in a multi-state grammar.	Simulator event time.	Competing terminal and nonterminal state transitions; censoring internal to simulator grammar.	Arm, profile/lane, event family, state class.	`tables/main_event_history_acl_con` `tables/female_aalen_johansen_summ`
Fine-Gray	Subdistribution pattern for target terminal proxy under competing events.	Simulator event time.	Competing terminal proxy families remain in the risk structure.	Arm, profile/lane, target-event family; no raw biometric covariates public.	`tables/female_finegray_summary.cs` main event-history summary.
Cause-specific Cox	Cause-specific process-intensity pattern for target family.	Simulator event time.	Other event families treated as competing/censoring pathways depending on model family.	Arm, profile/lane, target family, public state class.	Main event-history model stack and technical supplement.
Markov	One-step state-transition pattern.	Discrete simulator step.	Next-state movement among modeled pathway states.	Current state, arm, profile/lane.	Model stack public table; supplement transition definition.
Semi-Markov	Duration-aware transition / dwell interpretation.	Simulator state-holding time.	Holding time before transition to next state.	Prior state, next state, dwell-time family, arm/profile.	Supplement semi-Markov model family and pathway interpretation.

The public package does not report proportional-hazards diagnostics, individual-level residuals, or full covariate matrices because those objects would require non-public simulator internals or raw trajectory disclosure. The journal interpretation is therefore model-family coherence, not field-estimated hazard validity.

16.2 RMST

Restricted mean survival time is

$$\mathrm{RMST}(\tau) = \int_0^\tau S(t)\,dt.$$

SSI uses RMST because hazard ratios are difficult to interpret when hazards are non-proportional or when pathway retention is the practical question. RMST turns the comparison into retained modeled time over a fixed horizon. In this paper RMST-style quantities establish simulator-internal retained-time differences; they do not prove real-world career extension.

16.3 Markov transition model

For discrete state r and next state q,

$$P_{rq}(t) = \Pr(X_{t+\Delta t} = q \mid X_t = r), \qquad \sum_q P_{rq}(t) = 1.$$

Markov summaries show which states are occupied and how frequently adjacent-state transitions occur. This proves that the simulation did not only produce a terminal endpoint; it produced a pathway. It does not prove that real athletes follow the same transition matrix.

16.4 Semi-Markov model

A semi-Markov formulation adds holding time T:

$$Q_{rq}(u) = \Pr(X_{n+1} = q,\ T_{n+1} - T_n \le u \mid X_n = r).$$

SSI uses this because restriction and chronic-management states have duration, not just labels. Semi-Markov structure explains why burden can rise even when terminal proxy movement is favorable: subjects may spend longer in managed states before terminal exit.

16.5 Aalen-Johansen transition/CIF model

For transition intensity matrix $\Lambda(t)$, the Aalen-Johansen transition estimator is the product integral

$$\widehat{P}(s,t) = \prod_{(s,t]} \left(I + d\widehat{\Lambda}(u) \right).$$

SSI uses Aalen-Johansen because multiple transitions and terminal causes compete. The estimator shows cumulative state movement and cumulative incidence in the presence of censoring or competing terminal paths. It establishes event-history coherence; it does not establish clinical incidence.

16.6 Fine-Gray subdistribution model

For target cause k, Fine-Gray models the subdistribution hazard

$$\lambda_k^{FG}(t \mid Z) = \lambda_{k0}^{FG}(t) \exp(\beta_k^\top Z).$$

SSI uses this because one terminal family can prevent observation of another. Fine-Gray supports target-event cumulative-incidence interpretation in the presence of competing events. It does not prove SSI causes lower real-world event rates.

16.7 Cause-specific Cox model

For event family k,

$$\lambda_k(t \mid Z) = \lambda_{k0}(t) \exp(\gamma_k^\top Z).$$

SSI uses cause-specific Cox to distinguish process intensity from cumulative incidence. Fine-Gray asks how the cumulative incidence of a target family behaves when competing events remain in the risk structure; cause-specific Cox asks how the instantaneous target process behaves when other event types are treated as competing/censored pathways.

Figure ES-12A. Event-history interpretation of SSI pathway outputs

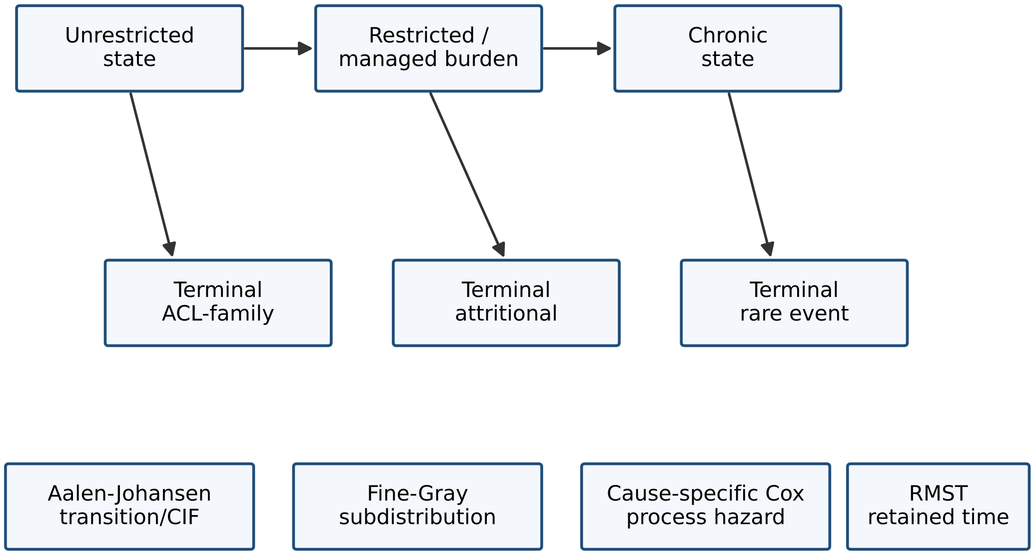

Event-history models separate retained time, transition occupancy, competing terminal causes, and cause-specific process intensity.

Figure 9: Event-history interpretation of SSI pathway outputs. The figure separates state transitions, retained time, cumulative incidence, subdistribution hazard, and cause-specific process intensity.

17 Sensitivity, robustness, and uncertainty model families

17.1 Bootstrap and convergence

Bootstrap/convergence analysis checks whether reported estimates are stable under resampling or increased computation. If $\hat{\theta}_b$ is the estimate from bootstrap replicate b, then uncertainty summaries use

$$\bar{\theta}_B = \frac{1}{B}\sum_{b=1}^{B}\hat{\theta}_b, \qquad \widehat{\mathrm{Var}}(\hat{\theta}) = \frac{1}{B-1}\sum_{b=1}^{B}(\hat{\theta}_b - \bar{\theta}_B)^2.$$

The purpose is not to validate the simulator externally. The purpose is to show whether the simulator's own estimates are numerically stable.

17.2 Morris elementary-effect screening

For parameter j and trajectory point x, the Morris elementary effect is

$$EE_j(x) = \frac{f(x_1, \dots, x_j + \Delta, \dots, x_p) - f(x)}{\Delta}.$$

SSI uses Morris because the model contains many parameters and not every parameter deserves full global sensitivity cost. Morris identifies high-leverage dimensions. It shows leverage; it does not show causality.

17.3 Sobol/global sensitivity analysis

For independent inputs $X_1, \dots, X_p$ and output $Y = f(X)$,

$$S_j = \frac{\mathrm{Var}_{X_j}(\mathbb{E}[Y \mid X_j])}{\mathrm{Var}(Y)}, \qquad S_{Tj} = 1 - \frac{\mathrm{Var}_{X_{\sim j}}(\mathbb{E}[Y \mid X_{\sim j}])}{\mathrm{Var}(Y)}.$$

Sobol/GSA tells reviewers whether output variance is dominated by one fragile assumption or distributed across parameters and interactions. It establishes global variance attribution inside the simulator, not field validity.

17.4 SCLG connection

SCLG converts sensitivity into claim discipline. A parameter may be influential, but a public claim remains eligible only if the claim's sign, precision, and source eligibility remain intact. This is why SCLG links Morris/Sobol rows to claim-boundary metrics such as the MEBVB lower bound.

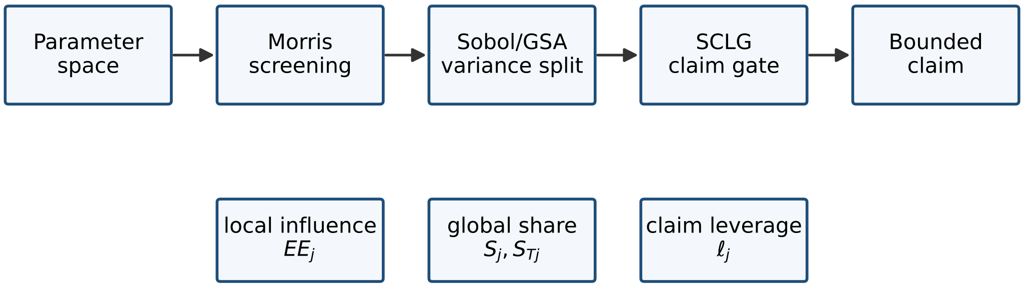

Figure ES-16A. Sensitivity-to-claim pipeline

Sensitivity is not only parameter ranking; SCLG links influence to whether a public claim remains stable under bounded uncertainty.

Figure 10: Sensitivity-to-claim pipeline. Morris and Sobol/GSA identify influential parameters; SCLG asks whether a public claim remains stable after sensitivity information is considered.

18 Tail, rare-event, and tradeoff model families

18.1 GPD/POT tail model

For exceedances $Y = X - u$ over threshold u, the generalized Pareto distribution is

$$\Pr(Y \le y \mid X > u) = 1 - (1 + \xi y/\beta)^{-1/\xi},$$

when $1 + \xi y/\beta > 0$. SSI uses GPD/POT to avoid hiding tail behavior behind means. In this package it is tail-adjudication governance and rare-event bookkeeping, not empirical catastrophic-event prediction.

18.2 Rare-event side-harness

Rare-event families are handled as side-harness outputs so they remain visible but do not dominate the primary estimand hierarchy. The model answers: did a rare terminal family appear under simulator priors, and was it denominator-paired? It does not answer: will this rare event happen in the field?

18.3 Pareto tradeoff model

Let r be retained-time improvement and b be burden change. A row is Pareto-dominated if another row has at least as much retained time and no greater burden with one strict improvement. The Pareto layer identifies non-dominated tradeoff structures:

$$(r_a \geq r_b,\ b_a \leq b_b) \quad \text{and at least one strict inequality.}$$

SSI uses this because safety systems can shift burden rather than eliminate it. Pareto analysis shows trade spaces; it does not declare a universal optimum.

18.4 Empirical-Bayes shrinkage

For a lane estimate $\hat{\theta}_s$ with sampling variance v_s, a shrinkage form is

$$\tilde{\theta}_s = w_s\hat{\theta}_s + (1 - w_s)\mu, \qquad w_s = \frac{\tau^2}{\tau^2 + v_s}.$$

This prevents thin or noisy lane cells from being overread. In SSI it is a stabilization model for interpretation, not a method for inventing missing data.

18.5 Leave-one-lane-out and multiverse stability

Leave-one-lane-out recomputes a quantity without lane s:

$$\Delta_{(-s)} = g(\mathcal{D} \setminus \mathcal{D}_s).$$

Multiverse stability evaluates whether conclusions survive admissible specification alternatives. These models show whether a public claim is robust to lane deletion and specification choice.

19 Lane-integrity and boundary models

19.1 Basketball female/male concordance

Basketball lane concordance is not a sex-comparison claim. It is a lane-preservation test. If basketball has sex-labeled lanes, the public evidence pipeline must not collapse those lanes into a generic basketball label unless the table is explicitly labeled as pooled or aggregate. The concordance model asks whether the lane identity, source label, estimand, and public table are consistent:

$$\text{Concordance} = \mathbb{I}\{\text{lane label}\}\mathbb{I}\{\text{source path}\}\mathbb{I}\{\text{estimand match}\}\mathbb{I}\{\text{claim boundary}\}.$$

The model proves lane integrity at the evidence-governance level. It does not assert biological equivalence or sex-comparison superiority.

19.2 Missingness and signal quality

Let $M_{j,s}$ indicate missingness for feature or output j in lane s. A public summary must not use a row if missingness changes the meaning of the estimand:

$$\text{Eligible}_{j,s} = \mathbb{I}\{M_{j,s} = 0\}\mathbb{I}\{\text{source reconciled}\}\mathbb{I}\{\text{pairing valid}\}.$$

This model prevents absence of data from being treated as evidence.

19.3 Biological-boundary index

The biological-boundary index is the modeled pressure ratio $\rho_{i,t}$ already defined above. Its role is to make the SSI model about structural pressure relative to modeled capacity, not raw external load alone. BLI supports a pathway-mechanics interpretation. It is not a measured biological limit and not a clinical diagnostic threshold.

19.4 Privacy/source/burden closure

The closure model asks whether three claims can coexist without contradiction: source eligibility, burden interpretation, and release boundary. A public claim is closure-eligible only if

$$\text{Closure} = \mathbb{I}\{\text{source eligible}\}\mathbb{I}\{\text{burden paired}\}\mathbb{I}\{\text{release permitted}\}\mathbb{I}\{\text{prohibited claim absent}\} = 1.$$

This is the governance mechanism that keeps the paper from overclaiming.

20 SSES overview: eight post-run evidence-governance models

SSES is not a single score. It is a stack of eight post-run models. Each model answers a narrow question. Together they allow a reviewer to verify whether the public claim surface is complete enough to be reviewed and bounded enough to avoid unsupported claims.

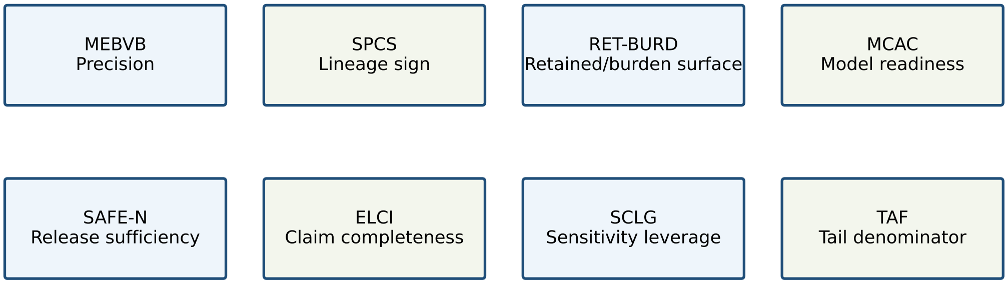

Figure ES-10A. SSES eight-model evidence-governance proof grid

Each layer proves a narrow public-review property; none converts simulation output into clinical, field, regulatory, or privacy-proof evidence.

Figure 11: SSES eight-model evidence-governance proof grid. Each model establishes a narrow public-review property; none converts simulator-internal results into clinical, field, regulatory, or privacy-proof evidence.

SSES layer	What it is	Why it is used	What it establishes
MEBVB	Empirical-Bernstein precision model over blocks.	Quantifies precision around burden-delta sign and magnitude.	Stable simulator-internal precision interval.
SPCS	Source-path sign-concordance model.	Prevents sign claims from being inferred from mere model presence.	Directional lineage coherence and contradiction count.
RET-BURD	Retained-time/burden quadrant surface.	Shows whether retained time coexists with lower or higher burden.	Tradeoff classification.
MCAC	Model-family concordance cube.	Prevents computed models from being mislabeled as absent.	Execution/readiness taxonomy.
SAFE-N	Evidence-release sufficiency score.	Scores whether public release is adequate and bounded.	Aggregate release sufficiency under exposure adjustment.
ELCI	Evidence-lineage completeness index.	Scores per-claim audit completeness.	Whether each claim has source, limitation, and boundary coverage.
SCLG	Sensitivity-to-claim leverage gate.	Links parameter sensitivity to claim eligibility.	Whether sensitive parameters threaten claim stability.
TAF	Tail denominator adjudication.	Ensures tail rows have paired numerator and denominator.	Tail-governance eligibility.

21 SSES portability beyond LAKANA SSI

SSES is introduced through LAKANA SSI because SSI provides the first full use case: model families, source artifacts, public claims, sensitivity summaries, tail denominators, governance scores, and release boundaries all exist in one simulation package. The method is not logically limited to LAKANA. SSES is a post-run evidence-governance pattern that can be applied to any simulation or quantitative study with five ingredients:

1. source artifacts or model-output tables;
2. a claim register or list of public statements;
3. figure/table manifests and captions;
4. sensitivity, tail, or denominator information where relevant; and
5. release-boundary rules defining what may be made public.

Under that abstraction, SSES proves a narrow property: whether a public claim is source-linked, denominator-eligible, sensitivity-aware, burden/benefit coherent, lineage-complete, and release-bounded under the rules of the study. It does not prove that the underlying domain model is true. It does not prove external validity. It does not replace peer review. It makes the evidence surface auditable.

21.1 Illustrative peer-reviewed-result example

Consider a peer-reviewed survival or competing-risk paper that reports an RMST difference, a cumulative-incidence curve, or a subdistribution hazard ratio. SSES could be applied after publication without changing the science. It would ask: Is the public claim tied to the correct estimand? Does the figure horizon match the table horizon? Are competing events and censoring rules stated? Are denominators present for rare-event or tail statements? Are sensitivity checks linked to claims rather than merely listed? Does the abstract overstate what the model estimates? If the answer is yes, SSES would mark the claim as eligible and lineage-complete. If the answer is no, SSES would flag the claim as incomplete, sensitivity-unclear, denominator-missing, or overextended.

Thus SSES would not prove that a drug works, that a device is safe, or that a sports intervention prevents injury. It would prove that the public claim is admissible under source, estimand, sensitivity, denominator, and release-governance rules. This is why SSES can be useful beyond LAKANA: it is an audit grammar for public evidence surfaces.

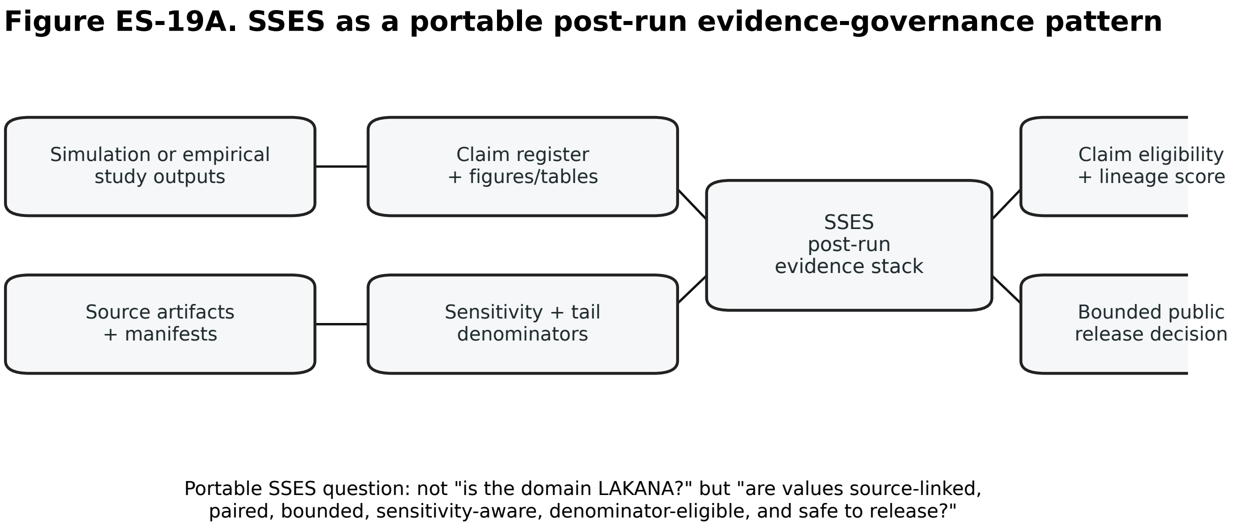
Figure ES-19A. SSES as a portable post-run evidence-governance pattern

Portable SSES question: not "is the domain LAKANA?" but "are values source-linked, paired, bounded, sensitivity-aware, denominator-eligible, and safe to release?"

Figure 12: SSES portability as a general post-run evidence-governance pattern. The method can sit above LAKANA SSI or any other simulation/quantitative study with source artifacts, model outputs, claim registers, sensitivity summaries, denominators, and release rules.

The current limitation is novelty. SSES has not yet been externally benchmarked across independent simulation programs. Future methods work should apply the stack to non-LAKANA simulations and peer-reviewed quantitative studies to test inter-rater reliability, scoring sensitivity, and whether SSES reduces overclaiming or improves reproducibility review.

21.2 SSES validation and benchmarking agenda

Because SSES is newly proposed, its strongest current claim is conceptual and procedural, not external validation. A journal-grade SSES validation program should test four properties across non-LAKANA studies: (i) inter-rater reliability when independent reviewers score the same paper, (ii) scoring sensitivity to admissible rule-weight changes, (iii) ablation behavior when source artifacts, denominators, or sensitivity links are removed, and (iv) portability to unrelated simulation or survival-analysis papers. Until those studies are completed, SSES should be described as a proposed post-run evidence-governance method with internal demonstration, not as a validated universal standard.

Validation target	Test design	Pass criterion	Current status
Inter-rater reliability	Multiple reviewers independently apply SSES to the same public claims.	High agreement on claim eligibility, lineage, denominator, and overclaim flags.	Future methods study.
Weight/rule sensitivity	Vary admissible scoring weights and gate thresholds.	Public claim decisions remain stable or instability is transparently flagged.	Future methods study.
Ablation behavior	Remove source tables, sensitivity outputs, denominators, or captions.	SSES downgrades or blocks affected claims.	Demonstrated conceptually; needs external benchmark.
Non-LAKANA portability	Apply SSES to unrelated simulations or peer-reviewed survival papers.	SSES identifies claim-to-evidence weaknesses without changing domain truth.	Proposed; not externally benchmarked.

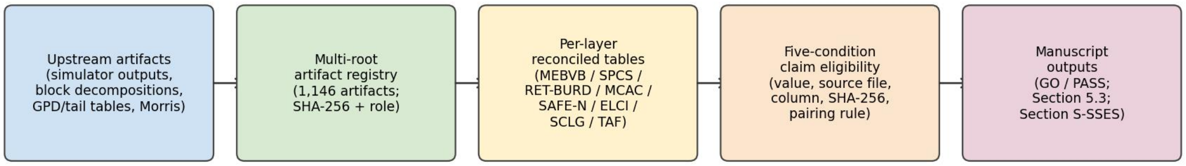

Figure ES-3. SSES evidence-governance pipeline (source-reconciled schematic).

Figure 13: SSES evidence-governance pipeline. SSES begins after simulation completion and evaluates whether source artifacts support public claims.

22 MEBVB: precision proof layer

MEBVB is the precision layer. It asks whether the simulator-internal block-level burden-delta estimate has a favorable sign that remains positive after a variance-aware empirical bound. For B blocks with mean $\bar{\mu}$, sample variance s^2, range R, and confidence level $1-\alpha$,

$$\varepsilon = \frac{\frac{2}{3}R\ln(2/\alpha) + \sqrt{(\frac{2}{3}R\ln(2/\alpha))^2 + 8Bs^2\ln(2/\alpha)}}{2B}, \qquad [\bar{\mu}-\varepsilon, \bar{\mu}+\varepsilon].$$

Quantity	Whole-frame value
Effective precision blocks B	390
Mean burden-delta unit	1,079,683.4568
Sample variance	706,192,094,291.83
Local range	3,237,388.1840
95% epsilon	126,239.2514
Lower bound	953,444.2055
Upper bound	1,205,922.7082
Favorable sign stable	True

What it proves: within the simulation and block design, the reconciled burden-delta sign remains favorable under the reported empirical-Bernstein bound. What it does not prove: real-world burden reduction, clinical benefit, or safety effectiveness.

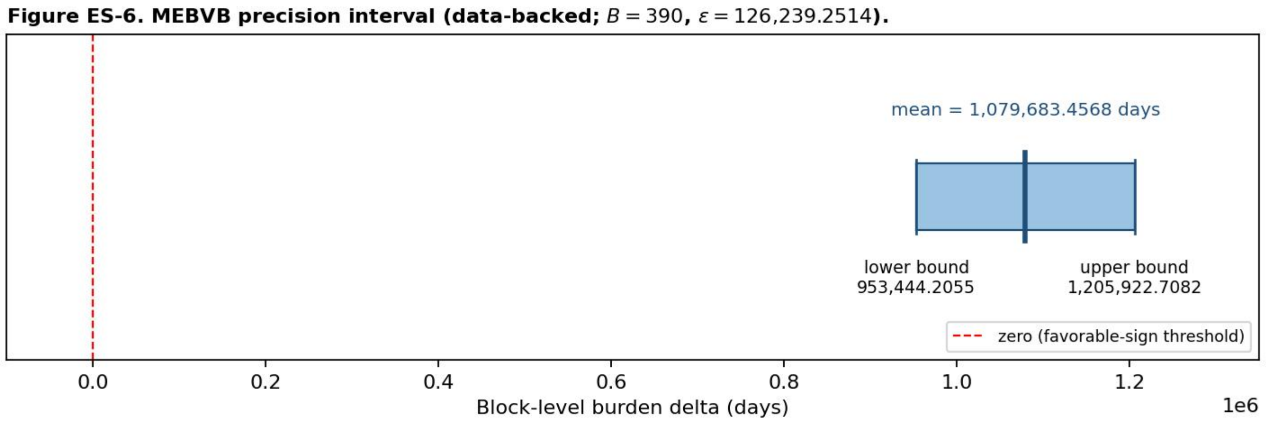

Figure 14: MEBVB precision interval. This is block-level simulator precision, not a field-data confidence interval.

23 SPCS: structural persistence concordance score

SPCS is the source-path sign-governance layer. It asks whether model-family rows and source-path entries contain explicit favorable or adverse direction language and whether contradictions exist. The canonical score has the form

$$\text{SPCS} = \frac{\sum_j w_j s_j}{\sum_j w_j} \cdot (1 - \text{ContradictionRate}),$$

where w_j is the source status weight and s_j is a signed evidence indicator. In the reported stack, no directional sign is inferred merely from model-family presence. This is why a zero SPCS can mean neutral source verification with full coverage and no contradiction, not an adverse result.

What it proves: the model-family/source-path evidence does not contain unsupported directional sign inference or contradiction under the SPCS rules. What it does not prove: outcome superiority.

24 RET-BURD: retained-time/burden surface

RET-BURD pairs retained-time movement with burden movement. For paired row r,

$$Q_r = \text{sign}(\Delta T_r, \Delta B_r),$$

with quadrants defined by favorable or adverse retained time and favorable or adverse burden. The result is not a single score. It is a surface. That matters because SSI may increase retained time while also shifting subjects into restriction or chronic-management states.

Quantity	Value
Quadrant-eligible rows	390
Excluded rows	0
Favorable retention / lower burden	360
Favorable retention / higher burden tradeoff	30
Adverse or ambiguous quadrants	0

What it proves: the retained-time/burden relation is mostly favorable in the reconciled surface, with a visible managed-burden tradeoff region. What it does not prove: clinical benefit or harm.

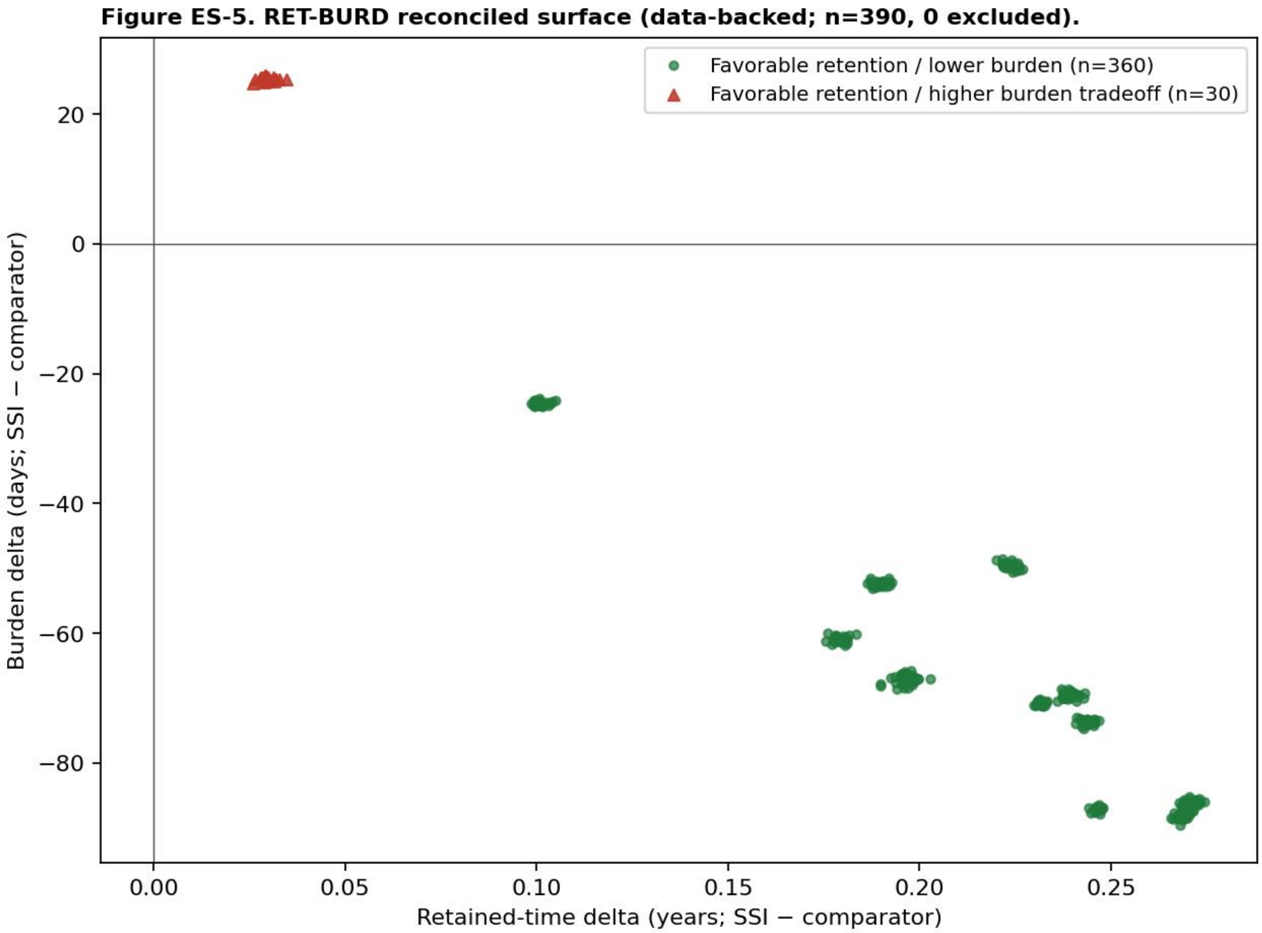

Figure 15: RET-BURD reconciled surface. Each point is a paired simulator row; the surface does not establish injury-burden reduction in any field population.

25 MCAC: model concordance adjudication cube

MCAC is the model-family census and status-weight matrix. It exists because a hostile reviewer will ask whether advanced model families were actually computed, absent, or merely not displayed. For model family j,

$$R_j = w(\text{status}_j), \qquad \bar{R} = \frac{1}{J}\sum_{j=1}^{J} R_j.$$

The readiness value is not evidence strength. It is a public-release/readiness posture.

Model family	Public status	Readiness
RMST	completed source-reconciled	1.0
Aalen-Johansen CIF	included as event-history family; source-reconciled	0.5
Fine-Gray	included as competing-risk family; source-reconciled	0.5
Cause-specific Cox	included as cause-specific hazard family; source-reconciled	0.5
Markov	completed source-reconciled	1.0

Model family	Public status	Readiness
Semi-Markov	included as duration-aware transition family; source-reconciled	0.5
Morris	completed source-reconciled	1.0
Sobol/GSA	included as global-sensitivity family; source-reconciled	0.5
Bootstrap/convergence	completed source-reconciled	1.0
Tail/GPD	completed source-reconciled	1.0
Rare-event side-harness	completed source-reconciled	1.0
Pareto	completed source-reconciled	1.0
Empirical-Bayes shrinkage	completed source-reconciled	1.0
Leave-one-lane-out / LOO	completed source-reconciled	1.0
Multiverse / specification stability	completed source-reconciled	1.0
Basketball female/male concordance	completed source-reconciled	1.0
Missingness / signal-quality	completed source-reconciled	1.0
Biological-boundary BLI	completed source-reconciled	1.0
Privacy / source / burden closure	completed source-reconciled	1.0

What it proves: the analysis stack is a model-family system, not a single-outcome report. What it does not prove: that every model family carries equal evidentiary strength or external validity.

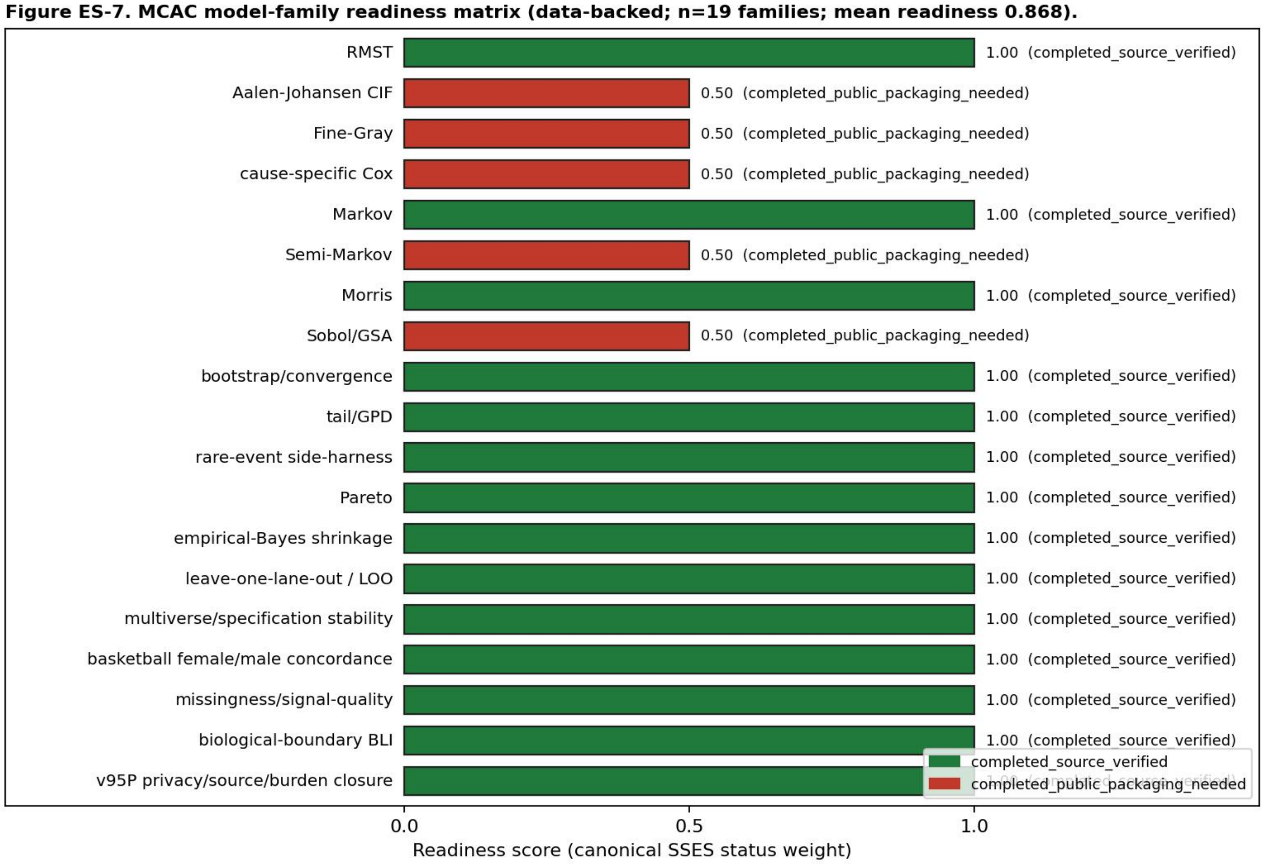

Figure 16: MCAC model-family readiness matrix. Readiness is a SSES status-weight score, not clinical, field, or external-validity evidence.

26 SAFE-N and ELCI: public evidence-release sufficiency and lineage completeness

SAFE-N scores whether the public evidence package is sufficient and bounded. It is not a privacy guarantee. ELCI scores whether claims have evidence-lineage completeness. It is not evidence strength.

Metric	Value
SAFE-N base score	0.813
Exposure penalty	0.029
SAFE-N adjusted	0.789
ELCI global mean	0.918

What SAFE-N proves: the public release has a quantified governance posture under the score definition. What ELCI proves: claims are mostly covered by evidence-lineage components. Neither proves privacy, safety, or real-world effectiveness.

SAFE-N and ELCI must not be interpreted as differential-privacy parameters (ϵ, δ), zero-knowledge proofs, secure multiparty computation guarantees, mobile-client side-channel resistance, de-identification proof, anonymity proof, no-reidentification-risk proof, or device-level cryptographic validation. They are metadata and release-governance scores: they evaluate whether the public evidence package is sufficiently source-linked, bounded, and reviewable under its own rule set. They do not certify any physical device, cloud service, firmware stack, mobile client, or adversary model.

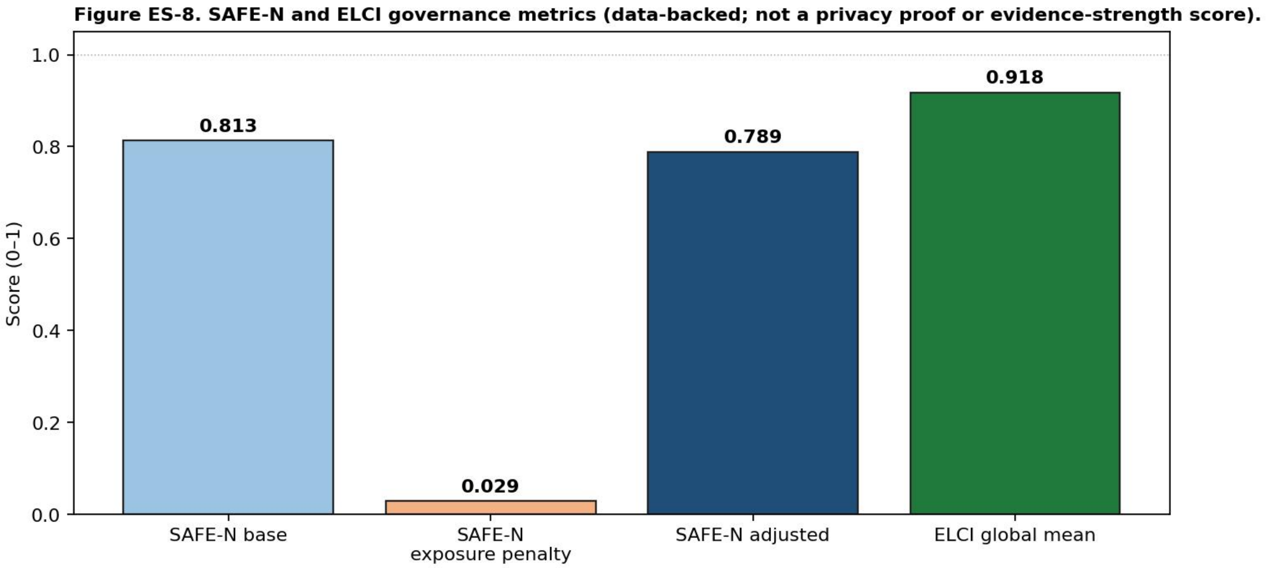

Figure 17: SAFE-N and ELCI governance metrics. These are evidence-release and lineage-completeness metrics, not privacy proofs or evidence-strength scores.

27 SCLG: sensitivity-to-claim leverage gate

SCLG is the model that prevents sensitivity analysis from floating away from claims. A high-leverage parameter is relevant only if it threatens a public claim. For parameter j and claim c,

$$\mathrm{SCLG}_{j,c} = \ell_j \cdot \mathbb{I}\{c \text{ depends on } j\} \cdot \mathbb{I}\{\text{claim boundary metric linked}\}.$$

Claim stability requires the linked precision or boundary metric to remain valid. In the reported SCLG layer, the dominant source metric is `restriction_burden_delta_days`; top normalized leverage is approximately 0.97–1.00; the leverage class is high; and claim stability remains stable because the relevant MEBVB lower bound is positive.

What it proves: high sensitivity was not ignored; it was tied back to claim stability. What it does not prove: causal parameter importance in real systems.

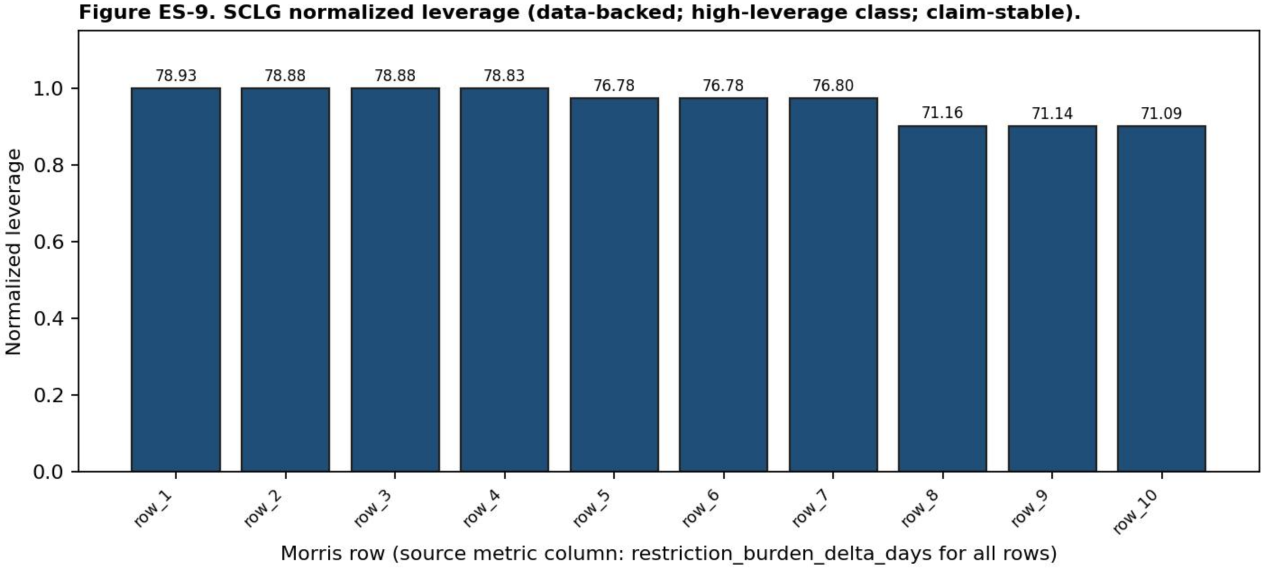

Figure 18: SCLG normalized leverage. This is simulator-internal Morris elementary-effect screening linked to claim boundaries, not causal attribution.

28 TAF: tail adjudication frontier

TAF makes rare-event rows admissible only when tail event count and exposure denominator are paired from the same upstream context. For row r,

$$\text{TAFEligible}_r = \mathbb{I}\{n_{exceed,r} \text{ present}\}\mathbb{I}\{n_r \text{ present}\}\mathbb{I}\{\text{source}(n_{exceed,r}) = \text{source}(n_r)\}\mathbb{I}\{n_r > 0\}.$$

The layer prevents rare-event reporting without denominators.

Quantity	Value
Claim-eligible rows	156
Excluded rows	0
Tail event count column	`n_exceed`
Exposure denominator column	`n`
Reading	Simulator-internal tail adjudication

What it proves: tail rows are denominator-paired and source-reconciled. What it does not prove: real-world rare-event incidence.

Figure ES-10. TAF denominator reconciliation (source-reconciled schematic).

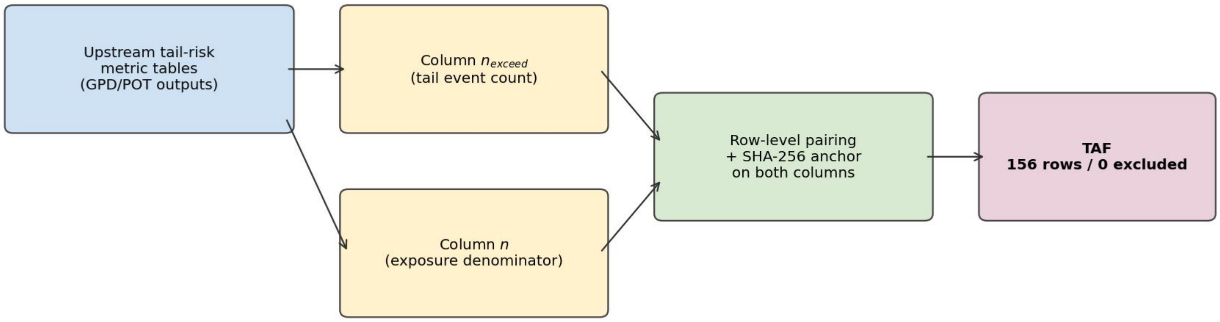

TAF is a simulator-internal tail-adjudication governance artifact; it is not real-world catastrophic-event prediction.

Figure 19: TAF denominator reconciliation. TAF is tail-adjudication governance, not real-world rare-event prediction.

29 Lane-specific verified contrasts

The female-lane contrasts below are included because they define the companion paper's source-reconciled lane structure and show how retained time, burden, and terminal-proxy movement are interpreted without claiming field validation.

Lane	Estimand	Comparator	SSI	Absolute Δ	Relative Δ
Women's basketball	Missed days	169.16	91.37	-77.79	-45.98%
Women's basketball	ACL-family proxy	1.876	0.596	-1.280	-68.25%
Women's basketball	Time-years	0.194	0.449	0.256	132.25%
Women's soccer	Missed days	158.29	89.98	-68.31	-43.15%
Women's soccer	ACL-family proxy	1.754	0.568	-1.187	-67.63%
Women's soccer	Time-years	0.250	0.470	0.220	88.33%
Softball	Missed days	101.50	63.06	-38.43	-37.87%
Softball	ACL-family proxy	1.121	0.117	-1.004	-89.58%
Softball	Time-years	0.446	0.595	0.149	33.30%

These values are simulator-internal pathway contrasts. They show lane-specific modeled movement. They do not establish demographic validation, injury prevention, or field performance.

30 Privacy, anti-surveillance, and release boundary

The public manuscript evaluates aggregates $z = \mathcal{A}(X)$, not unrestricted trajectory collections X. That distinction is mathematical and governance-relevant. It allows public review of reported aggregate quantities without releasing raw subject pathways or implementation-defining internals.

Aspect	Public claim	Bounded against
Raw trajectory collection X	Not publicly released.	No public raw-subject surveillance substrate.
Aggregate z	Reviewable through SSES tables.	Not clinical, field, or regulatory evidence.
Many-to-one $\mathcal{A}$	Structural non-invertibility orientation.	Not a formal privacy proof.
Release envelope $\mathcal{E}_k$	Monotone governance design target.	Not verified zero-knowledge implementation.
SAFE-N adjusted	Public-release sufficiency score.	Not differential privacy, de-identification, anonymity, or no-reidentification-risk.

Figure ES-11. Privacy release-boundary diagram (conceptual schematic).

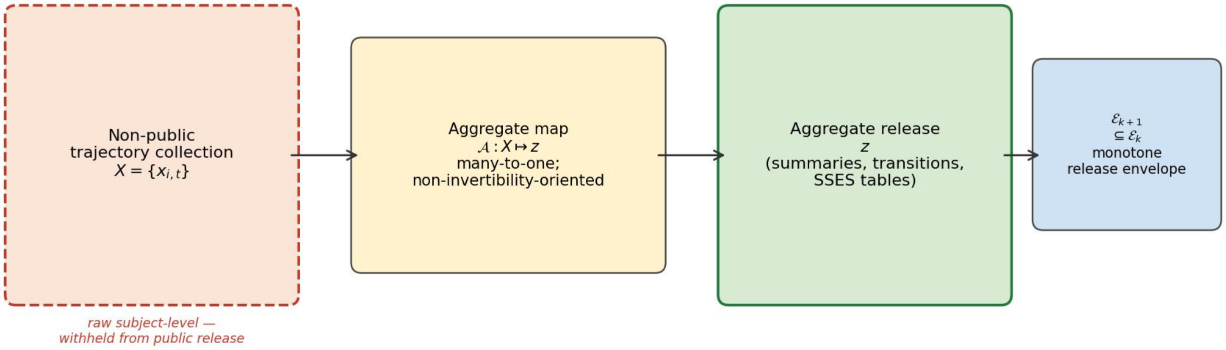

Structural property; not a formal privacy proof, differential-privacy claim, or no-reidentification guarantee.

Figure 20: Privacy release-boundary diagram. The public release boundary is a structural publication map and governance rule, not a formal privacy proof.

31 Reproducibility and artifact hierarchy

The reproducibility package supports verification of reported numeric outputs against SHA-anchored public artifacts. It does not support independent cloning of the simulator. That is intentional: cloning requires non-public coefficients, thresholds, weights, fusion logic, trigger rules, and raw trajectories.

Class	Examples	Public role
Manuscript outputs	Main paper, companion, engineering supplement.	Public interpretation and review.
Reconciled SSES tables	MEBVB, RET-BURD, MCAC, SAFE-N, ELCI, SCLG, TAF.	Numeric source verification.
Audit outputs	Manifests, registries, QA reports.	Provenance and claim boundary.

Class	Examples	Public role
Upstream computation inputs	Burden, tail, sensitivity, event-history, and model-output tables.	Source-anchored evidence; not all are public.
Non-public implementation	Simulator code, coefficients, thresholds, fusion internals, raw trajectories.	Withheld trade-secret and privacy-sensitive layer.

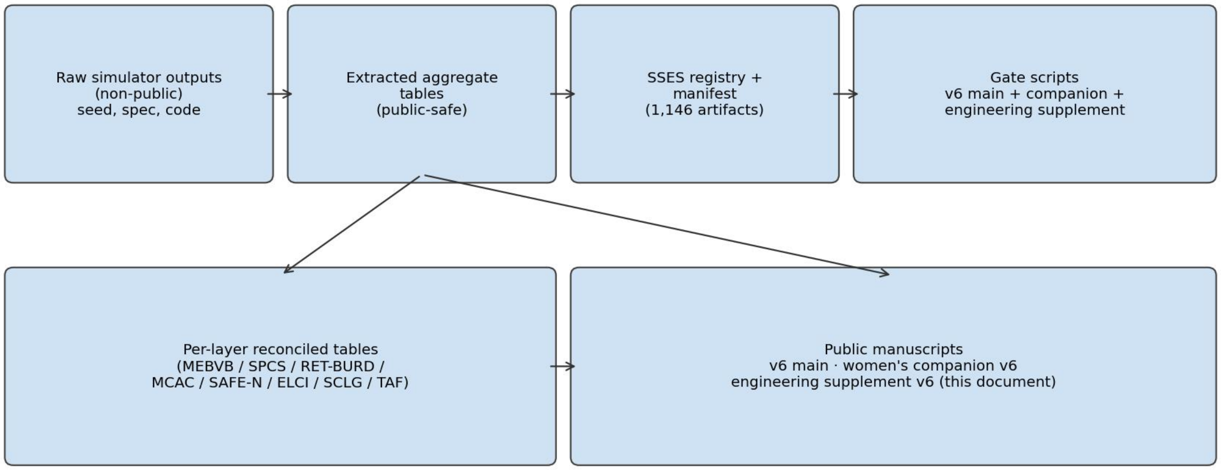

Figure ES-12. Reproducibility artifact flowchart (source-reconciled schematic).

Independent re-implementation of the simulator requires non-public coefficients, thresholds, weights, fusion details, and trigger logic.

Figure 21: Reproducibility artifact flowchart. Independent re-implementation requires non-public coefficients, thresholds, weights, fusion details, and trigger logic.

32 Engineering attack-surface analysis

The attack surface is not only code-level. It includes synthetic pathway assumptions, comparator construction, biological calibration, hidden proprietary thresholds, denominator pairing, MCAC packaging, SAFE-N interpretation, tail denominator handling, Morris/Sobol linkage, rare-event interpretation, privacy overclaim, regulatory overclaim, and external validation roadmap. The supplement addresses these as engineering constraints, not as afterthoughts.

32.1 Adversarial reviewer attack-surface register

The table below states predictable reviewer concerns as engineering attack surfaces. A concern is not treated as solved unless the present suite actually addresses it. Each row separates current mitigation from remaining future work.

Reviewer attack surface	Likely criticism	Current mitigation in this suite	Remaining future work
Simulation-only evidence	The results may not transport to real athletes or workers.	Claim boundary repeatedly states simulator-internal evidence only.	Prospective pilots with adjudicated outcomes and independent review.

Comparator abstraction	The comparator may not match current practice.	Comparator families are named as external-load, internal-recovery, ACWR-style, hybrid-SOTA, and industry-ablation abstractions.	Prospectively specified external comparators.
Novel SSES metrics	The governance stack is new and LAKANA-originated.	SSES is defined as post-run, non-generative, source-linked, and portable in principle.	Independent benchmarking on non-LAKANA studies.
Bounded-telemetry release is not privacy proof	Aggregate maps and role gates may be mistaken for formal privacy.	The supplement explicitly rejects differential privacy, anonymity, de-identification, zero-knowledge, and device-security claims.	Formal adversary models and security/privacy audits.
Trade-secret-bounded reproducibility	Public materials do not allow full independent simulator cloning.	Public-safe equations, tables, figures, checksums, and claim boundaries are released; clone-enabling internals are withheld.	Governed-access reproducibility package under purpose limitation.
Basketball female/male public table boundary	Sex-labeled basketball strata may appear under-reported.	`basketball_female` and `basketball_male` are preserved as design strata; incomplete side-by-side sex-comparison is avoided.	Complete common-provenance basketball sex-lane public table.
Female-stack transportability	Female-stack simulation does not prove real-world female-specific efficacy.	Companion keeps B=180 scope, lane boundaries, and non-claims explicit.	Prospective women's/female sport validation with role-specific data.
Tail/rare-event interpretation	Tail outputs may be overread as real-world catastrophe rates.	TAF and tail/GPD rows are framed as denominator governance and simulator-prior bookkeeping.	External rare-event data and predefined tail protocols.
Formal theorem/proof limitations	Public theorem boxes may be mistaken for production proof.	Theorems are narrow public-abstraction results with explicit non-claims.	Formal verification or implementation proofs where appropriate.
Figure/table public packaging	Figures may overstate or omit source context.	Captions include source basis and claim boundaries; checksums and registries support review.	Venue-specific figure QA and vector/source-file submission.
External calibration	Physics-pressure variables are not field-calibrated tissue laws.	Equations are labeled phenomenological and public-safe.	Biomechanics, physiology, recovery, and outcome calibration.
Device/hardware validation	TSARO/NICOLE are described without deployed hardware proof.	Architecture-level boundary is explicit.	Hardware-in-the-loop, latency, protected-memory, and side-channel validation.
SOS/CivOS adjacent architecture boundary	Adjacent architecture could be misread as current SSI evidence.	SOS/CivOS is future/adjacent only and not used to upgrade SSI claims.	Separate SOS/CivOS privacy, survivability, and field-validation studies.

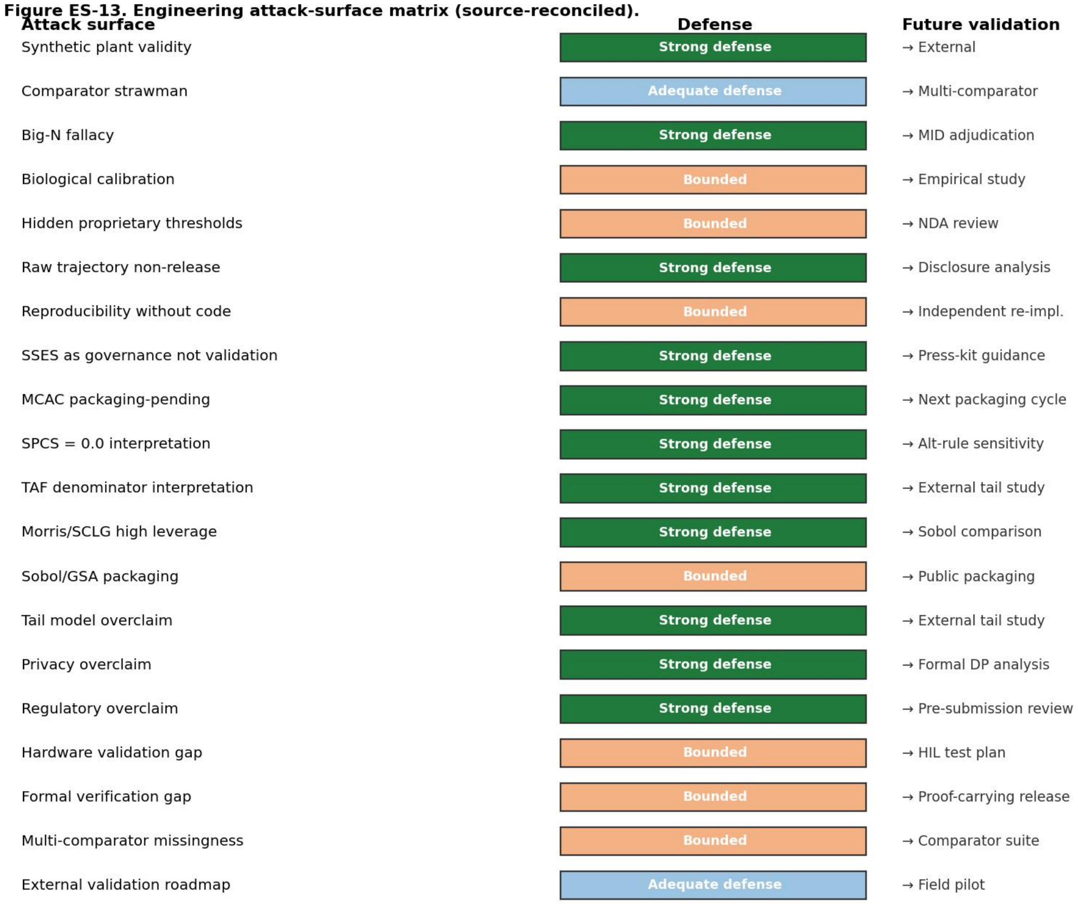

Figure ES-13. Engineering attack-surface matrix (source-reconciled).

Attack surface	Defense	Future validation
Synthetic plant validity	Strong defense	→ External
Comparator strawman	Adequate defense	→ Multi-comparator
Big-N fallacy	Strong defense	→ MID adjudication
Biological calibration	Bounded	→ Empirical study
Hidden proprietary thresholds	Bounded	→ NDA review
Raw trajectory non-release	Strong defense	→ Disclosure analysis
Reproducibility without code	Bounded	→ Independent re-impl.
SSES as governance not validation	Strong defense	→ Press-kit guidance
MCAC packaging-pending	Strong defense	→ Next packaging cycle
SPCS = 0.0 interpretation	Strong defense	→ Alt-rule sensitivity
TAF denominator interpretation	Strong defense	→ External tail study
Morris/SCLG high leverage	Strong defense	→ Sobol comparison
Sobol/GSA packaging	Bounded	→ Public packaging
Tail model overclaim	Strong defense	→ External tail study
Privacy overclaim	Strong defense	→ Formal DP analysis
Regulatory overclaim	Strong defense	→ Pre-submission review
Hardware validation gap	Bounded	→ HIL test plan
Formal verification gap	Bounded	→ Proof-carrying release
Multi-comparator missingness	Bounded	→ Comparator suite
External validation roadmap	Adequate defense	→ Field pilot

Figure 22: Engineering attack-surface matrix. Defense rating refers to manuscript discipline and source reconciliation, not field, clinical, regulatory, hardware, or privacy validation.

33 Technical limitations

1. The simulator's plant fidelity is not externally validated.
2. The modeled constraint region $\mathcal{C}$ is not an empirically calibrated tissue-failure surface.
3. The comparator is a modeled baseline, not a census of every real-world monitoring, coaching, wearable, clinical, or occupational practice.
4. Effect-size reporting does not rely on externally adjudicated minimum-important-difference thresholds.
5. Event-history models support pathway interpretation inside the simulation, not clinical incidence.
6. Tail models and rare-event harnesses are simulator-prior artifacts, not real-world catastrophic-event prediction.
7. The aggregate publication map is a structural property, not a statistical-disclosure or cryptographic adversary proof.
8. Independent re-implementation of the simulator is not possible from the public package alone.
9. Hardware-in-the-loop, sensor-latency, protected-memory, firmware, and transport-layer behavior are adjacent engineering problems and are not validated here.
10. Any future multi-transport survivability claim must include common-cause coupling, not only an independent-channel product term.

11. NICOLE and TSARO are described at architecture level; deployed cryptographic or hardware claims require separate validation.

34 Engineering validation roadmap

The roadmap is ordered by what would most improve external credibility without weakening the current claim boundary.

Priority	Item	Type
1	Hardware-in-the-loop test plan with sensor-latency characterization.	External validation.
2	Protected-memory and side-channel review for telemetry-ingesting devices.	External validation.
3	Formal statistical-disclosure analysis of the aggregate publication map.	Governance validation.
4	External clinical/sports-science validation cohort with adjudicated outcomes.	External validation.
5	Multi-comparator simulation suite.	Simulation extension.
6	Sex-stratified and community-reviewed validation design with consented variables.	External validation.
7	Governed-access reproducibility package under purpose limitation and revocation.	Reproducibility extension.
8	Field pilot design after institutional review with athlete/worker representation and adverse-use controls.	External validation.
9	Public packaging of event-history and global-sensitivity display tables where not yet included.	Packaging cycle.
10	Cross-method sensitivity comparison linking Morris, Sobol/GSA, and SCLG.	Sensitivity extension.
11	Protocolized RMST tau-grid sensitivity with predeclared horizon-selection rules.	Statistical methods extension.
12	SSES inter-rater, ablation, and non-LAKANA benchmarking study.	Governance-method validation.

35 Conclusion

This technical manuscript defines SSI as a deterministic human structural-health monitoring framework under bounded-telemetry release. It gives reviewers the mathematical and engineering spine of the system: a subject-time state vector, public-safe physics-pressure variables, admissible safe sets, saturation-bounded minimum-intervention projection, TSARO deterministic physics authority, NICOLE cryptographic governance, event-history model families, sensitivity and robustness model families, tail and rare-event models, lane-integrity models, and the eight-layer SSES post-run evidence-governance stack.

The central engineering result is not a clinical claim. It is a public-review architecture. SSI separates the simulation plant, the deterministic policy abstraction, the cryptographic governance boundary, and the post-run evidence-governance stack. SSES then determines what can be reported, how it is sourced, what it means, and what it cannot mean. That separation is what allows the work to be mathematically explicit and trade-secret safe at the same time.

The model families serve different proof functions. RMST and event-history models define retained time and pathway movement. Markov and Semi-Markov models define state dynamics and dwell structure. Fine-Gray and cause-specific

Cox models separate competing-risk cumulative incidence from process hazards. Morris and Sobol/GSA identify parameter leverage. GPD/POT and TAF prevent tail outputs from being denominator-free anecdotes. MEBVB establishes block-level precision. RET-BURD shows retained-time/burden trade space. MCAC prevents model-family misclassification. SAFE-N and ELCI govern public release sufficiency and evidence-lineage completeness. SCLG ties sensitivity back to claims. Together, these layers make the public evidence surface technically reviewable without disclosing proprietary implementation details or overstating simulator-internal results as field truth.

A Appendix A. Full model dossiers

This appendix gives a compact engineering dossier for every model family included in the SSI technical stack. The purpose is to make the review surface explicit: each model is listed with its object, reason for inclusion, public mathematical form, primary output, proof function, and boundary.

A.1 RMST dossier

Object. Retained modeled time over a fixed horizon. **Reason.** A time-scale estimand is more interpretable than a hazard ratio when the question is whether modeled pathways remain active longer. **Form.** $\mathrm{RMST}(\tau) = \int_0^\tau S(t)dt$. **Output.** Retained-time contrast between SSI and comparator. **Proof function.** Establishes simulator-internal retained-time direction and magnitude. **Boundary.** Not field longevity and not clinical survival.

A.2 Markov dossier

Object. Discrete transition dynamics. **Reason.** Terminal outcomes alone hide pathway structure. **Form.** $P_{rq} = \Pr(X_{t+\Delta t} = q \mid X_t = r)$. **Output.** State-occupation shares and adjacent transition counts. **Proof function.** Establishes that SSI changes pathway movement rather than only final labels. **Boundary.** Not an empirical transition matrix for real athletes.

A.3 Semi-Markov dossier

Object. Transition plus dwell-time structure. **Reason.** Restriction and chronic-management states have duration. **Form.** $Q_{rq}(u) = \Pr(X_{n+1} = q, T_{n+1} - T_n \leq u \mid X_n = r)$. **Output.** Duration-aware transition summaries. **Proof function.** Explains why longer retained time can coexist with increased managed burden. **Boundary.** Not empirical dwell-time validation.

A.4 Aalen-Johansen dossier

Object. Nonparametric multi-state transition probability and cumulative incidence. **Reason.** SSI has competing terminal paths and non-terminal state transitions. **Form.** $\widehat{P}(s,t) = \prod_{(s,t]}(I + d\widehat{\Lambda}(u))$. **Output.** Cumulative transition/CIF summaries. **Proof function.** Establishes event-history coherence under competing pathways. **Boundary.** Not field incidence.

A.5 Fine-Gray dossier

Object. Subdistribution hazard for a target event. **Reason.** A competing terminal event can prevent observation of the target terminal family. **Form.** $\lambda_k^{FG}(t \mid Z) = \lambda_{k0}^{FG}(t)\exp(\beta_k^\top Z)$. **Output.** Target-event cumulative-incidence association. **Proof function.** Shows target-family behavior while retaining competing events in the cumulative-incidence structure. **Boundary.** Not causal prevention.

A.6 Cause-specific Cox dossier

Object. Cause-specific process intensity. **Reason.** Fine-Gray and cause-specific Cox answer different questions. **Form.** $\lambda_k(t \mid Z) = \lambda_{k0}(t)\exp(\gamma_k^\top Z)$. **Output.** Process-hazard comparison by cause family. **Proof function.** Separates process intensity from cumulative-incidence interpretation. **Boundary.** Not clinical hazard calibration.

A.7 Bootstrap/convergence dossier

Object. Monte Carlo precision and stability. **Reason.** A simulation result is weak if it is numerically unstable. **Form.** $\widehat{\mathrm{Var}}(\hat{\theta}) = \frac{1}{B-1}\sum_b(\hat{\theta}_b - \bar{\theta}_B)^2$. **Output.** Resampling variability and convergence indicators. **Proof function.** Establishes numerical stability inside the run configuration. **Boundary.** Not external validation.

A.8 Morris dossier

Object. Elementary-effect screening. **Reason.** The parameter space is large and needs efficient screening. **Form.** $EE_j(x) = \{f(x + \Delta e_j) - f(x)\}/\Delta$. **Output.** Ranked local/global screening effects. **Proof function.** Identifies high-leverage inputs that require claim caution. **Boundary.** Not causality.

A.9 Sobol/GSA dossier

Object. Global variance decomposition. **Reason.** Reviewer confidence improves when output variance is not hidden behind one unexplored assumption. **Form.** $S_j = \mathrm{Var}_{X_j}(\mathbb{E}[Y \mid X_j])/\mathrm{Var}(Y)$ and $S_{Tj} = 1 - \mathrm{Var}_{X_{\sim j}}(\mathbb{E}[Y \mid X_{\sim j}])/\mathrm{Var}(Y)$. **Output.** First-order and total-order sensitivity. **Proof function.** Establishes global uncertainty attribution. **Boundary.** Not validation of parameter truth.

A.10 GPD/POT dossier

Object. Threshold-exceedance tail structure. **Reason.** Rare events must not be buried by means. **Form.** $\Pr(Y \le y \mid X > u) = 1 - (1 + \xi y/\beta)^{-1/\xi}$. **Output.** Tail-exceedance summaries. **Proof function.** Establishes denominator-paired tail bookkeeping. **Boundary.** Not real-world catastrophe prediction.

A.11 Rare-event side-harness dossier

Object. Side accounting for low-frequency terminal families. **Reason.** Rare outputs need visibility without becoming headline claims. **Form.** $\hat{p}_k = n_k/N_k$ with explicit denominator pairing. **Output.** Rare-event count, denominator, and fraction. **Proof function.** Keeps rare terminal families auditable. **Boundary.** Not empirical incidence.

A.12 Pareto dossier

Object. Tradeoff frontier. **Reason.** SSI can shift burden rather than eliminate all burden. **Form.** Row a dominates row b if $r_a \ge r_b$ and $b_a \le b_b$ with one strict inequality. **Output.** Non-dominated retained-time/burden rows. **Proof function.** Establishes tradeoff structure. **Boundary.** Not a universal optimum.

A.13 Empirical-Bayes shrinkage dossier

Object. Stabilized lane estimates. **Reason.** Thin lane cells should not be overread. **Form.** $\tilde{\theta}_s = w_s\hat{\theta}_s + (1 - w_s)\mu$. **Output.** Shrunk lane estimate. **Proof function.** Reduces noise-driven overinterpretation. **Boundary.** Not imputation of missing evidence.

A.14 Leave-one-lane-out dossier

Object. Lane-deletion robustness. **Reason.** A whole-frame claim should not secretly depend on one lane. **Form.** $\Delta_{(-s)} = g(\mathcal{D} \setminus \mathcal{D}_s)$. **Output.** Claim movement after lane removal. **Proof function.** Establishes lane dependence or independence. **Boundary.** Not subgroup validation.

A.15 Multiverse/specification stability dossier

Object. Robustness across admissible specifications. **Reason.** A result that appears under only one admissible specification is fragile. **Form.** $\{\Delta^{(v)} : v \in \mathcal{V}\}$ over admissible specifications. **Output.** Specification distribution. **Proof function.** Shows whether conclusion direction is stable. **Boundary.** Not proof of true mechanism.

A.16 Basketball female/male concordance dossier

Object. Sex-labeled lane integrity. **Reason.** Basketball must not be collapsed into a generic lane when sex-labeled outputs exist. **Form.** Product of lane-label, source-path, estimand-match, and claim-boundary indicators. **Output.** Concordance status. **Proof function.** Establishes lane identity preservation. **Boundary.** Not sex-comparison rhetoric and not biological equivalence.

A.17 Missingness/signal-quality dossier

Object. Missingness and signal integrity. **Reason.** Missing or malformed data can generate false confidence. **Form.** $\text{Eligible}_{j,s} = \mathbb{I}\{M_{j,s} = 0\}\mathbb{I}\{\text{source reconciled}\}\mathbb{I}\{\text{pairing valid}\}$. **Output.** Eligibility flag. **Proof function.** Prevents missingness from becoming evidence. **Boundary.** Not telemetry validation.

A.18 Biological-boundary index dossier

Object. Modeled pressure-to-capacity ratio. **Reason.** Structural-load safety depends on demand relative to capacity, not external load alone. **Form.** $\rho_t = (L_t + \Delta L_t + H_t + F_t + D_t)/(C_t + \epsilon_C)$. **Output.** Boundary pressure. **Proof function.** Supports pathway-mechanics interpretation. **Boundary.** Not a measured biological limit.

A.19 Privacy/source/burden closure dossier

Object. Joint source, burden, and release-boundary eligibility. **Reason.** A valid value can still be an invalid public claim if it violates release or overclaim boundaries. **Form.** Product of source-eligible, burden-paired, release-permitted, and prohibited-claim-absent indicators. **Output.** Closure flag. **Proof function.** Establishes public claim admissibility. **Boundary.** Not a privacy proof.

B Formal public-safe results

The following results are formal properties of the public abstraction. They are intentionally narrow. They do not prove clinical efficacy, field effectiveness, production safety, or deployed cryptographic security.

Theorem 1 (Saturation-bounded projection existence-or-fail-closed). *Let $\mathcal{U}_{\max}$ be compact and non-empty, let $\mathcal{U}$ be closed, and define the feasible set*

$$\mathcal{F}_t = \{u \in \mathcal{U} \cap \mathcal{U}_{\max} : F_\theta(x_t, u, e_t) \in \mathcal{C}_t\}.$$

If $\mathcal{F}_t$ is non-empty and closed and the cost $\|u\|_W^2$ is continuous, then the minimum-intervention projection attains a minimizer. If $\mathcal{F}_t$ is empty, the public policy cannot request an unbounded correction and must route to a fail-closed state family.

Proof. Because $\mathcal{U}_{\max}$ is compact and $\mathcal{U}$ is closed, $\mathcal{U} \cap \mathcal{U}_{\max}$ is compact when non-empty. A closed subset $\mathcal{F}_t$ of that compact set is compact. A continuous quadratic cost attains a minimum on a compact feasible set. If $\mathcal{F}_t$ is empty, no physiologically bounded action can satisfy the public constraint, so the policy cannot expand the feasible set and the defined route is fail-closed. This is a public abstraction result, not a production-controller proof. ☐

Theorem 2 (TSARO monotone fail-closed contraction)**.** *If uncertainty states satisfy $\omega_2 \geq \omega_1$ and TSARO defines admissible envelopes so that $\mathcal{C}_t(\omega_2) \subseteq \mathcal{C}_t(\omega_1)$, then increasing uncertainty cannot expand the permissible state/action set.*

Proof. For any candidate transition accepted under the higher-uncertainty envelope, the resulting state lies in $\mathcal{C}_t(\omega_2)$. Since $\mathcal{C}_t(\omega_2) \subseteq \mathcal{C}_t(\omega_1)$, that state also lies inside the lower-uncertainty envelope. The converse need not hold. Therefore uncertainty can preserve or reduce admissibility, but cannot expand it. ☐

Theorem 3 (SSES non-generation invariant)**.** *Let SSES layer outputs be $y_j = G_j(z, M)$, where z is an upstream public artifact and M is a manifest. If no G_j calls the simulator transition operator F_θ, samples trajectories, or mutates upstream model outputs, then SSES cannot create new Monte Carlo evidence.*

Proof. Each SSES layer is a deterministic map from completed artifacts and manifests to an eligibility, score, classification, or exclusion. Because G_j has no access to the simulator generator or subject-time sampler, its output is a function only of existing evidence. It can admit, exclude, qualify, classify, or bound evidence; it cannot generate new simulation evidence. ☐

Theorem 4 (Aggregate many-to-one release non-invertibility)**.** *Let $z = \mathcal{A}(X)$ be an aggregate public release where X is a subject-time trajectory collection and $\dim(z) \ll \dim(X)$. If $\mathcal{A}$ aggregates over subjects, lanes, trials, or horizons, then multiple distinct trajectory collections can map to the same public artifact z.*

Proof. Aggregation removes indices and degrees of freedom. For any aggregate mean, count, interval, or score, there are generally multiple underlying trajectory collections with the same aggregate value. Therefore $\mathcal{A}$ is many-to-one. This supports release-boundary discipline. It is not a formal privacy guarantee because no adversary model or privacy parameter is established. ☐

C Appendix B. Public-safe proof sketches

C.1 Saturation-bounded projection proof sketch

Let $\mathcal{U}_{\max}$ be compact and non-empty and let $\mathcal{U}$ be closed. If $\mathcal{U} \cap \mathcal{U}_{\max}$ is non-empty and the feasible set $\{u \in \mathcal{U} \cap \mathcal{U}_{\max} : F_\theta(x_t, u, e_t) \in \mathcal{C}_t\}$ is closed, then the continuous quadratic cost $\|u\|_W^2$ attains a minimizer on the feasible set. If the feasible set is empty, the policy does not extrapolate beyond biological action bounds; it returns the pathway to the fail-closed state family. This establishes existence-or-fail-closed behavior for the public abstraction, not a production controller proof.

C.2 TSARO monotone-contraction proof sketch

If uncertainty ω increases, the admissible envelope satisfies $\mathcal{C}(\omega_2) \subseteq \mathcal{C}(\omega_1)$ for $\omega_2 \geq \omega_1$. Therefore, a state/action pair admissible under higher uncertainty is also admissible under lower uncertainty, but not conversely. Thus uncertainty cannot expand permissible action space. This is the public mathematical form of fail-closed behavior.

C.3 NICOLE append-only proof sketch

If $\ell_k = H(\ell_{k-1} \parallel payload_k \parallel counter_k)$ and the counter is strictly monotone, then altering a past payload changes ℓ_k and every later digest. The public proof sketch establishes tamper-evidence of the chain structure, not deployed cryptographic security.

C.4 SSES non-generation proof sketch

SSES accepts upstream artifacts as input and applies deterministic eligibility maps G_j to produce layer outputs $y_j = G_j(z, \text{manifest})$. No G_j queries the simulator transition operator F_θ or samples new subject trajectories. Therefore, SSES cannot create new Monte Carlo evidence; it can only classify, bound, or reject existing evidence.

C.5 Aggregate non-invertibility proof sketch

The public release is $z = \mathcal{A}(X)$ with $\dim(z) \ll \dim(X)$ and many-to-one aggregation over subjects, lanes, trials, and horizons. Therefore, multiple trajectory collections map to the same public artifact. This supports release-boundary discipline, but it is not a formal privacy guarantee because no adversary model or privacy parameter is proven.

D Appendix C. Figure-to-model crosswalk

Figure	Model family or layer	Purpose
ES-1	System architecture	Shows simulator, policy, SSES, and manuscript-output layers.
ES-2	State machine	Shows admissible state classes and terminal families.
ES-2A	TSARO	Shows deterministic envelope enforcement and fail-closed contraction.
ES-3A	NICOLE	Shows ledger, access gate, DCL, and revocation mechanics.
ES-5A	SSI physics state-space	Shows load, reserve, thermal, fatigue, damage, and policy-action coupling.
ES-8A	Complete model stack	Shows all model families by engineering function.
ES-10A	SSES proof grid	Shows the eight SSES layers and their proof function.
ES-12A	Event-history models	Shows Markov, Aalen-Johansen, Fine-Gray, Cox, and RMST interpretation.
ES-16A	Sensitivity-to-claim	Shows Morris/Sobol/SCLG connection.
ES-18A	Trade-secret-safe determinism	Shows public form, non-public internals, aggregate outputs, and claim gate.
ES-5–ES-10	SSES data-backed layers	Show RET-BURD, MEBVB, MCAC, SAFE-N/ELCI, SCLG, and TAF.

E References

Aalen, O. O., and Johansen, S. (1978). An empirical transition matrix for non-homogeneous Markov chains based on censored observations. *Scandinavian Journal of Statistics, 5*(3), 141–150.

Andersen, P. K., Borgan, O., Gill, R. D., and Keiding, N. (1993). *Statistical Models Based on Counting Processes.* Springer.

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., and 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., and 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., and 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., and 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`

Geskus, R. B. (2016). *Data Analysis with Competing Risks and Intermediate States.* CRC Press. `https://doi.org/10.1201/b18695`

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., and Succop, P. (2005). 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., and Ford, K. R. (2005). 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., and 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., and 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., and 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., and Buchheit, M. (2013). Training adaptation and heart rate variability in elite endurance athletes. *International Journal of Sports Physiology and Performance, 8*(6), 688–691. `https://doi.org/10.1123/ijspp.8.6.688`

Putter, H., Fiocco, M., and 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., and 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., and Tarantola, S. (2008). *Global Sensitivity Analysis: The Primer.* Wiley. `https://doi.org/10.1002/9780470725184`

Sandve, G. K., Nekrutenko, A., Taylor, J., and 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., and Duran, B. (2010). Community-based participatory research contributions to intervention research. *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., and 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`