



VANTHEON

A Protocol for Verifiable Healthcare

Governed care, shared value, and learning healthcare intelligence

Alejandro S. Méndez Valentín

x/acc labs · amendez@xacc.io

Whitepaper v2.7 (publication candidate) · October 18, 2026

Abstract

Healthcare must coordinate care, sustain its resources and learn from experience while preserving the authority of affected people. These responsibilities often cross institutions with different records, incentives and decision powers. We propose VANTHEON, a Computed Learning Health Protocol that connects complete care, collective governance, shared value and authorized learning through versioned evidence and explicit obligations. Patients retain consent, representation and recourse. Humans, agents and robots contribute accountable services; service-specific provider trust can inform separately adopted economic and voting contracts. Codex governs rule improvement, Trust-Time coordinates deliverable capacity, and the Pool and Treasury organize care funding, eligible value sharing and settlement. THEON is the proposed healthcare foundation-model family and multi-agent research program, including parallel model-learning and recursive self-improvement studies. The contribution has three levels: an institutional reference architecture, selected formal and executable profiles, and a discriminating research program. Bounded reference and conventional workflows agree on 116 tested command attempts but both miss an omitted dependency. A restricted exact trust gate repairs five numerical false approvals, while synthetic estimators still under-cover uncertainty. Paired Pool simulations expose reduced resilience when surplus is distributed. These findings define tests of integration, incentives and access alongside computational correctness. THEON studies remain unexecuted. The architecture makes consequential interfaces inspectable; clinical benefit, governance superiority and model or research-procedure improvement require separate evaluation.

Keywords: learning healthcare; governance; clinical evidence; trust; economic participation; agents and robots; healthcare foundation models; recursive self-improvement.

Manuscript status. Publication candidate prepared for private author review. The full-system design and THEON program are proposals. Formal results and synthetic experiments retain their explicit assumptions; this edition reports no new clinical study.

Contents

1	Introduction: healthcare that can learn from its work	3
2	Participants, rights and accountable authority	7
3	A shared record with separate decision powers	10
4	Codex: governing better care	13
5	Evidence, attribution and service-specific trust	19
6	Coordinating humans, agents and robots	23
7	Trust-Time: allocating a deliverable service	25
8	Pool and Treasury: funding care and sharing value	26
9	THEON: learning models and improving research	31
10	Privacy, verification and resilient operation	38
11	What the current evidence establishes	39
12	Limits and the development program	44
13	Conclusion	47
	Reproducibility and declarations	48
	References	48
A	Formal Event and Transition Model	52
B	Service-Specific Trust	55
C	Approval-or-Review Rule	57
D	Settlement and Reserve Accounting	58
E	Restricted Integer Threshold Certificate	60
F	Conditional Fraud Calculation	62
G	Source Reconciliation and Preservation	63
H	Shared Interface Contracts	66
I	Retained Evidence and Methods	68
J	RP-01: A Bounded Institutional Reference Profile	70
K	TP-01: Exact bounded gate and synthetic trust comparisons	73
L	Economic reference profile and paired stress study	76
M	THEON model and procedure study charters	78
N	Candidate institutional contracts and discriminating evaluations	82
O	Derived views of retained synthetic results	88

Reading routes. For the purpose, evidence and limits, read Sections 1, 11 and 12–13. For institutional design, read Sections 2–8. THEON’s model and research-procedure program is in Section 9. Appendices A–F give the restricted formal results; the remaining appendices supply contracts, methods and provenance.

1. Introduction: healthcare that can learn from its work

Healthcare must do more than deliver an isolated intervention. People need timely, effective care they can afford and influence; professionals need the capacity and support to complete it; populations need sustainable financing and a way to improve what happens next. The scale is substantial in very different settings. WHO estimates that 4.6 billion people worldwide lacked full coverage of essential health services in 2023, while 2.1 billion experienced financial hardship from health spending in 2022; these overlapping populations measure different problems [1]. In the United States, national health expenditure reached \$5.3 trillion, or 18.0% of GDP, in 2024 [2]. Neither spending nor recorded activity alone establishes that care met its purpose.

Failures have several causes and require different remedies. A service may be unavailable, denied, unaffordable, performed unsafely or left without follow-up. WHO's 2023 global synthesis associates unsafe care with more than three million deaths annually; this is neither a U.S. estimate nor mortality attributed to care denial [3]. Administrative review supplies a more specific access example: among June 2024 skilled-nursing admission requests at 19 Medicare Advantage organizations, 18% of denials were appealed and 95% of those appealed denials were overturned [4]. The selected appeal population cannot establish the error rate of all denials. It does show why reasons, timely review and remedies belong alongside initial decisions.

Resources and incentives create a related problem. A 2019 U.S. review estimated annual waste at \$760–935 billion across delivery, coordination, low-value care, pricing, fraud and abuse, and administration [5]. That historical synthesis is not a current loss total or an amount recoverable by VANTHEON. CMS separately estimated FY2025 Medicare fee-for-service improper payments at \$28.83 billion (6.55%, down from 7.66% in FY2024); this payment-compliance measure includes documentation failures and underpayments and is explicitly not a fraud measure [6]. Sustainable improvement must therefore distinguish better care, reduced avoidable work, valid payment and fraud control instead of assigning them one target.

The workforce also needs a system that learns from improvement. In the AMA's participating-health-system survey, the share of responding physicians reporting at least one burnout symptom declined from 43.2% in 2024 to 41.9% in 2025; these responses are not a random national prevalence estimate [7]. Such progress is compatible with serious unmet need. The design problem is to support useful work, retain failures and successes as evidence, and give affected people a credible role in changing the conditions of care.

1.1. The institutional question

Existing approaches already address much of this problem. Learning health systems connect care with evidence generation and practice improvement; value-based care relates outcomes to resources; interoperability standards preserve exchangeable records and provenance [8, 9, 10]. Clinical AI research evaluates medical reasoning, multimodal information and long clinical contexts [11]. These traditions supply foundations, not capabilities VANTHEON can claim by association. Their joint institutional question remains concrete: who is responsible when evidence, authority, capacity, financial obligations and learning cross organizational boundaries, especially when the evidence later changes?

VANTHEON proposes a Computed Learning Health Protocol and healthcare ecosystem organized around those responsibilities. Its design thesis is to share *versioned evidence and consequences* while keeping *decision powers separate*. Patients retain consent, delegation, appeal and independent representation. Humans, institutions, software agents and robots contribute accountable services. Codex governs shared rules; Trust-Time coordinates deliverable capacity; the Pool organizes care funding, capital participation and eligible shared value; Treasury carries obligations through payment, reconciliation and remedy. THEON develops healthcare foundation models and research methods from authorized experience [12, 13, 14, 15, 16].

These functions belong together because their choices affect one another. A scheduling rule changes access and workload. A reward rule can encourage useful effort or avoidance of difficult cases. A distribution can reward contributors while weakening resources for later care. A model trained only on completed episodes may miss people whose care never began. The Provider Trust Index (PTI) is one proposed connection: better demonstrated, service-specific decisions and outcomes by human or machine providers should earn stronger incentives and greater influence on relevant proposals through separately adopted contracts. Patient rights have an independent basis.

The overarching research question is whether connecting these responsibilities improves completed care, participant control, sustainable contribution and useful learning relative to competent separate systems, without increasing exclusion or weakening accountability. The scale of the healthcare problem motivates that test; it does not establish the necessity or superiority of the proposed architecture.

1.2. Three connected cycles

Consider a primary-care visit for new concerns alongside ongoing conditions. The visit ends with a plan for tests, medicines and follow-up, but no follow-up slot is available. Completing the appointment has not completed the care. In the first cycle, *completing care*, the patient

states preferences and chooses any representative. The clinical team agrees the plan and names the responsible contributors. A coordination agent seeks feasible capacity; qualified humans, agents or robots perform their authorized tasks. The clinician remains responsible for clinical review. In the illustration, a primary-care nurse owns contacting the patient and arranging the missing next step. The unfinished obligation remains visible alongside completed work. Evidence can support a recognized payment obligation and a service-specific trust update. The system must also record people who receive no feasible offer, not only successful completions.

The second cycle is *collective improvement*. Repeated gaps can motivate a proposal to reserve results review and follow-up when tests are ordered. Affected patients and relevant providers use Codex to compare alternatives, deliberate and decide what merits evaluation. Reserving scarce capacity may improve completion while delaying others or increasing cost. A valid decision, evaluation and separately authorized staged activation connect the proposal to future practice. Monitoring can justify revision or retirement; voting alone cannot establish benefit.


The third cycle is *authorized learning*. THEON would develop reusable representations and methods for understanding longitudinal patient state, reasoning under uncertainty, coordinating work, and studying care and policy. A first model study asks whether changed weights improve source grounding. A parallel research-procedure study asks whether a changed improvement method produces better later candidates. Independent assessment and authorized promotion connect useful research to permitted use. These first tasks enable a wider foundation-model, multi-agent and recursive self-improvement program rather than define its full scope.

The Complete Institutional Event records the evidence, authority and obligations shared by these cycles. A clinical interpretation, a patient's consent, a payment instruction and a model release can refer to one episode without granting one another power. A correction returns affected claims to their responsible owners. Figure 1 locates this record within the institution. *Learn* $\rightarrow$ *Verify* $\rightarrow$ *Settle* expresses the relationship among learning, accountable decisions and consequences, rather than one fixed temporal order. The primary-care episode is an illustration. RP-01 supplies a separate eye-examination care/correction fixture, CP-01 a nonbinding vote fixture, and THEON study charters specify experiments that have not run.

Primary care that continues beyond the visit

New concerns. Ongoing conditions. A completed visit can still leave care unfinished.


Trust-Time coordinates accepted care with named owners and capacity



Complete Institutional Event

Permissions, work, results + obligations
Corrections return to affected owners

Sustain the work



Provider Trust Index (PTI)

Service-specific performance across episodes

Follow-up slot missing

Owner: primary-care nurse
Contact patient + arrange next step

Codex

Reserve result review and follow-up when ordering tests?

Patient voice + expertise
Compare, evaluate + vote
Authorize, monitor + revise

Incentives under adopted reward terms

Relevant proposal weight under voting rules

THEON Healthcare foundation models + multi-agent research




Figure 1. A conceptual primary-care episode within VANTHEON. Tests and medicines can proceed while a follow-up slot remains missing. The named nurse owns contact and coordination; the clinician retains clinical review. The Complete Institutional Event retains evidence and unfinished obligations. Pool reserves precede surplus, Treasury reconciles payment, and issuer backing stays separate. Provider Trust Index (PTI) evidence informs separately governed incentives and relevant proposal weights; patient rights remain independent. Codex can evaluate a follow-up policy. Dashed routes denote adopted policy and purpose-granted research use. THEON separates model candidates from retained research-procedure improvement, with independent evaluation and authorized release. Human, agent and robotic contributions require qualification. No clinical result or trained capability is asserted.

1.3. Contributions and present evidence

This whitepaper makes three kinds of contribution, with different evidence obligations:

1. **An institutional reference architecture.** It connects patient authority, human and machine participation, collective rulemaking, capacity, shared value and healthcare intelligence through explicit interfaces and recourse (Sections 2–9).
2. **Selected formal and executable profiles.** Restricted proofs, deterministic traces and synthetic comparisons make particular transition, trust, accounting and correction properties inspectable (Section 11 and Appendix A).
3. **A discriminating research program.** Candidate contracts and study charters specify comparisons for governance, incentives, access, integrated operation, model learning and research-procedure improvement (Appendix N and Section 9).

The labels **C1–C3** name selected cross-cutting obligations within this hierarchy. C1 propagates unfinished or corrected work through admitted dependencies; C2 separates issuer, care Pool and Treasury resources; C3 separates model and research-procedure evaluation and promotion. They neither enumerate the full vision nor establish a useful institution by themselves. Table 1 connects the architecture to current evidence and the tests still needed.

The present findings deliberately retain unfavorable results. Beyond the omitted dependency and trust undercoverage, sharing surplus raises base-case negative-equity paths from 37 to 47 per 1,000 and from 622 to 770 under the selected cost shock. The mathematical appendix establishes local properties under stated assumptions; no proof covers the full institution. The claim ledger and technical supplement bind each result to its source, method and scope.

A competent alternative is a versioned workflow database with scoped authority, provenance, explicit obligations, compensation and ordinary model-development controls. It deserves the same information, resources and evaluation. The local comparison establishes parity. Section 12 credits the inherited mechanisms and identifies the remaining integration question. A network, token or recursive procedure must demonstrate a useful additional property or tradeoff to justify its cost.

2. Participants, rights and accountable authority

The proposed institution begins with people and their rights. A patient should be able to understand a decision, control the applicable uses of their information, appoint a representative, correct a record and pursue an appeal. A contributor should be able to identify the rule under which work is recognized and an obligation is paid. Making those rights opera-

Table 1. Contribution hierarchy, evidence and discriminating tests. Proposed institutional contracts remain distinct from executed profiles and conditional proofs.

Contribution / scope	Present evidence	Next discriminating question
Architecture: rights, Codex and PTI consequences	Typed authority and candidate contracts; CP-01 executes voting arithmetic, not legitimate adoption. Sections 2, 4.	Do representation and service-relevant incentives improve proposals and access without capture or case avoidance?
Architecture: capacity and heterogeneous work	Bounded assignments and a local bundle scheduler. Sections 6, 7.	Does integrated coordination improve completed care, waiting and burden against an equal-information scheduler?
Profiles: care and correction (C1)	Reference/workflow parity on 12 workloads and 116 command attempts; both miss an omitted dependency. Section 11.	Can independent dependency discovery and recovery prevent false closure across actual adapters?
Profiles: trust and shared value (C2)	Exact restricted gate repairs five false approvals; synthetic uncertainty under-covers. Three-entity accounting exposes distribution costs. Sections 5, 8.	Are service decisions calibrated, and do adopted incentives, reserves and exit terms preserve access and payment?
Research: integration and THEON (C3)	Matched-comparison contracts and separate model/procedure charters; deterministic harness only, zero model calls. Section 9; Appendix N.6.	Does the composition add joint benefit, and do learned weights or retained research changes improve held-out tasks at matched cost?

tional requires an accountable recipient, response deadline, remedy and continuation path for unresolved care.

Identity, competence, authority and economic position are separate attributes. A clinician may also hold Pool capital; the capital does not expand clinical scope. A patient may contribute research data without becoming a capital investor. Human and machine providers can earn recognition and proposed governance influence through their demonstrated service performance. Voting eligibility attaches to the registered, accountable principal under the relevant charter; copying an agent or robot creates no additional electorate. Participants can hold several roles, but each consequential action identifies the role and principal under which it occurs.

2.1. Individual control and collective stewardship

Care authorization, research permission and economic participation use separate records. A majority vote cannot silently turn consent to a service into unrestricted permission for model

training. Refusal of optional research must preserve the rights attached to care. Patients retain explanation, representation, refusal, appeal, correction, portability and exit under the applicable individual and institutional authority [17, 18, 19].

Collective learning also creates artifacts whose control cannot be inferred from possession of a source record. The proposed model is stewardship with explicit contribution and use terms. A patient’s permission limits an authorized use; it does not by itself determine ownership of every derivative model. Conversely, an institution’s custody of a model cannot erase promised contributor rights. Table 2 separates these objects. Appendix N supplies candidate contracts and the decisions needed before adoption.

Table 2. Objects of control and benefit. These proposed rights require explicit contractual and institutional terms; the table does not assert adopted ownership or compensation.

Object	Participant purpose and responsible steward	Contract that completes the right
Care records	Access, correction, purpose-limited use and portability; record custodian	Retention, delegate recovery, revocation and unresolved obligations after exit
Contributed knowledge	Attribution and governed reuse; research/content steward	License, recognized contribution and compensation, if any
Collective models	Member-serving use and accountable control; designated model steward	Access, license, representation, revenue rights if adopted, withdrawal effects
Research outputs	Traceable contribution; research owner	Publication and reuse rights; treatment of completed derivatives
Pool participation	Eligible care, capital participation and shared benefit; Pool authority	Loss allocation, distribution, withdrawal and accrued claims
Producer work	Compensation for recognized obligations; adjudicator and Treasury	Evidence, amount, challenge, recipient allocation and recovery

Portable identity and reputation require more than exporting a score. An export must carry issuer, service, evidence and policy versions, correction lineage and uncertainty under controlled disclosure. A receiving institution verifies those claims and applies its own rules. Exit similarly separates stopping new participation from resolving outstanding care, payment and correction obligations. Control remains meaningful only if participants can recover credentials, revoke delegation and obtain usable records.

2.2. Authority at the moment of action

A grant identifies its holder, accountable owner, purpose, scope and expiry. Authentication, credentials, consent, revocation and the applicable policy determine whether it is effective (Definition A.2). The check occurs when the consequential action is attempted. An accu-

rate recommendation cannot renew expired delegation; restarting an agent cannot supply a missing permission. Emergency access requires a bounded purpose, duration, accountable holder and subsequent review.

In the running episode, a patient agent can negotiate an appointment within the patient's preferences; a clinician authorizes the service; an imaging device operates within its task qualification; and Treasury acts under a distinct financial grant. A community or research contributor can earn recognition only under the applicable contribution contract. Common attribution records connect their work without merging their powers.

The intended machine-authorization rule returns approval only when its evidence, guideline, effective-authority and service-trust conditions hold; otherwise it returns human review. Care denial remains a human decision [18]. Proposition C.2 proves that this function has the stated approval-or-review output set. It does not establish the clinical sufficiency of its predicates. The surrounding system must keep review available, preserve professional responsibility and expose the reason for an unresolved request.

These rights identify who can decide and who must respond. The next requirement is a shared account of the work, evidence and obligations to which those decisions refer.

3. A shared record with separate decision powers

The shared object must connect care, policy, finance and research without making any one of them the owner of every decision. VANTHEON assigns network execution to L0, data and semantics to L1, Codex governance to L2, attribution and trust to L3, agent execution to L4, Treasury to L5, and interfaces to L6. THEON provides learning and research capabilities across these responsibilities; it does not acquire a new authority layer [13, 20, 21, 16].

The Complete Institutional Event carries the applicable evidence, authority, obligations and contests. Closure means that the obligations and remedies within its admitted scope have authorized dispositions. Finishing an agent task, receiving a report or recording a payment can satisfy one part while the episode remains open. This makes unfinished work visible to the responsible people rather than hiding it behind a completed transaction.

3.1. Records, observations and consequences

The retained formal profile writes an event as

$$e = (R, V, A, D, X, O, C, L, S), \quad (1)$$

where R is the versioned rule, V source-signed evidence, A authority including consent and grant, D the decision, X the action and performer, O the observed outcome, C verification results, L a proposed rule update, and S settlement. Definition A.3 gives the types. That profile restricts O to a declared binary endpoint. A wider clinical record can contain multiple

observations and horizons; a registered endpoint contract must justify any reduction to the binary model.

An observation envelope binds source identity, revision, event and receipt times, units, permitted purpose and transformation lineage. An action receipt identifies the principal, effective grant, attempted operation, tool or device version and execution result. An outcome contract supplies population, endpoint, maturity window, reference standard and uncertainty. Sharing an episode identifier does not make these objects interchangeable. A successful imaging invocation can yield a receipt before a qualified interpretation or communicated follow-up exists.

The integration design includes HL7/FHIR records, X12 financial transactions, device observations, patient reports and relevant social context [17]. Adapters preserve terminology and unit mappings, missingness and derivation history. FHIR provenance is an existing representation for source and transformation relationships [10]; policy determines how those relationships qualify a decision. The architecture specifies these interfaces rather than claiming universal adapter coverage.

3.2. Replay, clocks and branching correction

For ordinary admitted events, the state transition records applied identifiers and uses fixed versioned inputs. Theorem A.7 establishes determinism and at-most-once application under its assumptions. It supports replay of the same identifier, not discovery that two different identifiers describe one physical act. External tool retries, payment receipts and corrections therefore need separate contracts.

Three clocks matter: when care occurred, when evidence arrived and when a consequence became authorized. Network finality can order a record while a clinical outcome is still immature. A recognized care liability can remain due while an outcome bonus awaits follow-up. Neither delayed evidence nor incomplete network processing is evidence that the obligation disappeared.

RP-01 implements a bounded local correction example (Appendix J). Its observation OBS-01 feeds the toy score SC-01, the policy-reassessment object PC-01 and the research dataset DS-01; the dataset in turn feeds MODEL-REF-01. These are branching dependencies, not successive stages of one pipeline. In the fixture, withdrawing the disputed contingent linkage changes the amount due from 100 to 80 fictional units and recomputes a toy posterior mean from $1/2$ to $1/3$. The dataset is reassessed and its dependent model reference receives a release-review disposition. No trained checkpoint is evaluated.

The correction is appended to the original history. Each affected object receives a responsible owner and disposition: recomputation, compensation, quarantine, reassessment or an explicitly unresolved effect. A compensating payment does not reverse physical care, and a

withdrawn dataset row does not prove that a model has forgotten it. Figure 2 separates the dependency graph from the order of authorized remedies.




Figure 2. A corrected source creates several obligations; recovering money alone does not close the episode. In the formal RP-01 fixture, (A) source observation OBS-01 parents score SC-01, policy candidate PC-01 and dataset DS-01; the dataset parents MODEL-REF-01. (B) Open-to-filled intervals show pending-to-terminal obligations on a linear logical-tick axis. Cash recovery finishes at 13; closure waits for the model-reference assessment at 14. The reference is rejected; no trained weights are evaluated. Confirmed paid history stays 100 while corrected due falls to 80. (C) C-02 reopens closure at 15; C-03 opens at 16. Their resolutions at 17 and 18 restore closure only at 18. Filled closure marks mean closed. These are logical ticks and fictional accounting units, not clinical latency or external payments.

The local reference and a conventional workflow produce identical evaluated states on the tested commands. Both also fail to discover an intentionally omitted dependency. Thus the present closure claim concerns the admitted graph; an integration must independently mea-

sure how it discovers and reconciles affected external objects. This limitation is a concrete reason to retain open obligations rather than equating internal consistency with completeness.

3.3. Open applications over common contracts

Patient, provider and institutional workspaces, developer APIs and third-party applications share projections of the authoritative state. They expose source evidence, permitted and blocked actions, outstanding obligations and the accountable next step. A patient can inspect a claim and appeal; a producer can submit scoped evidence; an independent application can propose an interpretation. The interface does not directly manufacture closure or payment authority.

An extension manifest binds application/version, accountable principal, represented person, record scopes, purposes, schema and policy compatibility, operation identifiers, receipt destination and revocation route. Stale authority, unavailable authoritative state or a schema mismatch blocks consequential execution with a typed reason. This contract permits diverse applications while maintaining common institutional semantics. Registration, interoperability and withdrawal must be exercised through deployment conformance, including recovery from partially completed external actions.

A shared account makes repeated failures visible. Deciding whether those failures justify changing a care pathway requires a collective process: who can propose a change, whose expertise and experience count, and how its effects will be evaluated. Codex supplies that proposed process.

4. Codex: governing better care

Codex is VANTHEON's proposed institution for turning experience into better rules of care. Patients, clinicians, institutions, agents and robots can identify problems and develop proposals within their participation rights. Providers with stronger demonstrated decisions and outcomes are intended to earn greater influence over proposals relevant to their expertise. Patients contribute experience, preferences and representation in their own right. Together, these roles connect learning from care to decisions about how the health system should work [12, 14].

A proposal explains what should change, who would be affected and why it should help. Codex connects evidence and alternatives to deliberation, voting, evaluation, activation and revision: both a decision and an obligation to learn whether it works. The architecture is proposed; CP-01 exercises a nonbinding voting contract.

4.1. Rules, jurisdictions and delegated power

The constitutional level establishes protected rights, responsibilities and limits on power. The policy level specifies evidence and conduct for a class of decisions. The operational level implements an adopted policy for a task and environment. Authority flows through this hierarchy; successful computation does not ratify the rule being computed.

A policy therefore has a readable statement and a bound machine representation. Its record identifies purpose, jurisdiction, eligible population, responsible authority, evidence, permitted actions, exceptions, effective interval and predecessor. Changing an executable threshold that changes clinical or economic meaning is a policy revision, even when the software change is small. Adoption must connect text and code, with conformance tests for representative and adverse cases.

Table 3 separates the powers that would otherwise be conflated in a single score or token vote. A clinician who supplies capital declares which capacity they act in and any conflict. An institution operating many agents remains their accountable principal; software replication does not create independent representation. Agents and robots can contribute evidence, simulations and proposals under defined eligibility rules.

Table 3. Four authority domains. Each requires a charter, eligible actors and recourse; authority in one domain does not confer another’s powers.

Domain	Decision object	Constraint
Clinical policy	Evidence standards, pathways and justified exceptions	Relevant expertise, affected-patient representation and professional responsibility
Network operation	Upgrades, validators and execution parameters	Consensus orders admitted transitions; it does not determine clinical truth
Economic policy	Fees, reserves, capital and distribution terms	Financial votes preserve valid care obligations and individual consent
Individual care	Consent, authorization, delegation and review	Population rules do not supply a person’s authorization

4.2. Provider trust, incentives and voting power

The Provider Trust Index (PTI) evaluates a human or machine provider’s demonstrated performance in a specified service. Better supported decisions and outcomes are intended to earn stronger economic incentives and greater service-relevant voting weight. Thus a clinician, diagnostic agent or robot can gain influence through useful healthcare work. Clinical credibility is earned through evidence; buying tokens or capital does not supply it. Section 5 separates the performance estimate from the rules that use it.

The connection has two branches: an adopted compensation or financing contract sets eligible economic terms, and a proposal's voting contract sets expertise weight. Neither silently sets the other. Machine participation follows declared identity, eligibility and representation rules; replication supplies no extra independent votes. Patients are not assigned a provider score for being patients: their representation, consent and rights remain distinct. A high imaging score supplies no automatic expertise in interpretation or financial policy.

A proposal's snapshot freezes the score definition, service domain, evidence window, uncertainty treatment, eligible principals, conflicts, weights and cutoff time. Its rule also declares patient representation, quorum, threshold, abstention treatment, veto grounds, appeal and activation authority. Missing fields block activation. Changing an electorate or its weights after observing votes requires a new snapshot and procedure. Appendix N.2 records a candidate contract; it leaves legitimacy and adoption to the appropriate constituencies.

A nonbinding human-machine example.

Suppose clinician H and interpretation agent A have qualified lower PTI bounds of 0.90 and 0.80 for the same declared interpretation service, endpoint and population. H is a registered provider principal; A acts through operator P, whose designated accountable representative casts P's vote. An illustrative service contract could attach bonuses of 10 and 5 accounting units to their respective eligible, verified contributions. A separate proposal charter could assign expertise weights of 3 and 2 to H and P on an interpretation-quality proposal. These authored terms show the two branches; they are not a calibrated PTI conversion, CP-01 execution or adopted entitlement. Payment still requires its own recognized obligation and authority.

A robotic imager R operated by P has its own acquisition-performance record. It can support compensation for distinct acquisition work and P's eligibility on a relevant acquisition proposal; it supplies no interpretation weight merely because imaging and interpretation occur in one episode. An operator eligible through several services receives the single represented entitlement defined by that proposal's charter, with a declared aggregation rule rather than summed machine votes. Replicas of A or R add neither principals nor votes, and copying a source outcome adds no independent observation or duplicate contribution payment. Patient representatives retain their separate chamber and individual patients retain consent and appeal regardless of H's, A's or R's scores. This connection is the candidate contract to evaluate, not a grant of machine personhood.

Influence, opportunity and income can reinforce one another: assignments generate evidence; evidence affects provider trust; adopted rules connect that trust to later opportunities, rewards and votes. This is a proposed incentive for better care and a feedback loop to test. A provider could instead improve recorded success by avoiding difficult cases. Evaluation must

include people offered no service, new providers seeking experience, concentrated influence and conflicting interests, alongside completed outcomes. PTI, CCS, CXS and CGS answer separate questions and require explicit contracts for each use (Section 5).

4.3. A substantive proposal: completing the service

The introductory example asks whether nonurgent eye examinations should reserve imaging, qualified interpretation and communicated follow-up together. A patient group, clinician, coordination agent or robotic-service operator could submit the problem and evidence through an eligible principal. A candidate rule would describe the bundle, responsibilities and exceptions. Alternatives include current scheduling, coordinated reservations and increased interpretation capacity. The tradeoff is explicit: reserving scarce capacity might reduce unfinished care while delaying other patients, increasing costs or leaving reserved slots unused. No alternative is presumed superior.

The proposed lifecycle has six responsibilities:

1. **Define the problem.** Register the affected population, eligible referrals, missing hand-offs, patient preferences, baseline practice and accountable owner.
2. **Compare alternatives.** Assess expected benefit, costs, subgroup harms, capacity, uncertainty and people absent from completed records.
3. **Deliberate and decide.** Disclose interests, retain dissent and freeze the electorate, patient representation and relevant PTI-based provider weights. State whether the decision supports further study, a bounded pilot or a policy version.
4. **Evaluate before wider activation.** Use simulation and, where warranted, a separately authorized pilot with a competent comparator, outcomes and stop criteria. Voting does not itself authorize a clinical study.
5. **Activate within scope.** Check evidence, rights, text/code conformance and resources; identify the authority, effective interval, staged population and fallback.
6. **Monitor, revise or retire.** Compare access, waiting, burden, costs and harms with registered expectations. Retain failures and appeal outcomes for the next proposal.

The 30-day service-completion estimand in Section 5 supplies one candidate outcome. Performed work, outcome-contingent incentives and research costs require separate terms. This scenario establishes neither an adopted pathway nor a clinical result.

4.4. A bounded voting example: CP-01

CP-01 tests a narrower question: should report-to-service linkage be checked before evidence admission? This nonbinding clinical-information fixture demonstrates voting arithmetic, sep-

arately from the complete-service proposal above. It does not authorize treatment or denial. The patient chamber contains three represented principals with one vote each. The service-expertise chamber contains three clinicians with authored PTI lower bounds 0.90, 0.80, 0.70 and fixture weights 3, 2, 1, calculated by $w = 1 + (q - 0.70)/0.10$ only at those values.

Each chamber requires participation by two of its three principals. Adoption also requires two patient yes votes and clinical yes weight of at least four of six. Abstentions count toward participation but not yes support; all denominators are frozen. Patients P1/P2 and clinicians C1/C2 voting yes produce 2/3 patient support and 5/6 clinical weight. This passes the numerical gate; it cannot outweigh a failed patient chamber or supply missing rights review, text/code conformance, clearance of active veto/appeal, or an activation grant. These authored weights are not calibrated entitlements or adopted VANTHEON governance.

An independent ethics reviewer can suspend activation on recorded rights or safety grounds. After three logical ticks, the fixture suspension moves to required review, never automatic activation. A distinct adjudicator records uphold, dismiss or remand with reasons. The machine-readable CP-01 profile retains the rule, votes and unresolved legitimacy, incentive and concentration questions. Figure 3 places this worked gate within the full policy lifecycle.

4.5. Adoption, dissent and institutional learning


For the CP-01 linkage change, evaluation would compare ambiguous reports, false exclusions and human review effort. For the broader service proposal, it would also examine whether care becomes easier to complete. The same lifecycle can govern clinical pathways, coordination methods, contribution rules and appropriately scoped research proposals; the electorate, evidence and activation authority depend on the decision domain.

A changed score or rule can alter assignments, compensation, voting eligibility and training labels. Activation must identify those dependencies and the treatment of already authorized episodes. An individual appeal supplies evidence for policy research; it does not automatically rewrite all future care or reinterpret historical entitlements.

Clinical exceptions remain legitimate objects of review. A responsible professional records the patient context, rationale, authority and follow-up for a justified deviation. Participants may contest an observation, score, policy application or procedure. Independent appeal requires notice, deadlines, an available adjudicator and an executable remedy, including compensation where owed, with a final route to external arbitration or applicable proceedings [14]. The proposed ethics veto likewise requires defined reviewer eligibility, limited grounds, notice, expiry and review [12]. Emergency suspension has bounded scope and a subsequent review obligation.

Codex turns care failures into evaluated policy

Primary-care proposal: reserve results review and follow-up when tests are ordered.



Machine evidence has an accountable principal; copies do not create extra votes.




Figure 3. Codex connects a primary-care problem to collective institutional learning. The illustrative proposal reserves results review and follow-up when tests are ordered, while the primary-care nurse remains responsible for arranging the missing follow-up. The proposal has its own named owner. Alternatives, resource tradeoffs and dissent precede separate patient and service-expertise gates. A failed chamber returns the proposal for revision. Passage is followed by evaluation and rights review, separately authorized activation, then monitoring of access, burden, costs and harms; a withheld grant leaves the change inactive. Machine evidence remains attributable to an accountable principal and does not create extra votes through copies. The inset is a different, nonbinding CP-01 fixture about report-to-service linkage. Two of three patient votes and five of six expertise weight units meet their respective thresholds of two of three and four of six. Clinician PTI lower bounds 0.90, 0.80 and 0.70 receive authored weights 3, 2 and 1. Each chamber separately requires two participating principals. The arithmetic neither establishes calibrated entitlements nor adopts the primary-care proposal.

The initial constitution must identify its ratifying constituencies, adopting procedure, authority holders, appointments, terms, replacement and independent appeal. It must also define succession from founder-led development to stakeholder governance. Authorship alone does not ratify a charter. Codex’s purpose is a durable way for participants to improve rules and obtain remedies. The next question is how evidence can identify the provider expertise and contribution on which part of that institution relies.

5. Evidence, attribution and service-specific trust

VANTHEON intends providers who demonstrate better decisions and outcomes to earn greater trust, stronger incentives and more influence on relevant better-care proposals. The Provider Trust Index (PTI) expresses that provider-performance role for clinicians, institutions, software agents and robots. It is not a score of patients’ worth or entitlement to care. The relevant question is which service a provider can perform well, for which population, under which conditions, and with how much uncertainty.

5.1. Define the quantity before using the score

PTI supplies performance evidence; an allocation rule, compensation contract or Codex voting contract determines its consequence. A stronger service-specific estimate may support a preferred assignment, more favorable eligible reward or financing terms, or greater expertise weight on a proposal. Those mappings must be adopted and evaluated separately. No universal multiplier makes a PTI increment automatically equal to money or votes. Patient representation has its own basis, even when a patient also acts in a separately qualified provider role.

Codex Clinical Score (CCS) describes clinical assessment; Codex Execution Score (CXS) describes task execution; Composite Governance Score (CGS) records permissible influence on a specific proposal [12, 14]. Each needs its own evidence and use contract. Trust propagation has established reputation-system precedents [22]; their mechanisms and assumptions do not make distinct clinical constructs interchangeable. Table 4 fixes the questions; Appendix N.1 supplies the typed candidate contract, and Table 18 identifies the boundaries that prevent misuse.

Every score version needs units, population, inputs, uncertainty, permitted uses and transition rules. A reliable laboratory-service history may inform assignment, its incentive terms and a related Codex proposal; it does not establish competence to interpret a clinical result or decide a financial dispute. Rare services, cold starts, difficult cases and recovery after error require declared policies. A new provider needs a supervised or otherwise accountable route to acquire evidence.

Table 4. Score meanings and use constraints. Mission fidelity and transparency are separately specified institutional objectives; neither is a hidden term in the PTI posterior.

Name	Question	Required qualification
PTI	Human or machine provider performance in a named service	Endpoint, risk model, uncertainty, observation window and validated use in assignment, incentives or voting
CCS	Structured assessment of clinical evidence	Endpoint, horizon, patient context, missingness and qualified review
CXS	Quality of execution in context	Separate adherence, justified deviation, coordination and outcome alignment
CGS	Permissible influence on a particular proposal	Jurisdiction, representation, conflicts, frozen weights and recourse

5.2. A bounded performance mechanism

The retained model illustrates one bounded PTI endpoint. It links a declared binary outcome to latent service quality and risk (Equation 4), then updates a time-weighted posterior (Equation 5). Its mean summarizes the fitted quantity; a lower quantile enters the approval-or-review predicate. This does not reduce all clinical judgment to that endpoint. Theorem B.2 shows how an additional successful observation changes that quantity at fixed time and model. Theorem B.3 concerns concentration of the continuous, non-decayed posterior under its stated regularity conditions. Theorem B.4 identifies ineligibility for the specified prior and threshold. These results explain local behavior; they do not establish calibration in a changing clinical population.

Numerical implementation also matters at the decision boundary. TP-01 replaces a finite-grid decision with exact integer arithmetic for one reference-risk, no-decay Beta count profile with at most 1,000 observations. It avoids five demonstrated false approvals near the threshold. Inputs outside that profile require review or a separately justified method; the correction does not certify the general risk-conditioned, decayed implementation (Appendix K).

The broader synthetic comparison finds near parity between the 199-point estimator and a conventional 1,001-point risk-adjusted Bayesian comparator. Their nominal 90% intervals cover about 85% under the matched-likelihood generator. This is a useful failure to expose before clinical use. Generalized Bayesian updating [23] can organize weighted evidence, but a chosen risk model, decay and interval still require service-specific assessment of noise, delayed outcomes, selective missingness and changing practice.

Performance also depends on who receives an opportunity. A provider who accepts only easy cases can appear more reliable without making better decisions for the eligible population. The proposed evaluation therefore records eligible referrals, offers, withheld services, refusals, abstentions, escalations, completion and missing follow-up alongside observed out-

comes. Patient risk and access enter this evaluation as context for provider performance, not as a patient PTI. Risk adjustment alone cannot recover outcomes that were never observed or establish the causal value of a provider’s contribution.

5.3. Causal effects, contribution and justified exceptions

A signature supports provenance under issuer and key assumptions. A prediction estimates a modeled outcome. Causal attribution additionally needs an estimand, comparison and identifying assumptions, including consistency, positivity and causal sufficiency where invoked [12]. Causal graphs and interventions require an identification argument beyond association [24]. Selection, confounding and interference may defeat an attribution method even when all inputs are authentic.

A proposed coordination estimand makes this obligation concrete. Consider consenting adults referred for a nonurgent eye examination at one participating site during a registered enrollment period. Let $Y(a)$ indicate completion of qualified interpretation and communicated follow-up within 30 days of referral under assignment a : a coordinator offering a reserved service bundle ($a = 1$), or conventional scheduling with the same clinical eligibility and human support ($a = 0$). The target is $\mathbb{E}[Y(1) - Y(0)]$ in that enrolled population. Randomized assignment would target the intention-to-treat effect; an observational comparison additionally requires conditional exchangeability and overlap. Both require well-defined interventions, consistent outcome ascertainment and an interference model: shared capacity, crossover and missing follow-up can invalidate a naive comparison. The site, period, safety protocol and handling of these failures must be registered before study. This proposed service-completion effect is neither an estimated clinical benefit nor a rule dividing contributor credit.

An identified effect also does not uniquely divide value among a clinician, patient, coordinator, device and institution. A contribution rule combines evidence with an explicit agreement about credit. The rule should preserve nonclinical contributions, including patient knowledge, research and community work, when the applicable contract recognizes them. High adherence is not always good clinical judgment: a justified deviation must retain rationale, accountable authorization and follow-up rather than being mechanically scored as failure.

5.4. Time and heterogeneous evidence

Humans, agents and robots generate observations at different frequencies and delays. The temporal addendum proposes an Entity-Type Temporal Profile with maximum frequency, mean interval, half-life and governed calibration [25]. For exponential decay, $\lambda = \ln(2)/t_{1/2}$

A Observation is not yet independent evidence

CONCEPTUAL

Human clinician


Assessment + later outcome

Agent

Copies of one source record

Robot

Linked telemetry + endpoint




B Declare the counting unit

TP-01 synthetic participant

40 records

20 known episode IDs




Grouping by known identity only.

Independence needs its own justification.

C Evaluate the declared gate

Separate numerical illustration



5th percentile = 0.74915 < 0.75

REVIEW

A score grants no authority.

Figure 4. Observation frequency, evidence weight and authority answer different questions. (A) Conceptual human, agent and robot streams distinguish missing outcomes, repeated source records and dependent telemetry. Time is unscaled. (B) One retained TP-01 synthetic participant contains 40 rows and 20 known episode IDs; each pair of diamonds groups into one circular ID mark. This is known-ID grouping, not a general proof of independence. (C) A separate numerical illustration uses Beta(34,6). Its lower fifth percentile is approximately 0.74915, below 0.75; equivalently $F(0.75)$ exceeds 0.05, so the exact declared gate returns REVIEW. The panels do not constitute one executed clinical pipeline. Neither a statistical estimate nor a passing numerical gate supplies clinical sufficiency or permission to act.

only when units agree. A profile must distinguish observed rate from capacity limit and identify its calibration population and version.

Figure 4 places source-specific qualification before aggregation. A hundred robot telemetry samples can describe one procedure; they are not automatically a hundred independent clinical outcomes. Normalizing values to $[0, 1]$ does not make them clinically comparable. TP-01's known-identifier deduplication raises duplicate-regime interval coverage from approximately 73% to 82%, while worsening MAE from 0.077 to 0.085 under its fixed skeptical

prior. The result exposes both the dependence problem and the estimator tradeoff. Real identity resolution, residual correlation, endpoint delay and any sampling-theorem analogy require their own assumptions and tests.

Qualified evidence can inform trust, reward and influence. Delivering the service still requires the right contributors, capacity and authority at the right time. The next sections connect these judgments to coordinated work and complete-service allocation.

6. Coordinating humans, agents and robots

Coordination turns an authorized objective into completed work. It includes preferences, planning, negotiation, reservations, tool use, handoffs and escalation. A patient representative can seek appointments; a provider agent can assemble records; an institutional agent can reconcile payment; a robot can perform a bounded physical task. Their shared protocol identifies the principal, grant and expected evidence for each action rather than assigning permissions from a role name [13, 26].

A durable work item binds objective, principal, effective grant, inputs, policy version, budget, deadline and expected observations. A planner may propose subtasks. Each consequential tool or device invocation still checks its current scope and records an attempt and receipt. What followed is a separate observation. This distinction lets a clinician recognize a missing report even when the scheduling and imaging APIs both returned success.

6.1. A complete team rather than isolated tasks

In the running episode, the coordinator must assemble an appointment, imaging, qualified interpretation and follow-up. The patient can inspect the proposed plan and decline an unacceptable slot. The selected producers receive scoped assignments; the responsible reviewer remains visible when a report is missing or ambiguous. A disputed interpretation can hold the relevant score or payment projection while coordination continues for outstanding care needs. Figure 5 traces these execution, receipt and escalation paths.

6.2. Recovery and physical accountability

An outbox/inbox records message intent and deduplicates processing. An external operation still needs a stable identity, receipt lookup and a distinction between known failure and unknown outcome. Retrying an unknown payment or physical operation without reconciliation can duplicate its effect. A remedy may be a corrected record, compensation, rescheduling or a human intervention; some acts cannot be reversed.


Robots add device identity, task qualification, physical operating limits, supervision, safe stop and recovery. Evidence of a movement or image is not evidence that a clinical endpoint

was met. Agent and device replicas retain their version and accountable principal; copying them creates neither independent evidence nor additional voting rights.

A An available visit is not a complete bundle

RP-01


Eye-examination fixture · logical ticks



Reserve B-02 only after authority, patient acceptability and trust eligibility are rechecked.

B Execution still requires authority and a known receipt

PROPOSED



Unknown receipt stays unresolved; retry queries the existing operation, not a new effect.
A valid software grant does not establish physical safety.

Figure 5. An earlier appointment can still leave care incomplete. (A) The retained RP-01 eye-examination fixture places B-01 at visit/completion ticks 2/5, B-02 at 3/6 and B-03 at 4/7. B-01 lacks interpretation; its candidate endpoint is unavailable. The reference selection rechecks authority, patient acceptability and trust eligibility, reserves B-02, and leaves B-03 as the next choice. Only visit and candidate completion times are supplied; imaging and readout marks indicate availability, not intermediate timing. (B) The proposed execution sequence reserves a stable operation identity, rechecks the current grant, and stops if authority is denied. Unknown receipts remain unresolved; retries query the same identity. A robot additionally requires an independent physical controller, limits, safe stop and escalation. These boundaries establish neither distributed exactly-once effects nor validated physical safety.

The proposed benefit is fewer incomplete handoffs and less avoidable coordination burden. It requires comparison with a competent human-assisted workflow and a strong single-agent baseline under matched access, time and resources. Measure completed obligations, escalation quality, duplicate effects, recovery from ambiguous receipts, evidence quality and hu-

man effort. Multi-agent dialogue or throughput alone cannot establish a better care process. The integrated-study candidate in Appendix N.6 makes these intended benefits measurable alongside correction and failure recovery.

Accountable execution still needs a feasible schedule. Trust-Time connects qualified contributors to the scarce, compatible capacity needed to finish the service.

7. Trust-Time: allocating a deliverable service

Trust-Time represents the proposal to reserve and coordinate trustworthy capacity across human and machine contributors [12, 27]. A capacity offer identifies a service, delivery location or channel, available interval, prerequisites, responsible principal, applicable trust estimate and cancellation terms. A request adds patient preferences, urgency, eligibility and authority to commit resources. The object being allocated is a feasible service, not an abstract score or unit of time.

The eye-examination episode makes the distinction concrete. An early imaging appointment without a qualified interpreter or follow-up capacity may be unusable. The bounded RP-01 allocation profile therefore filters visit–imaging–interpretation bundles by effective authority, patient acceptance, eligibility and available capacity. It ranks the feasible bundles by completion time using declared tie-breaks and rechecks capacity at reservation (Appendix J). A reservation remains distinct from performed work. Follow-up capacity, cancellation and reassignment are proposed integration obligations beyond this local profile; they must release or reassign capacity without recording a successful service.

In the broader design, service-relevant PTI can favor a more reliable human or machine provider among eligible choices, creating an opportunity to earn through better care. That preference belongs to an adopted allocation rule and remains distinct from the local completion-time ranking just described. Evaluating it requires the same eligible patient opportunities across alternatives, including cases a provider declines and services never offered. More favorable outcomes among accepted cases alone cannot establish a better allocation.

The broader design includes auctions and other matching mechanisms. The existing ranked-assignment experiment is not a price-discovery or incentive-compatibility result. Appendix N.4 distinguishes a candidate initial allocation rule from the economic choices still to test: private information, bids, prices, payment, false availability, strategic withholding and collusion. Concurrent reservations across institutions also require an atomicity or recovery contract.

Urgent care and access require policy before ranking. A high score or price cannot silently determine priority among incomparable clinical needs. The mechanism must specify how new producers acquire evidence, how rare specialties and underserved locations enter the feasible

set, and how patient preference, eligibility and capacity conflicts are resolved. Its value should be assessed through complete service bundles, waiting time, cancellations, access, cost and participant burden against simpler scheduling and matching baselines. Deliverable care also needs resources and credible terms for its contributors; Pool and Treasury connect those obligations to funding and shared value.

8. Pool and Treasury: funding care and sharing value

The Pool organizes care funding, capital participation and eligible surplus. Treasury reserves, pays, reconciles and corrects authorized financial obligations. Together they are intended to let members finance care, depositors supply loss-bearing capital, and producers earn for recognized work. Participation is intended to be open under the final authorized terms [18, 19, 15]. Those roles may coexist in one person, but providing capital is not performing care, and a producer need not buy a capital position to be paid.

Recognition can include clinical teams, coordinators, authorized software and robotic services, and research, development, community or public-health work under an adopted contribution contract [15, 28]. The value proposition is broader participation in the resources and benefits of healthcare. Its financial durability depends on explicit obligations, loss allocation and custody rather than on a transfer mechanism alone.

8.1. Making better care economically worthwhile

Service-relevant PTI is intended to help human and machine providers earn more favorable terms through better decisions and outcomes. A contribution contract distinguishes performed-work payment, outcome-contingent reward and trust-informed financing. It names the beneficiary, evidence, observation period, risk treatment, amount or rule, funding and recourse. These proposed uses establish no universal PTI-to-payment formula or promised return.

Performing imaging can create a valid service obligation while a separately contracted completion reward awaits interpretation and communicated follow-up. Evaluation asks whether this encourages cooperation or instead extra billable steps and avoidance of complex patients; it includes care that never starts, review costs and delayed outcomes. Codex governs the incentive within its proposed economic jurisdiction; Treasury carries out the authorized obligation. Compensation and voting weight remain separate rules.

8.2. Three owners of financial resources

The payment-token issuer holds backing against redemption liabilities. The care Pool holds assets against incurred care, liquidity requirements and designated capital risk. The network

Treasury funds protocol responsibilities and retains restrictions on development resources. A Pool holding of VVAL is a claim on the issuer; it is not an additional copy of the issuer's backing. Consolidation must eliminate that inter-entity asset and liability. Figure 6 makes the ownership and permitted flows visible.

An unpaid or disputed care claim remains an obligation according to its governing terms. A capital subscription is financing, not earned income. A Trust-Indexed Funding (TIF) advance exchanges cash for a receivable; principal repayment is not revenue, and impairment consumes capital. Restricted development resources do not become care reserves merely because both balances appear in Treasury. These distinctions determine which resources can absorb a loss or fund a payment.

8.3. A worked ledger and a distribution gate

EP-01 makes those boundaries executable with illustrative terms (Appendix L). Let C denote available Pool cash, H its VVAL holding, F the carrying value of advances, and U recognized unpaid care. The illustrative net capital position is $E = C + H + F - U$. Operating income, expenses, incurred care and impairment determine accumulated result G ; new subscriptions increase contributed capital separately. Paying an already recognized claim reduces cash and U together.

The profile permits an aggregate distribution only after three bounds: earned result, capital above a floor K , and cash above unpaid care plus buffer B :

$$D = \alpha \max\{0, \min(G, E - K, C - U - B)\}, \quad \alpha = \frac{1}{2}. \quad (2)$$

After payment, C , E and G each fall by D . Minting tokens or subscribing capital cannot create earned surplus. The gate protects the stated accounting quantities; it cannot correct omitted liabilities, mispriced assets or an inadequate capital target.


The worked trace starts with issuer backing and liabilities of 1,000, Pool capital of 300 and restricted Treasury resources of 50, in fictional dollars. The Pool converts 100 cash into VVAL, earns 120 of member income, recognizes 115 of care, pays 60 of that care, incurs 8 of administration and a 2 network fee, and advances 20. Five of the advance is impaired and 15 repaid. A subsequent 50 capital subscription leaves the operating loss of 10 intact. Redemption of the Pool's remaining 40 VVAL gives Table 5; 55 of care remains unpaid.

Sharing surplus transfers resources that otherwise remain available for future care. In EP-01's paired synthetic study, distributions increase negative-equity paths from 37 to 47 per 1,000 in the base case, and from 622 to 770 under a 35% care-cost increase in months 13–24. Temporary unavailability of 80% of cash produces due-care shortfalls on 588 versus 706 paths. These adverse results expose a distribution–resilience tradeoff under authored

A Three ledger owners retain separate obligations

EP-01


Step 9 snapshot · assets = liabilities + equity



Pool: 60 VVAL paid for care; TIF cash advance 20, principal recovered 15.

B Earned result excludes financing

Fictional USD



$120 - 115 - 10 - 5 = -10$. Positive equity does not turn the operating loss into surplus.

Figure 6. Positive equity and new capital can coexist with an operating loss. (A) The formal EP-01 worked ledger has separate issuer, Pool and Treasury owners. Step-9 snapshots reconcile assets = liabilities + equity: issuer $1,100 = 1,100 + 0$; Pool $345 = 55 + 290$; Treasury $52 = 0 + 52$, including 50 restricted development capital. The numbered journal legs are separate transactions, not simultaneous balances. Solid open arrows transfer cash; dashed arrows transfer or return VVAL. Minting 100 creates matching issuer backing and liability. Returning and burning 40 VVAL at step 12 accompanies 40 cash redemption and reduces both. Paying 60 VVAL for care changes the holder, not the issuer liability. A TIF advance of 20 becomes a receivable; recovery of 15 principal is not income. (B) The proportional waterfall is $120 - 115 - 10 - 5 = -10$ fictional USD. New capital of 50 is financing outside that earned result. The ledger is separate from reserve simulation and does not establish solvency, distribution entitlement or legal priority.

assumptions (Section 11); they neither forecast future losses nor invalidate shared value as an institutional objective.

The retained accounting propositions clarify narrower properties. Proposition D.2 conserves entries in an event-keyed payment ledger. Proposition D.3 establishes positive op-

Table 5. Closing balances of EP-01’s worked ledger, in fictional dollars. The Pool’s 350 contributed capital is reduced by a 10 operating loss. Treasury’s development restriction stays separate.

Entity	Assets	Liabilities	Net resources	Composition
VVAL issuer	1,060	1,060	0	backing/token
Care Pool	395	55	340	350 – 10
Network Treasury	52	0	52	50 + 2

erating drift under its stated model, not solvency. Proposition D.4 bounds an immediate distribution by the declared reserve buffer. None establishes sufficient capital, enforceable individual claims or real custody. Appendix N.3 makes candidate participation, loss and exit choices explicit for further evaluation.

8.4. An obligation from recognition to remedy

An entitlement proceeds through authorization or reservation, evidence, determination, challenge, release instruction, confirmed payment and reconciliation. Its record binds the economic obligation, policy/evidence versions, any score snapshot, recipient allocation, asset, amount, custody and maturity. Treasury executes the approved obligation; it does not choose a new clinical eligibility rule while transferring funds.

A stable obligation identity connects retries, splits, batches and corrections to what is authorized, reserved and discharged. Calculation, submitted instruction and confirmed movement are separate states. An unknown external response requires reconciliation before retry. RP-01 exercises these distinctions with fictional receipts, including partial discharge, challenge and recovery; it cannot establish that a custodian moved or recovered funds.

The running episode can owe payment for a performed service while a separate outcome-contingent amount awaits evidence. Correcting report attribution may change the contingent amount without erasing the service obligation. An unreleased amount can be adjusted under the governing contract; a completed transfer may instead require compensation or a recoverable claim. A promised prior challenge must precede release, with deadlines, notice and remedies for unavailable adjudicators or recipients. A ledger clawback alone cannot retrieve money.

TIF preserves early financing before final outcome evidence [29]. Its contract states purpose, recipient, amount, exposure cap, pricing if any, retained portion, repayment or reconciliation, and loss allocation. A validated trust measure may inform authorized terms; it does not turn an advance into earned compensation. Under EP-01, losses first reduce operating result and then economically erode aggregate capital. Exit is deferred unless recognized care, the buffer and capital floor remain covered. A deficit preserves unpaid

claims instead of manufacturing an offsetting token balance. Individual units, entry/exit prices, seniority, unreported claims and recapitalization remain decisions to complete.

8.5. Token roles and institutional failure

VVAL has a proposed reserve-backed, USD-redeemable payment role; the direction is for VANTHEON to develop its own permitted issuance capability. VHTH has proposed network participation, staking, fee and defined governance roles [19, 30]. Holding a payment unit differs from owning a Pool capital position, and network stake does not supply clinical competence. Neither token is presented here as an issued asset or demonstrated redemption service.

A complete VVAL issuance contract must identify eligible holders, accepted backing, custody and independent verification, mint authority, and redemption fees and timing. Redemption binds the returned or burned units to the corresponding liability reduction and confirmed backing transfer. A delayed transfer, unavailable custody and an impaired backing asset require different dispositions and recourse. This paper selects none of those operating terms and promises neither instant redemption nor an established peg. Pool capital exit remains a separate claim with its own loss-bearing terms.

The recorded VHTH allocation intention is 50% sale, 20% verification rewards, 15% foundation and 15% contributors, with contributor vesting over four years and a one-year cliff [19]. These are design parameters, not an offering or validated incentive system. Sale proceeds retain their authorized use; an allocation to Pool capital needs explicit loss exposure and accounting and is not operating surplus.

The Dai precedents make liability, backing, parameter authority, shortfall allocation and unwind explicit [31, 32]. Their collateralized-debt/soft-peg mechanisms differ from the proposed VVAL redemption promise. EP-01's emergency restriction pauses new advances, discretionary distributions and capital exit while preserving care claims and funded care payments. A candidate crisis record binds seven fields: triggering evidence and policy version; the authorized holder and affected domain; restricted actions; continuing care and financial obligations with their owners and resources; notice, independent review and appeal; expiry and review deadline; and conditions plus authority for resumption, succession or wind-down. Expiry requires review, not automatic resumption. A restriction on new exposure cannot authorize care denial, extinguish incurred claims or repurpose issuer backing. A continuity plan must identify the accountable human service path and available funding; recording a plan alone cannot guarantee continuity. Thresholds, holders, funding and enforceable priorities require separate adoption.

Scores, allocations, funding and votes can form a concentrated feedback loop. Evaluation must test strategic reporting, collusion, adverse selection, capture, correlated shocks and

delayed evidence alongside accounting. EP-01 excludes VHTH appreciation, newly issued rewards and hoped-for subscriptions as loss-absorbing resources. A calibrated model must include administrative and verification costs, delayed valid claims, false rejection, liquidity and participant behavior. The success criterion is useful care and accountable shared value sustained through adverse experience. What is learned from that work can inform future policy and intelligence, provided the people and contributions behind it retain their agreed rights.

9. THEON: learning models and improving research

THEON is VANTHEON’s proposed healthcare foundation-model family and research program. It would learn reusable healthcare representations from authorized experience and support research, planning and accountable work by people, agents and robots. VANTHEON supplies rules for data use, contribution and promotion; THEON supplies candidate models and methods. This division lets intelligence improve without giving a learner unrestricted clinical, financial or constitutional authority.

The program encompasses causal reasoning, Bayesian uncertainty, medical knowledge representation and stakeholder coordination [27, 16]. Foundation-model development and multi-agent/recursive research proceed in parallel with the unified protocol. The first two study profiles deliberately separate changing model weights from changing the procedure that finds improvements. Neither study has run; the supplied deterministic harness invokes no model.

The research path begins with understanding what evidence supports and what is missing. Grounding is an enabling task: longitudinal reasoning, causal analysis and coordination all need to distinguish a supported observation from a plausible invention. It is not the whole model-family objective. Table 6 connects the broader program to useful healthcare work and the first evidence each claim would need. The tracks can develop in parallel; their ordering does not reinstate a phase-based research deferral. In the complete-service example, THEON might compare coordination policies or investigate lost access, while Codex and the relevant authorities decide whether an evaluated proposal should be used.

These aims guide architecture and modality selection; no model internals or scale are selected by the map. Clinical tasks need appropriate professional and patient authority, while population research needs authorized data and defensible external validity. Sovereign control over artifacts, execution and recovery applies across the tracks and is evaluated separately from predictive performance.

Table 6. THEON capability-to-purpose map. Every row is a research aim; no trained capability or clinical benefit is established here. The first grounding studies specify only a subset of these obligations.

Research capability	Healthcare purpose	First appropriate evaluation
Longitudinal representations	Reconstruct patient state, changes and missing follow-up across episodes.	Temporally held-out trajectories; source grounding, missingness and subgroup error against conventional longitudinal methods.
Uncertainty and medical knowledge	Support interpretable reasoning, calibrated abstention and consistent clinical concepts.	Calibration and decision consequences; clinician-reviewed ontology relations with source provenance and ambiguity tests.
Causal and policy research	Compare plausible interventions and learn why access or outcomes differ.	Declared estimand, identifying assumptions, confounding/selection sensitivity and appropriate prospective evaluation.
Multimodal and physical work	Connect language, images and device observations to qualified human, agent and robotic tasks.	Modality alignment and task transfer; simulated then controlled physical-safety tests before patient-facing robotic use.
Coordination and multi-agent work	Help teams complete patient-accepted services and investigate failed handoffs.	Completion, access, burden and resources versus competent schedulers and single agents; correlated-error tests.
Research improvement	Produce better later models or methods from retained changes to research procedures.	Fixed/adaptive paired histories, held-out transfer, critical failures and matched total resources.

9.1. From an authorized episode to a research object

A research release begins with permitted inputs: longitudinal clinical records, language and knowledge, patient-reported context, workflow traces, device observations or robotic evidence, chosen for the task. Its manifest records provenance, inclusion/exclusion, permitted purposes, transformations, labels, missingness, population and license. Care permission is separate from model-training, secondary-research and publication permission. Contribution terms separately specify attribution or compensation [17, 16].

The first model study asks whether a model can connect claims about an episode to their supporting records and abstain when support is missing. Its concrete object is a typed source-grounded episode graph. Inputs are synthetic FHIR resources with source identifiers, times, service context and explicit missingness. Targets identify supported evidence-to-service links, contradictory or missing dependencies, correction descendants, source citations and abstention. A model that returns a plausible link with an invented source has failed the grounding

task even if its prose sounds clinically reasonable. Figure 7 follows these objects through the proposed study and promotion boundary.

Synthetic patient and template families are separated before generation; counterfactual siblings remain together. The proposed feasibility split is 6,000 training, 1,000 development and 2,000 protected assessment episodes. Generator truth supplies the task reference, not clinical truth. A separately controlled evaluator checks that the labels match the written task and keeps assessment inaccessible to candidate search. The exact generator, checkpoint, rights, budget and evaluation artifacts must be frozen before execution (Appendix M).

In the running episode, correcting the misattributed report identifies affected dataset and model-release references. The research owner determines whether to quarantine a row, rerun evaluation, retrain, replace a checkpoint or narrow its use. This lineage supports an accountable decision; it does not demonstrate automatic unlearning. Withdrawal terms distinguish future collection and use from completed derivatives and actions, including retention and remediation limits.

Authorized data can still be selectively observed. If incentives favor easy cases, difficult cases may disappear from provider evaluation and future training. A separate proposed study would link eligible opportunities, use, observed labels and later model performance, with assessment independent of care selection. Group-specific coverage, abstention, calibration and benefit must accompany aggregate scores. Neither first study tests this feedback or establishes that provenance removes selection bias.

9.2. Model learning and procedure learning are different experiments

TH-FM-01 proposes supervised low-rank adaptation of a licensed open-weight text checkpoint for the grounding task. The base checkpoint, tokenizer, prompt, tools and decoding remain fixed for the base-versus-adapter comparison; development-only work selects and freezes the training configuration. Baselines include deterministic extraction and linking, the identical frozen base, and conventional full fine-tuning within a matched compute limit where feasible. This first study evaluates standard adaptation, not a claimed new neural architecture.

TH-RP-01 instead holds the model checkpoint fixed and changes how candidate evidence-linkage methods are proposed or selected. It compares deterministic linking, a stronger graph method, a single agent, and fixed versus adaptive research procedures. Its authored synthetic graphs are partitioned into 600 development, 200 admission and 400 final-assessment episodes. Input projections exclude the completed gold graph. The small harness fixture exposes its graph for bookkeeping checks and is not a protected task benchmark.

A better grounding score can support a bounded model claim. It does not establish clinical benefit, broad foundation-model capability or recursive improvement. The larger program

THEON: healthcare understanding and research improvement

Proposed foundation-model family and research program




Figure 7. THEON is a proposed healthcare foundation-model family and research program. A primary-care episode supplies history, symptoms, observations, tasks and missingness only under purpose-specific research permission and source lineage. The capability map groups six research obligations: longitudinal and multimodal understanding; medical knowledge, uncertainty, causal and policy reasoning; and coordination, robotics and retained research-procedure improvement. It describes intended capabilities rather than selected neural-network internals. The subordinate authored FHIR example preserves an explicit DiagnosticReport R1.result reference to Observation O1 while treating the S1-to-R1 service relation as a proposed link requiring review or abstention. The two first studies change different objects: TH-FM-01 trains a standard low-rank adapter while retaining its frozen base; TH-RP-01 holds the model fixed while changing an eligible research procedure. Both require independent assessment and separately authorized release. No trained THEON capability, measured improvement or clinical benefit is shown; sovereign operation and task-specific validation remain separate obligations.

Table 7. Parallel first studies and their distinct changes. All counts and budgets are proposed feasibility parameters; neither column reports a model experiment.

Object	TH-FM-01: model learning	TH-RP-01: procedure learning
Changed artifact	Adapter weights; frozen base retained	Prompt sequence, role allocation, retrieval order or candidate-selection procedure; checkpoint frozen
Task	Source-grounded episode graph	Evidence linking and correction-descendant reassessment
Controlled comparison	Base versus adaptation, deterministic extractor and compute-matched fine-tuning where feasible	Deterministic, single-agent, fixed and adaptive research procedures
Assessment	Edge F1, unsupported links, abstention, calibration, service/template counts and costs	Paired histories, later-cycle artifacts, critical errors, abstention and all resources
Pre-run dependency	Licensed checkpoint, reviewed generator/targets, budget rehearsal, protected evaluator	Registered checkpoint, reviewed labels, protected assessment, executor and protocol freeze

therefore keeps separate study obligations for longitudinal and multimodal representation, causal estimation, uncertainty, ontologies and robotics. Causal work needs an estimand and identifying assumptions; Bayesian intervals need calibration; ontology relations need provenance and meaning; robotic learning needs task and physical-safety evidence. Architecture, modality selection and scale follow those task requirements rather than a diagram of unselected model internals.

9.3. An operational test for recursive improvement

Recursive self-improvement (RSI) concerns the improvement procedure itself. In TH-RP-01, a retained procedure change must help later cycles produce useful task improvements relative to a fixed procedure. Both arms start with matched model, data access, tools and resources. Only the adaptive arm may edit the registered candidate-generation or candidate-selection procedure. Gold labels, evaluator, thresholds, budgets, tool permissions and Codex policy remain outside that scope.

The exploratory design has five independently seeded histories and three cycles per history. Each arm receives at most 60 model calls, 120,000 input-plus-output tokens and 120 execution minutes per history, with up to three agent roles. Critics, selectors, retries, rejected candidates and human assistance are accounted for. A candidate procedure is retained only after the fixed admission check accepts its artifact, critical-error record, resources and at least 0.01 macro-edge-F1 improvement on the admission partition. This is development feedback; it authorizes another sandbox cycle, not clinical deployment.

Testing retained research-procedure improvement

TH-RP-01 · five independent paired histories; three planned cycles in each

Shared frozen model, task stream and tools




Figure 8. The proposed TH-RP-01 study repeats one fixed/adaptive procedure pair across five independently seeded histories, with three cycles per history. Both arms share a frozen model, task stream and tools; $P(t)$ denotes a possible procedure version, not an observed improvement or an accepted edit. Development-only admission requires a valid artifact, no added critical error, compliant resources and at least 0.01 macro-edge-F1 gain on the admission partition. Rejection retains the current procedure and records the candidate, failure and resources. Equal limits per arm and history are 60 model calls, 120,000 input-plus-output tokens, 120 execution minutes and up to three agent roles. All attempts, rejected candidates and human assistance count; unknown spend halts. Both arms retain every cycle’s task artifacts. After all five histories end, both arms freeze and the independent custodian evaluates the stored artifacts once on protected assessment. No protected cases, labels or final scores return to search. Five pairs support exploratory analysis only, and no experimental outcomes are shown.

Final evaluation occurs after all histories end, without feedback to search. It assesses final and prespecified earlier-cycle artifacts, critical errors, abstentions and full resource use. Figure 8 separates the mutable research procedure from evaluation custody. A five-history stage reports the individual paired outcomes and makes no confirmatory success or

nominal interval-coverage claim. Development-only pilot variability and simulated operating characteristics must justify a later confirmatory design.

A rising task score alone cannot show recursion: additional compute, accumulated examples, evaluator leakage or manual repair could explain it. The comparison must identify what procedural change persisted and whether it improves subsequent search under matched resources. Failed candidates and unstable results remain evidence. Automated agent design and self-modification already have research precedents [33, 34, 35]; THEON must establish its particular healthcare method and contribution rather than claiming novelty for recursion alone.

9.4. Multi-agent research, promotion and sovereignty

Hypothesis generation, evidence appraisal, experiment design, implementation, analysis, critique and replication are research responsibilities. They need not correspond one-to-one to separate agents. A strong single agent may carry several responsibilities; a multi-agent design must earn its additional communication and resource costs. Medical specialization and role-based reasoning have prior examples [11, 36]. Agreement among agents sharing sources or errors is not independent corroboration.

Four change classes remain separately registered: model weights; memory, tools or coordination workflows; clinical/economic policy proposals; and the research improvement procedure. A useful change in one class supplies no automatic authority or evidence for another. Each candidate binds task, population, artifact/data/tool versions, budget, failed attempts and human assistance. Independent evaluation means separate access and responsibility, not merely naming a model an evaluator.

Promotion then follows candidate, independent assessment, authorized scope, bounded integration, monitoring and fallback. Policy proposals return to Codex; economic changes to the relevant financial authority; execution changes to their runtime owner. Current acceptance rules revalidate older candidates when dependencies or permitted uses change. Useful research may continue while deployment remains unwarranted.

Sovereign operation adds testable control over training and inference, custody of keys/artifacts, inspection, export, restoration, replacement and external-provider dependencies. A proprietary checkpoint alone cannot establish independence. THEON must demonstrate recovery and operation within its claimed scope. This research program aims for useful healthcare intelligence and better research methods while retaining participant authority, contributor rights and institutional accountability through each iteration.

10. Privacy, verification and resilient operation

The proposed architecture keeps protected clinical records in access-controlled stores and uses signed references or commitments for integrity and network state. The surrounding metadata can still disclose relationships. Key custody, linkage, revocation, retention and authorized access therefore belong to the privacy design; a signature establishes neither appropriate collection nor clinical truth.

RP-01 has a narrower trust boundary: one deterministic process, admitted fixture actors, append-only local history, separate care and Treasury grants, and fictional receipts. Its tests exercise typed inputs, stale authority, conflicting identifiers, duplicate discharge and unknown-effect retries. It assumes the admitted dependency set and endpoint qualification. Appendix N.5 selects a nonbinding laboratory integration target and identifies the authentication, fault, adapter and recovery obligations needed to replace those assumptions.

10.1. Computational proofs and real-world evidence

The restricted Beta certificate in Lemma E.1 certifies a quantile inequality under its specified prior, tallies, ranges and arithmetic. It does not authenticate the underlying care or grant permission to act. The development circuit uses Groth16 over BN254 with a local setup; it is neither a production ceremony nor a post-quantum system. Deployment needs event binding, source validation, replay controls, current grants and compatible verifier execution. The proposed custom-VM and Solidity/EVM paths remain distinct network choices [30, 15].

Theorem F.1 bounds the chance of positive profit for fixed independent attempts with a declared undetected-settlement probability and collectible realized losses. It can inform sensitivity analysis of that model. Correlated sources, adaptive stopping, collusion, new identities or unenforceable penalties violate its scope. The statement supplies neither a strategic equilibrium nor a production deterrence result.

10.2. Attack the feedback paths, not only individual messages

Threats include forged or duplicated evidence, colluding sources, Sybil identities, strategic case selection, reward and policy capture, data poisoning, prompt injection, evaluator leakage and tool misuse. Allocation, scores, rewards and votes create feedback: an attacker who gains favorable evidence may gain both resources and influence over later rules. THEON must also prevent generated proposals from becoming their own acceptance labels. Scoped authority, conflicts and independent checks need adversarial evaluation across these paths.

Recovery reconciles durable state, external action receipts and financial ledgers. Restoring a database cannot reveal which external operations completed during an outage. Contracts must identify backup windows, recovery-point/time objectives, key recovery, operators and

unresolved-state handling [37]. A failed or uninitialized detector must report unavailable evidence rather than an empty reassuring result; logs must redact sensitive fields and message content consistently.

Clinical and physical work require a continuing human service path during degradation. The interface should show what is unknown, which actions remain permitted and who is responsible. Safe stop, care continuity, escalation and financial reconciliation are observable recovery behaviors, not consequences of labeling the system autonomous or decentralized.

11. What the current evidence establishes

The contribution map in Table 1 distinguishes architectural proposals, local proofs, finite executions and synthetic studies. This section asks what each experiment teaches about its mechanism. The comparisons use competent conventional alternatives and retain failures. None reports an operational healthcare economy, patient benefit or trained THEON result.

11.1. Does the reference workflow preserve its admitted obligations?

RP-01/0.2 validates types before hashing, enforces capacity timing, supports authorized terminal contest dispositions and traverses registered dependents after correction. It recomputes a toy posterior mean from $1/2$ to $1/3$ and reassesses dataset-lineage eligibility. An independent conventional workflow/event-ledger reducer receives the same admitted graph and authority information. Across twelve finite workloads and 116 paired command attempts, they make identical admission decisions and produce identical evaluated projections; neither records a specified admission or closure violation. This supports tested implementation correspondence, not equivalence on arbitrary traces.

The omission control intentionally leaves PC-01 unregistered. Both implementations then close without discovering it. Thus a correct reducer can act on an incomplete institutional model. C1 remains bounded to admitted dependencies. Discovery completeness, integration cost, recovery effort and any distributed-network benefit require separate measurements. Appendix J identifies the workload, implementation and comparison artifacts.

11.2. When do numerical and statistical trust decisions fail?

TP-01's exact integer gate repairs five demonstrated false approvals from the 199-point grid near the threshold. Its domain is a reference-risk, no-decay Beta count posterior with at most 1,000 observations. The gate uses an established incomplete-beta identity; it is not a new statistical theorem. Out-of-domain inputs require review, and the result does not validate the general decayed estimator or authorize care.

TP-01 / SYNTHETIC TRUST

Numerical precision does not resolve miscalibration


24 populations $\times$ 128 participants per regime $\cdot$ 40 recorded observations per participant

Figure 9. Finer numerical grids barely change average error, yet rounding can change a threshold decision; statistical undercoverage remains. The synthetic TP-01 study has 24 populations of 128 participants per regime, each with 40 recorded observations. A–B show mean absolute error and coverage of nominal 90% intervals; whiskers are 95% t intervals across populations. C shows 24 paired grid differences per regime (96 total), deduplicated 199-point minus conventional 1,001-point, with mean diamonds and 95% t intervals. D uses filled circles for continuous Beta lower 5% quantiles and open squares for grid values at .75; exact gates return REVIEW. Known-ID deduplication improves coverage but increases point error under the fixed prior. D uses a separate reference-risk/no-decay profile; neither clinical calibration nor estimator superiority follows.

The accompanying study generates 24 independent populations of 128 participants in each of four regimes: matched likelihood, shifted risk, misspecified risk slope and duplicated records. Each participant has 40 observations; the duplicate regime copies twenty independent episodes twice. Four methods receive identical records: the retained 199-point estimator, its known-identifier deduplicated version, a conventional 1,001-point risk-adjusted Bayesian estimator with the same prior/likelihood, and an unadjusted Beta estimator. Iden-

tifier truth is supplied; identity resolution is not tested. All methods use no decay. Latent quality follows $\text{Beta}(4, 2)$ while estimators use $\text{Beta}(2, 3)$, so “matched” refers to likelihood rather than population prior.

Figure 9 reports mean absolute error (MAE) and empirical coverage of nominal 90% intervals. The conventional risk-adjusted estimator nearly matches the retained estimator’s MAE, approximately 0.058 in the matched regime. Their coverage is approximately 85%, below nominal. Under shifted risk, the unadjusted MAE rises to approximately 0.264 versus 0.060 for the risk-adjusted methods. These results support risk adjustment under this generator, not a novel estimator or calibrated clinical interval.

Known-identifier deduplication raises duplicate-regime coverage from approximately 73% to 82% while MAE worsens from 0.077 to 0.085 under the fixed skeptical prior. Both effects matter. The reported population summaries use 95% Student- t intervals across the 24 population metrics. Participant records retain approval and false-approval counts with explicit denominators and Wilson intervals; zero counts retain uncertainty. General dependence, delay, selective missingness, changing quality and real clinical endpoints remain outside the experiment (Appendix K).

The paired-resolution analysis compares the deduplicated 199-point estimator with the conventional 1,001-point estimator on identical records. Across all 96 paired population metrics, the absolute MAE difference is below 9.36×10^{-7} . This makes their numerical parity visible in this experiment; it is not an error certificate for other inputs or the general posterior. Appendix O defines this reanalysis of retained records.

11.3. What does shared surplus cost in resilience?

EP-01 separates the issuer’s backing and redemption liabilities, Pool assets and care obligations, and Treasury resources. Its worked ledger exercises minting, redemption, cash restriction, impairment, capital participation and gated distributions. The separate stress study compares *cash-only Pool* policies; it does not simulate a circulating token economy. Figure 10 summarizes the paired stress study described below.


Four scenarios each contain 1,000 paired paths over 36 months. The conventional policy retains gains; the shared-surplus policy distributes half of eligible quarterly surplus after the same reserve and withdrawal checks. Each pair receives the same exogenous path. The retained 288,000 monthly rows support recalculation; an independent parser checks 1,728,000 accounting identities. These checks establish consistency of the records, not solvency.

Figure 10 distinguishes negative equity from an unpaid amount already due. Base-case negative-equity paths rise from 37/1,000 under retention to 47/1,000 with distributions. Under the 35% cost shock they rise from 622/1,000 to 770/1,000. Temporary cash unavailability raises due-payment-shortfall paths from 588/1,000 to 706/1,000. The artifact supplies two-

EP-01 / SYNTHETIC POOL

Surplus distribution increases modeled disruption

1,000 paired paths per policy and scenario · 36 months · All paths remain in the denominator



Monte Carlo intervals only · Model/parameter uncertainty and issuer redemption are outside this study

Figure 10. Distributing eligible surplus increases modeled shortfall frequency and duration in these synthetic scenarios. EP-01 compares 1,000 paired paths per policy/scenario over 36 months. A shows cumulative matured-care payment shortfalls with pointwise 95% Wilson intervals on common 0–100% axes; all four scenarios also give month-36 percentages, counts out of 1,000 and Wilson intervals. Shocks are care cost +35% in months 13–24, 80% cash unavailable in 12–15, and 40-unit TIF impairment at 12. B shows shared-minus-retained endpoint differences with retained 2,000-replicate paired bootstrap intervals. C shows the paired difference in months with due arrears $>10^{-8}$, including zero-shortfall paths, with 2,000 path-bootstrap draws. Intervals quantify Monte Carlo sampling, not model uncertainty. The cash-only Pool model excludes VVAL redemption and establishes neither solvency nor clinical benefit.

sided 95% Wilson intervals for frequencies and paired path-bootstrap intervals for differences. These quantify Monte Carlo variation conditional on the model, not uncertainty in actual healthcare losses.

A cash restriction can block payment with positive net assets; an advance impairment instead destroys value. Neither capital subscriptions nor issuer backing automatically re-

place lost Pool earnings. The results therefore motivate distinct capital, liquidity, exit and crisis-authority contracts. They expose a cost of the selected distribution policy, without forecasting VANTHEON finances or establishing a return (Appendix L).

The retained paths also show when payments fall behind and how long that burden persists. Under the cost shock, mean months with due-care arrears rise from 5.779 to 7.793 across all 1,000 paths, a paired increase of 2.014 months (95% paired-bootstrap interval [1.811, 2.230]). Under temporary cash unavailability, the corresponding increase is 0.377 months ([0.333, 0.421]). These counts include zero-arrears paths and count affected recorded months, not necessarily consecutive episodes. Figure 10 and Appendix O expose timing, duration and severity alongside endpoint frequencies; a balance-sheet identity alone would not reveal this care-payment burden.

11.4. When would stronger incentives produce better care?

The current experiments hold care opportunities or cash paths fixed; they do not show how providers respond to PTI-linked opportunities, rewards and voting influence. Luan, Singh and Dai demonstrate, within an analytical model, how a liability rule can change both physician use of AI and firms' investment in accuracy [38]. The relevant methodological lesson is to evaluate participants' responses, not only the rule with their behavior held fixed.

BI-01 is a proposed mechanism-study charter, not an executed result. It asks when outcome-linked rewards and PTI-based assignment increase useful care and when they encourage avoidance of difficult patients. A small model would distinguish two case-risk groups, provider acceptance and effort, and allocation and compensation rules. It would compare fixed assignment with endogenous acceptance under per-use, outcome-linked and prespecified risk-adjusted rewards, using common eligible opportunities, capacity and information. Provider participation, effort costs and valid performed-service obligations remain explicit.

The primary quantities would be completion among all eligible opportunities and worst-group patient net benefit under declared benefit and cost assumptions. Instrumentation follows eligibility, offers, patient refusal, provider withholding, abstention, escalation, completion and missing follow-up as timed transitions with explicit denominators. A higher PTI or recorded success rate accompanied by worse eligible-population completion would count against the mechanism. The charter requires payoff assumptions, response analysis, parameter ranges and failure criteria before execution; no optimal rule or clinical benefit is claimed.

Separate studies must compare PTI-informed Codex voting with alternative patient/expert rules on the same proposal evidence, measuring proposal quality, harm detection, concentration and appeal. THEON also requires a longitudinal test of selective care use, permitted training data and subsequent model performance under a fixed

independent assessment. These studies test whether earned incentives and expertise improve the learning health system, beyond correctly implementing their contracts.

11.5. The next discriminating studies

The next integrated experiment must measure useful coordination as well as correct recovery. Matched workflows should attempt complete care bundles, produce an inspectable better-care proposal and prepare a purpose-authorized research artifact. Outcomes include access and waiting, completed obligations, patient and producer effort, blinded proposal quality, correct promotion boundaries, and contribution recognition. Delays, revoked grants, newly discovered dependents, failed handoffs and unknown external effects then test the same paths under stress. Appendix N.6 specifies an equal-information comparison; it is not an executed study.

Corrections after payment or model training require discovery of affected objects, valid/false closure, correction latency, recovery completeness and human effort, with episode denominators, exclusions, tail behavior and unresolved effects. The conventional comparator can use provenance, compensation and replay; withholding them would mistake a weak baseline for an architectural advantage.

THEON's TH-FM-01 adaptation and TH-RP-01 procedure studies proceed in parallel. Their protected evaluations distinguish model gain, coordination gain and retained-procedure gain. Five exploratory research histories cannot supply a confirmatory success claim or nominal bootstrap coverage. A later frozen analysis must use development-only evidence to establish operating characteristics. Clinical benefit requires a separate claim-appropriate study.

The frozen v1.2 evidence remains in the supplement: a 24-epoch trust simulation, compound-Poisson reserve sweep and descriptive CMS analyses with their original cohorts and periods. Those generators do not calibrate the current studies, and their plots are not Figures 9–10. Full methods, hashes and adverse results remain available for reproduction.

12. Limits and the development program

The full proposition is judged by what participants can do: obtain and coordinate care, shape and challenge the rules, contribute under intelligible terms, share eligible value and retain useful learning. The current architecture connects those purposes and exposes selected mechanisms to inspection. Its local workflow comparison shows parity, its trust study shows undercoverage, and its economic study shows a distribution-related resilience cost. These are evidence about the chosen profiles, not an argument that a network or token is already necessary.

12.1. Relation to existing mechanisms

The closest-mechanism comparison in Table 8 establishes a researched integration question rather than priority for familiar components. Learning-health-system and value-based-care work already connect improvement with evidence and resources [8, 9]. Distributed state-transition systems supply relevant computational primitives [39, 40, 41, 42]. VANTHEON must demonstrate a useful healthcare property under matched assumptions and costs after those inherited capabilities are credited. The reviewed sources do not constitute an exhaustive priority or patentability search.

Chainlink’s original architecture is a particularly relevant precedent for connecting external evidence to consequential execution. It separates oracle selection, off-chain reporting and result aggregation; its proposed reputation system records assignment-specific completion, acceptance and response metrics, and its security discussion treats Sybil identities, shared sources and collusion [43]. These mechanisms preclude novelty claims for performance-informed provider selection or a feedback loop from verified work to later assignments alone. VANTHEON’s additional question concerns healthcare semantics and institutional authority: which endpoint is measured, whose rights constrain use, who corrects a clinical attribution, and which separate compensation or policy consequence follows. An oracle agreement or aggregated answer cannot by itself settle those questions. The comparison is to the 2017 design, not a claim about current Chainlink deployments or adoption by VANTHEON.

12.2. What remains a design decision

The candidate contracts in Appendix N separate research choices from adopted powers. They specify score meanings and uses, expertise with patient representation, participation and exit, capacity allocation, a laboratory deployment boundary and discriminating comparisons. CP-01’s weights are illustrative. DP-RIGHTS-01’s individual units and valuation are proposed beyond EP-01’s aggregate ledger. A named protocol is not an adopted constitution, enforceable redemption service or operating healthcare arrangement.

Consequential choices remain: institutional ratification and succession, service-specific calibration, contribution and compensation rules, custody and crisis authority, individual participation terms, unreported liabilities, cross-entity semantics and the network/VM boundary. THEON must select authorized data, licensed checkpoints, resource limits and independent assessment custody before model runs. Its broader multimodal, causal, uncertainty, ontology and robotics program needs its own studies beyond the first synthetic grounding task.

Table 8. Closest mechanisms and the remaining integration question. VANTHEON proposes a healthcare-specific composition and evaluation contract; these precedents preclude claiming novelty for its constituent techniques.

Precedent	Existing mechanism	VANTHEON obligation
Learning health systems [8]	Care-generated evidence informs practice and research.	Bind learning to patient authority, collective rule revision and accountable consequences.
FHIR provenance [10]	Versioned targets, source entities, activities and responsible agents.	Determine who reassesses dependent claims after correction and how omitted dependents are discovered.
Event sourcing; sagas [44, 45]	Reconstruct state from history; amend partial execution through compensating work.	Preserve effective policy and rights across care, payment and learning; demonstrate value beyond a competent conventional workflow.
Chainlink oracle design [43]	Select providers using service requirements and performance; aggregate external reports with explicit source and node-fault assumptions.	Define clinical endpoints, patient authority and corrections; keep evidence delivery distinct from care authorization and settlement.
Institutional and executable governance [46, 47]	Collective rule arrangements; voting snapshots, thresholds and delayed execution.	Separate expertise, patient representation, conflicts and activation; test capture and participation burden.
Forecast and decision evaluation [48, 49]	Evaluate forecast quality and decision consequences under stated losses.	Validate service-specific PTI and its decision rule under delay, selection and repeated use.
Economic mechanisms [31, 32, 50]	Financing and unwind rules; allocation with strategic participants.	Separate issuer, Pool and Treasury claims; test surplus, exit and Trust-Time against equal-information alternatives.
Model and agent research [51, 36, 33, 34, 35]	Model reporting, collaborative inference, agent-design search and editable improvement procedures.	Test healthcare weight learning and retained procedure improvement separately, with held-out evaluation and governed promotion.

12.3. Development by evidence dependency

Three workstreams can progress together. Institutional development defines and tests rights, voting, capital, compensation and recourse. Systems development implements and compares complete care, policy and research paths with explicit faults and recovery. THEON research tests weight learning and research-procedure improvement under protected assessment. Their interfaces are shared, but no stream must claim another’s result to advance.

The first integration should exercise versioned identities, grants, records and obligations through actual adapters. Independent calibration then determines which evidence and trust rules can support which service decisions. A prospective setting adds accountable clinicians, participant/data authority, success and stop criteria, and actual operational work. Research gates determine permitted scope; the retired phase schedule does not defer all advanced research by default [21, 16].

The preservation ledger retains all fifteen vision pillars, eight additional obligations and the corrected local proofs. It also retains source limits: the earlier corpus audit did not fully read two long clinical-rubric volumes, could not access four remote canonical documents and could not parse two historical PDFs. Appendix G identifies these exclusions; their contents are not assumed reconciled. Preserving a purpose keeps it visible and accountable; demonstrating it requires the contracts and studies identified here.

13. Conclusion

VANTHEON asks whether connecting care, collective rules, shared value and learning can improve their joint outcomes while people retain authority over their care. Its proposed answer is an institution with common versioned evidence and consequences, and separate decision powers. Patients retain consent, representation and recourse. Human and machine providers contribute accountable work; service-specific PTI can inform separately adopted reward and voting contracts. Codex governs collective improvement, Trust-Time coordinates capacity, the Pool funds care and eligible value sharing, and Treasury reconciles obligations. THEON supplies the healthcare foundation-model and research program.

The paper's contribution has three levels: the institutional reference architecture, selected formal and executable profiles, and a discriminating research program. The shared event makes obligations and corrections inspectable within this architecture. Current executions establish bounded workflow correspondence while exposing an omitted dependency; restricted proofs establish conditional properties. Synthetic studies expose uncertainty undercoverage and a resilience cost of distributing surplus. C1–C3 organize selected obligations, not a complete validation of the institution. THEON's proposed studies separately evaluate learned weights and retained research-procedure changes; neither trained capability nor recursive improvement has been demonstrated.

The next milestone combines a matched comparison across care, correction and policy or research action with parallel model-learning and procedure-learning studies. Incentive evaluation must include people offered no care, new providers and concentrated influence; improving recorded success by excluding difficult cases would defeat the purpose. Institutional adoption, clinical calibration and financial arrangements require their own evidence

and authority. The intended value is care people can complete, rules they can improve, contributions that are recognized and rewarded, and intelligence that learns responsibly from the experience.

Reproducibility and declarations

Availability

The accompanying versioned package contains manuscript source, vector figures, synthetic data, executable reference profiles, environment locks, proof hashes, candidate contracts and preservation crosswalks. It also records the proposed BI-01 provider-incentive study and the literacy review that motivates this revision; neither is an executed experiment. The PDF is authoritative; the source-derived HTML companion provides semantic headings, tables, figure descriptions and native MathML. The PDF is not tagged. Reproduction commands and artifact hashes are recorded in the package README and manifests; no permanent repository identifier has been assigned. Historical v1.2 and v2.4 artifacts remain frozen.

Author and interests

Alejandro S. Méndez Valentín is the sole author and founder of VANTHEON. The author reports no external funding, no additional financial interests and no additional contributors. The founder role is disclosed independently of that declaration. AI assistance supported drafting, source comparison, figure production and technical checks; the author retains scientific and editorial responsibility.

Ethics and scope

This revision collected no patient data and reports no human-participant study. Experiments use fictional or synthetic records under the stated research profiles. They do not establish clinical safety, valid consent, enforceable financial rights or deployment readiness. Subsequent studies require their own data authority, participant protections, independent evaluation and applicable review.

Rights and release

The manuscript license remains pending the author's choice. Third-party fonts retain their included licenses. This package is a technical-whitepaper publication candidate; it is not a journal submission, external publication or authorization to deploy the proposed institution.

References

- [1] World Health Organization. Universal health coverage (UHC). Fact sheet, 5 December 2025. Global service-coverage estimates for 2023 and financial-hardship estimates for 2022. [WHO fact sheet](#). Accessed 17 October 2026.
- [2] Centers for Medicare & Medicaid Services. NHE Fact Sheet: Historical NHE, 2024. U.S. national health expenditure accounts. [CMS NHE fact sheet](#). Accessed 17 October 2026.
- [3] World Health Organization. Patient safety. Fact sheet, 11 September 2023. Global evidence synthesis; underlying estimate periods vary. [WHO patient-safety fact sheet](#). Accessed 17 October 2026.
- [4] U.S. Department of Health and Human Services, Office of Inspector General. *Medicare Advantage Organizations Overturned Nearly All Appealed Prior Authorization Denials for Skilled Nursing Facility Admission, Raising Concerns About Initial Denials*. Report OEI-09-24-00331, June 2026; June 2024 requests at 19 organizations. [Original report](#).
- [5] W. H. Shrank, T. L. Rogstad and N. Parekh. Waste in the US health care system: estimated costs and potential for savings. *JAMA* 322(15):1501–1509, 2019. doi:10.1001/jama.2019.13978. Corrected online abstract checked 17 October 2026; cited waste total is a historical estimate.
- [6] Centers for Medicare & Medicaid Services. Fiscal Year 2025 Improper Payments Fact Sheet. 15 January 2026. [CMS improper-payments fact sheet](#). Accessed 17 October 2026.
- [7] American Medical Association. Physician burnout rates are falling, specialty gaps remain. 16 April 2026. Organizational Biopsy findings: nearly 19,000 physician responses across 106 participating health systems in 38 states. [Original AMA report](#). Accessed 17 October 2026.
- [8] Institute of Medicine. *The Learning Healthcare System: Workshop Summary*. National Academies Press, 2007. doi:10.17226/11903.
- [9] M. E. Porter. What is value in health care? *New England Journal of Medicine* 363:2477–2481, 2010. doi:10.1056/NEJMp1011024.
- [10] HL7 International. FHIR R4: Provenance. hl7.org/fhir/R4/provenance.html.
- [11] K. Saab, T. Tu, W.-H. Weng et al. Capabilities of Gemini models in medicine. 2024. [arXiv:2404.18416](https://arxiv.org/abs/2404.18416).
- [12] A. S. Méndez Valentín. *VANTHEON Gray Paper: the Computed Learning Health Protocol*, v1.81, 24 October 2025. Internal architectural source; source reconciliation in Appendix G.
- [13] VANTHEON. *System Architecture Specification*, cover v2.0 (90 pages), and *System Architecture Specification (SAS)*, v1.0 (65 pages). 8 January 2025. Internal specifications; conflicting closing metadata is recorded in the source manifest.
- [14] VANTHEON. *Codex Governance & Scoring Logic Specification*, v1.0. 8 January 2025. Internal specification, 57 pages.
- [15] VANTHEON. *Treasury & Settlement Protocol Specification*, v1.0. 8 January 2025. Internal specification, 105 pages.
- [16] A. Méndez Valentín. DEC-017: THEON—Healthcare Foundation Models and Research. 15 October 2026. Confirmed direction and research charter, not a model release.
- [17] VANTHEON. *Data & Integration Specification (DIS)*, v1.0. 8 January 2025. Internal specification, 73 pages.
- [18] A. Méndez. DEC-015: VANTHEON Is the New Ecosystem (Programmable Healthcare). 10 October 2026. Confirmed internal direction record.
- [19] A. Méndez. DEC-016: The Pool Is Open to Participation, and Surplus Is Shared. 13 October 2026. Confirmed direction; operating arrangements remain unresolved.
- [20] A. Méndez. DEC-002: Layer Numbering Resolution. 22 July 2026. Confirmed internal direction record.
- [21] A. Méndez. DEC-005: Unified Protocol Build—The Codex Computational Heart. 31 July 2026. Confirmed internal direction record.
- [22] S. D. Kamvar, M. T. Schlosser and H. Garcia-Molina. The EigenTrust algorithm for reputation management in P2P networks. *Proceedings of WWW*, pp. 640–651, 2003. doi:10.1145/775152.775242.
- [23] P. G. Bissiri, C. C. Holmes and S. G. Walker. A general framework for updating belief distributions. *Journal of the Royal Statistical Society, Series B* 78(5):1103–1130, 2016. doi:10.1111/rssb.12158; arXiv:1306.6430.

- [24] J. Pearl. *Causality: Models, Reasoning, and Inference*, 2nd edition. Cambridge University Press, 2009. doi:10.1017/CBO9780511803161.
- [25] VANTHEON. *SAS Addendum 1: Adaptive Trust-Time Calibration Method for Heterogeneous Care Delivery Entities*, v1.0. 24 March 2026. Internal source, 13 pages.
- [26] VANTHEON. *ENO A2A Architecture v1*. Internal architecture note, 3 pages; date not established in the retained source record.
- [27] VANTHEON. *ASCLEPION: System Design and Technical Architecture Specification*, v2.0. 18 January 2025. Historical internal source, 16 pages. Current program name: THEON, under DEC-017.
- [28] VANTHEON. *Economic Architecture & Value Capture Model*, v1.0. Internal source, 33 pages; historical economic assumptions are not adopted as current Pool forecasts.
- [29] VANTHEON. *ENO: Trust-Indexed Funding (TIF) Logic*. Internal source, 3 pages; date not established in the retained source record.
- [30] VANTHEON. *Subnet Architecture & Smart Contract Engineering Specification*, v1.0. 17 January 2025. Internal specification, 68 pages.
- [31] Maker Team. *The Dai Stablecoin System*. December 2017. [Original whitepaper](#).
- [32] MakerDAO. *The Maker Protocol: MakerDAO’s Multi-Collateral Dai (MCD) System*. White paper, 2020. Supplied 25-page edition, filename revision 021720. [Original whitepaper](#).
- [33] S. Hu, C. Lu and J. Clune. Automated design of agentic systems. *ICLR*, 2025. arXiv:2408.08435v2; first preprint 2024. Mechanism comparison uses Sections 2–3 and Appendix H.
- [34] J. Zhang, S. Hu, C. Lu, R. Lange and J. Clune. Darwin Gödel Machine: open-ended evolution of self-improving agents. 2025. arXiv:2505.22954v1; original version used for the mechanism comparison.
- [35] J. Zhang, B. Zhao, W. Yang, J. Foerster, J. Clune, M. Jiang, S. Devlin and T. Shavrina. Hyperagents. 2026. arXiv:2603.19461v1.
- [36] X. Tang, A. Zou, Z. Zhang, Z. Li, Y. Zhao, X. Zhang, A. Cohan and M. Gerstein. MedAgents: large language models as collaborators for zero-shot medical reasoning. 2023; revised 2024. arXiv:2311.10537v2.
- [37] VANTHEON. *Operational Playbook*, v1.0. January 2025. Internal specification, 101 pages.
- [38] S. Luan, S. Singh, and T. Dai, “AI Model Design and Physician Liability,” *Management Science*, forthcoming. Author accepted manuscript, October 2026. Working-paper record: doi:10.2139/ssrn.5046254.
- [39] S. Nakamoto. *Bitcoin: A Peer-to-Peer Electronic Cash System*. 2008. bitcoin.org/bitcoin.pdf.
- [40] V. Buterin. *Ethereum: A Next-Generation Smart Contract and Decentralized Application Platform*. 2014. ethereum.org/whitepaper.
- [41] G. Wood. *Ethereum: A Secure Decentralised Generalised Transaction Ledger* (Yellow Paper). 2014–. ethereum.github.io/yellowpaper/paper.pdf.
- [42] Team Rocket, M. Yin, K. Sekniqi, R. van Renesse and E. G. Sirer. Scalable and probabilistic leaderless BFT consensus through metastability. 2019. arXiv:1906.08936.
- [43] S. Ellis, A. Juels and S. Nazarov. *ChainLink: A Decentralized Oracle Network*. Version 1.0, 4 September 2017. Mechanism comparison uses Sections 2.1, 5.2 and 5.3. [Original whitepaper](#).
- [44] M. Fowler. Event Sourcing. 12 December 2005. Author’s pattern description; particularly Event Replay, External Updates and Code Changes. martinfowler.com/eaDev/EventSourcing.html. Accessed 15 October 2026.
- [45] H. Garcia-Molina and K. Salem. Sagas. *Proceedings of ACM SIGMOD*, pp. 249–259, 1987. doi:10.1145/38713.38742; original paper, university-hosted copy.
- [46] E. Ostrom. *Governing the Commons: The Evolution of Institutions for Collective Action*. Cambridge University Press, 1990. doi:10.1017/CBO9780511807763.
- [47] OpenZeppelin. How to set up on-chain governance. *Contracts 5.x documentation*. Sections Setup and Timelock. docs.openzeppelin.com/contracts/5.x/governance. Accessed 15 October 2026; documentation comparison, not a selected VANTHEON dependency.
- [48] T. Gneiting and A. E. Raftery. Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association* 102(477):359–378, 2007. doi:10.1198/016214506000001437.
- [49] A. J. Vickers and E. B. Elkin. Decision curve analysis: a novel method for evaluating prediction

- models. *Medical Decision Making* 26(6):565–574, 2006. doi:10.1177/0272989X06295361.
- [50] A. E. Roth. The economist as engineer: game theory, experimentation, and computation as tools for design economics. *Econometrica* 70(4):1341–1378, 2002. doi:10.1111/1468-0262.00335.
- [51] M. Mitchell, S. Wu, A. Zaldivar, P. Barnes, L. Vasserman, B. Hutchinson, E. Spitzer, I. D. Raji and T. Gebru. Model cards for model reporting. *Proceedings of FAT**, pp. 220–229, 2019. doi:10.1145/3287560.3287596; arXiv:1810.03993v2.
- [52] A. W. van der Vaart. *Asymptotic Statistics*. Cambridge University Press, 1998, Chapter 10. doi:10.1017/CBO9780511802256.
- [53] NIST Digital Library of Mathematical Functions. §8.17, Incomplete Beta Functions, equation 8.17.5. dlmf.nist.gov/8.17.E5. Accessed 15 October 2026.
- [54] The MITRE Corporation and contributors. *Synthea: Synthetic Patient Population Simulator*. Official repository and FHIR export documentation. github.com/synthetichealth/synthea. Accessed 15 October 2026.
- [55] E. J. Hu, Y. Shen, P. Wallis, Z. Allen-Zhu, Y. Li, S. Wang, L. Wang and W. Chen. LoRA: low-rank adaptation of large language models. 2021. arXiv:2106.09685.

A. Formal Event and Transition Model

The following definitions and conditional results retain the corrected technical core of whitepaper v1.2. They formalize selected event, trust, authorization and accounting mechanisms. They do not constitute a proof of the complete governance, care-delivery, financial or learning system. In particular, the binary endpoint below is one research profile; it is not a universal representation of clinical benefit. The broader architecture is described in Section 3.

A.1. Participants, authority and event commitment

Let $\mathcal{A}$ denote the participant set and $\mathcal{K}$ the service set. The consent registry K is distinct from $\mathcal{K}$.

Definition A.1 (Participant).

A participant is a tuple $a = (id, class, cred, key)$ with $class \in \{\text{clinician, organisation, agent, robot, patient, pool, steward}\}$, verifiable credentials $cred$ (e.g. licence and NPI for clinicians, device certification for robots, model and version attestation for agents) and a public key key .

Definition A.2 (Grant).

A grant is $\gamma = (holder, purpose, scope, owner, t_{exp})$. An action x by holder a at time t is *authorized* under the proposed policy only if the grant is authenticated, unrevoked, issued by an entitled authority, matches the holder and purpose, covers x , and has $t < t_{exp}$. The accountable owner must hold the required authority and credentials; applicable patient consent and service restrictions must also hold at the policy-defined decision point.

Definition A.3 (Verified event).

A completed verified event is a tuple $e = (R, V, A, D, X, O, C, L, S)$: $R \in \mathcal{R}$ is the versioned rule that applied; $V = \{(v_j, src_j, t_j, sig_j)\}$ the evidence, each value signed by its source; $A = (\kappa, \gamma, role)$ the authority, including consent; $D \in \{\text{APPROVE, REVIEW, DENY}\}$ the decision and X the action, with its performing participant a_X ; $O \in \{0, 1\}$ the observed outcome; $C = (c_1, \dots, c_m)$ the verification results; L the proposed rule update, if any; S the settlement.

Stable identity and commitment

[Specified] Each care event has a stable identifier $\iota(e)$ assigned by its authorized event owner. Its immutable pre-settlement body $B(e)$ contains that identifier, the event/rule versions, service and event time, and the closed values of (R, V, A, D, X, O, C, L) . It excludes S and all receipts derived from the body's commitment. We write $H(e)$ as shorthand for

$H(\text{canon}(B(e)))$, rather than a hash of a tuple containing its own settlement receipt. A settlement receipt records $(\iota(e), H(e), a_X, v_C(e))$. Corrections have distinct identifiers, reference the prior identifier and digest, and follow an explicit correction/contest transition; they cannot be counted as another independent success or payment through ordinary ingestion. The authority registry preserves versioned effective-time history; access and new actions must also satisfy current consent and policy. This identity rule prevents replay of a recorded identifier; detecting two different identifiers for the same real-world care is a separate evidence and identity-resolution obligation.

A.2. Network and adversary

The protocol model assumes (A1) an append-only event log and an anchor mechanism satisfying its stated safety/finality conditions; (A2) at least one outcome-attestation path independent of the participant whose trust it affects; (A3) the specified outcome model and the statistical conditions of the particular estimator theorem; and (A4) collision resistance and signature unforgeability under the threat model of the selected primitives. These are assumptions to instantiate and validate, not properties established by naming a ledger, signing a record or running the simulation. In particular, (A2) does not imply mutual independence among every check or independence across repeated attempts. The stronger independence assumptions used in the fraud calculation are stated separately.

The modeled adversary may control multiple participant identities, select patients, choose documentation and attempt re-entry. It cannot forge evidence from an uncompromised external signer or break (A1)/(A4). Signed inaccurate evidence, shared source failures, collusion and compromised issuers remain outside those exclusions and require additional controls/evaluation. The current local proof prototype uses classical Groth16; post-quantum signature migration and a post-quantum proof construction are separate specified/research work, not guarantees of the present artifact.

A.3. Protocol state and transition

Definition A.4 (Protocol state).

$\sigma = (\mathcal{W}, \mathcal{T}, \Pi, \mathcal{P}, \mathcal{J})$, where $\mathcal{W} = (\mathcal{R}, \Gamma, K, \Theta)$ holds rules, grants, consents and parameters; $\mathcal{T} : \mathcal{A} \times \mathcal{K} \rightarrow \Delta[0, 1]$ maps each participant a and service k to a posterior over true quality; Π is the settlement balance; $\mathcal{P}$ is the event-keyed attribution ledger; and $\mathcal{J}$ records finalized epochs, input digests and receipts. The Pool experiment separately tracks incurred obligations and net assets (Appendix D.3).

Here “quality” denotes the model’s latent success probability for the named service endpoint at reference risk; it is not an all-purpose measure of clinical quality. The mathematical distribution $\mathcal{T}$ is distinct from a finite-grid numerical representation.

Definition A.5 (State-transition function).

For a new epoch t , let $\mathcal{E}_t^{\text{new}}$ be the admissible new-event set defined below. With contribution input C_t and the canonical pre-epoch governance queue fixed, the specified transition is

$$\begin{aligned}\tilde{\sigma}_t &= \Phi(\Delta_\lambda(\sigma_t), \mathcal{E}_t^{\text{new}}), \\ \sigma_{t+1} &= \Upsilon_t(\sigma_t, \mathcal{E}_t, C_t) \\ &= \text{Commit}_t[\Omega(\Lambda_{C_t}(\tilde{\sigma}_t))].\end{aligned}\tag{3}$$

Here Δ_λ decays accumulated log-likelihood evidence, leaving the prior fixed; Φ updates trust (Def. B.1); Λ_{C_t} applies contributions, fees and payment entries (Def. D.1); and Ω applies the canonical steward-approved queue. The outer operation commits those updates and the epoch record atomically. Matching retries return the recorded receipt without changing the current state; conflicting inputs are rejected without state change.

Definition A.6 (Verification predicate).

$\mathbb{V}(e, \sigma) = \bigwedge_{j=1}^4 v_j(e, \sigma)$: v_1 , R is satisfied by V under its computable definition (CQL); v_2 , every evidence element required by R is present and source-signed; v_3 , A is valid at t_e (consent unrevoked, grant in scope and unexpired); v_4 , the outcome O is attested through a path independent of the performing participant (A2).

Epoch and event identity

[Specified] Every closed event has a stable event identifier $\iota(e)$ and an immutable content digest $H(e)$. Corrections reference the earlier identifier and follow an explicit correction/contest transition; they are not silently introduced as another independent success or payment. Extend the state with bookkeeping $\mathcal{J}$ recording finalized epochs and a digest binding the canonical event batch, contribution input, rule/authority snapshot and governance queue. The attribution ledger $\mathcal{P}$ is a finite map keyed by event identifier, including zero-value entries. Within a batch, identical repeated identifiers are deduplicated; conflicting contents for one identifier reject the batch without a state change. An identifier already in $\mathcal{P}_t$ must match its stored digest; otherwise the batch is rejected. An epoch is applied at most once. A matching retry returns the existing receipt without changing the current state or reapplying decay, contributions, trust or settlement; a conflicting retry is rejected. New epochs are processed in canonical order.

For a new epoch let

$$\mathcal{D}_t = \{e \in \text{dedup}(\mathcal{E}_t) : \text{ordinary}(e)\}.$$

The new-event set is

$$\mathcal{E}_t^{\text{new}} = \{e \in \mathcal{D}_t : \iota(e) \notin \text{dom}(\mathcal{P}_t), \\ \mathbb{V}(e, \sigma_t) = 1\}.$$

Here $\text{ordinary}(e)$ excludes correction and contest records, which require their separately specified transition. Use this same set for trust update and settlement, evaluate verification against the frozen pre-epoch rule/authority snapshot, and record the epoch atomically with both updates. Steward changes use a fixed canonical queue order independent of the input enumeration.

Theorem A.7 (Determinism and at-most-once application). *For fixed pre-state, epoch, contribution input and canonical steward queue, the specified transition yields the same post-state for any enumeration of the same admissible event batch. A finalized epoch or already credited event cannot cause a second trust or settlement update through this transition.*

Proof. Canonical identity and deduplication yield the same new-event set. Verification uses the same frozen inputs. In the mathematical model the likelihood product and settlement sums commute. The decay step is applied once, and the governance queue has the same fixed order. The atomic epoch record prevents repeat application; the domain check on $\mathcal{P}_t$ excludes previously applied events from later batches. $\square$

This is a specification result over the stated arithmetic. Executable implementations additionally require canonical serialization, versioned inputs, a specified numerical representation/reduction order and durable transaction semantics. The simulation checks do not establish those deployment properties.

B. Service-Specific Trust

The results in this appendix concern a named binary endpoint under an explicit probabilistic model. They do not identify causal contribution, calibrate a governance voting weight or establish that the endpoint is clinically sufficient. Section 5 places this profile within the wider evidence and score architecture.

B.1. Outcome model and posterior

For the research model, let $r \in (0, 1)$ be a service-specific risk covariate and $q \in (0, 1)$ the latent success probability at reference risk for a named binary endpoint. The illustration does not substitute CMS-HCC payment-risk scores for a validated clinical risk model. For service k ,

$$\pi(q, r) = \sigma(\text{logit } q - \beta_r(r - \bar{r})), \quad (4)$$

with σ the logistic function, $\bar{r}$ the mean risk and $\beta_r > 0$.

Definition B.1 (Trust index, VIP-0001).

For participant a and service k with verified events $\mathcal{E}_{a,k}$,

$$\begin{aligned} \mathcal{T}_{a,k}(q) &\propto \text{Beta}(q; \alpha_0, \beta_0) \\ &\times \prod_{e \in \mathcal{E}_{a,k}} [\pi(q, r_e)^{O_e} (1 - \pi(q, r_e))^{1-O_e}]^{w_e}, \end{aligned} \quad (5)$$

with $w_e = e^{-\lambda(t-t_e)}$. For $\lambda > 0$ this is a power-likelihood update [23]; its interval probabilities are model-conditioned and do not automatically have nominal repeated-sample coverage. The index and its lower credible bound are $\text{PTI}_{a,k} = \mathbb{E}_{\mathcal{T}_{a,k}}[q]$ and $\text{PTI}_{a,k}^- = F_{\mathcal{T}_{a,k}}^{-1}(\delta)$.

B.2. Conditional properties

Theorem B.2 (Monotonicity at fixed time and model). *At the same evaluation time, with prior evidence and model parameters fixed, adding one verified success with positive weight never decreases $\text{PTI}_{a,k}$ or $\text{PTI}_{a,k}^-$; adding one failure never increases them.*

Proof. For a fixed risk r , $\pi(q, r)$ is strictly increasing in q . The success update reweights the posterior by the increasing function $\pi(q, r)^w$, $w > 0$, producing monotone likelihood-ratio ordering and therefore first-order stochastic dominance. Its mean and quantiles weakly increase. Reweighting by $(1 - \pi(q, r))^w$ gives the reverse order for failure. $\square$

The result does not compare different epochs, changed risk models or changed policy parameters; decay can change a score even without a new event.

Theorem B.3 (Concentration of the continuous posterior). *Let the true quality $q_0 \in (0, 1)$ be fixed, set $\lambda = 0$, and suppose the risks are i.i.d. and outcomes are conditionally independent Bernoulli draws from equation (4). The risk slope and reference risk are known and fixed, the model is correctly specified, and the prior has a continuous positive density near q_0 . Under the usual local regularity conditions for the one-parameter model, the posterior concentrates at q_0 ; the width of its central $(1 - 2\delta)$ interval is $\Theta_P(n^{-1/2})$ for fixed $0 < \delta < 1/2$.*

Proof sketch. The model is a regular Bernoulli regression in logit q , with per-event Fisher information $I(q) = \mathbb{E}_r[\pi(q, r)(1 - \pi(q, r))]/[q^2(1 - q)^2] > 0$. The Bernstein–von Mises approximation [52] gives posterior variance $(nI(q_0))^{-1}$ and central width $2z_{1-\delta}(nI(q_0))^{-1/2}(1 + o_P(1))$. $\square$

This theorem concerns the continuous, correctly specified, non-decayed model. The simulation uses a finite grid and can only evaluate its numerical approximation at the reported sample sizes. With decay and mean monthly event count $\bar{m}$, expected total discounted event weight is bounded by $\bar{m}/(1 - e^{-\lambda})$; this is neither a deterministic bound for Poisson counts

nor automatically an information-based effective sample size. Estimation of risk parameters, misspecification, selective or missing outcomes and repeated decision-making require separate evaluation.

Theorem B.4 (Cold-start ineligibility). *If $F_{\text{Beta}(\alpha_0, \beta_0)}^{-1}(\delta) < \theta_k$ for every service k , a newly initialized identity is ineligible for instant authorization before accumulating verified evidence.*

Proof. With no evidence the posterior equals the prior, so the trust predicate a_4 in Rule C.1 is false. $\square$

With $(\alpha_0, \beta_0, \delta) = (2, 3, 0.05)$ the prior lower bound is 0.098 against $\theta = 0.75$. A new identity is therefore ineligible without verified evidence. This does not imply that re-entry always lowers a score: an adverse history can have a posterior below the prior. Credential-linked history and reset controls are separate protocol requirements. Cold-start ineligibility also does not guarantee that an under-threshold participant never crosses the threshold after observing data.

C. Approval-or-Review Rule

This is the retained decision-function profile. It addresses the function's output set; operational authorization and write permissions are separate requirements.

Definition C.1 (Authorization rule).

For request e by participant a for service k :

$$\text{Auth}(e, \sigma) = \begin{cases} \text{APPROVE} & \text{if } a_1 \wedge a_2 \wedge a_3 \wedge a_4, \\ \text{REVIEW} & \text{otherwise,} \end{cases}$$

with a_1 : R satisfied by V ; a_2 : required evidence complete with validated provenance/signatures under the frozen policy; a_3 : holder, role, purpose, scope, credentials and consent valid at the action decision time; a_4 : $\text{PTI}_{a,k}^- \geq \theta_k$.

Proposition C.2 (Approval-or-review output). $\text{DENY} \notin \text{im}(\text{Auth})$ for the function in Definition C.1.

Proof. Its two branches return only APPROVE or REVIEW. $\square$

The wider policy requires an identified, qualified human acting under the applicable role and grant to record any denial with signed attribution. The two-branch expression alone does not enforce runtime write permissions or establish compliance with a particular coverage rule. Those are integration requirements [Specified]. A proposed fraction $\rho_{a,k} = \rho_0(1 -$

$\text{PTI}_{a,k}^-$), $0 \leq \rho_0 \leq 1$, is sampled for retrospective audit. An adverse audit follows the event's correction/contest contract, retaining provenance and avoiding duplicate counting. The fixed-time monotonicity theorem does not prove how an arbitrary correction, decay step or policy change affects a later score.

D. Settlement and Reserve Accounting

These accounting statements apply to the specified transition and bounded Pool model. They do not prove the sufficiency of capital, the finality of external payment rails or a sustainable depositor return. Section 8 describes the broader institutional and economic design.

D.1. Settlement on a new event

Definition D.1 (Settlement on a new verified event).

For a contract $\mathcal{C}$ mapping admissible verified events to nonnegative values $v_{\mathcal{C}}(e)$, let $\mathcal{E}_t^{\text{new}}$ be the new-event set defined for the epoch transition. The payment ledger $\mathcal{P}$ is a finite map keyed by stable event identifier, including zero-value entries. The settlement stage applies

$$\Pi' = \Pi + C_t(1 - f) - \sum_{e \in \mathcal{E}_t^{\text{new}}} v_{\mathcal{C}}(e),$$

and inserts, for each such event,

$$\mathcal{P}'[\iota(e)] = (H(e), a_{X,e}, v_{\mathcal{C}}(e)).$$

All earlier entries are retained. Contributions, trust updates, new payment entries and the finalized-epoch record are committed atomically under the specified transition; retries do not repeat them.

Proposition D.2 (Payment-entry conservation). *In a fresh admissible epoch, total payment debits from Π equal the total value of newly inserted entries in $\mathcal{P}$. Each new entry references one new event identifier and its immutable body commitment.*

Proof. New keys are disjoint from $\text{dom}(\mathcal{P}_t)$, and each appears once. Both payment debit and newly inserted value sum the same $v_{\mathcal{C}}(e)$ over the same set $\mathcal{E}_t^{\text{new}}$. Contributions and fees are separately accounted inflows/outflows. $\square$

This is an accounting property of the specification. It does not establish adequate reserves, clinical validity, legal authority or finality of a transfer on an external payment rail. For work attributed to an agent or robot, its accountable owner holds the contractual financial claim.

D.2. Reserve dynamics

Let legitimate incurred care cost be compound Poisson, $L_t = \sum_{j \leq N_t} Y_j$, with $N_t \sim \text{Poisson}(\nu)$ and $\mathbb{E}[Y] = \mu$. Assume $\nu > 0$, nonnegative identically distributed severities independent of the claim count, and $0 < \mu < \infty$; take $\eta \geq 0$, $0 \leq f < 1$ and $u \geq 0$. These are the domain assumptions for the following drift equivalence; if $\nu = 0$, the expected increment is zero. Contributions are $C = (1 + \eta)\nu\mu$ and the protocol fee is fC . In a deliberately simplified immediate-payment model, let uL_t be an additional payment loss, so the retained net-asset process is $R_{t+1} = R_t + C(1 - f) - (1 + u)L_t$. Here u is an assumed loss multiplier relative to legitimate care cost, not a fraction of total payments or an estimate from CMS data.

Proposition D.3 (Positive operating drift). *Before distributions, with fixed parameters and no additional expenses, $\mathbb{E}[R_{t+1} - R_t] = \nu\mu[(1 + \eta)(1 - f) - (1 + u)]$. The expected increment is positive exactly when $(1 + \eta)(1 - f) > 1 + u$.*

Proof. Take expectations using $\mathbb{E}[L_t] = \nu\mu$. □

With $\eta = 0.10$ and $f = 0.02$, the expected increment is 7.8% of expected legitimate care cost at $u = 0$, and -2.2% at $u = 0.10$. Positive drift does not prevent a finite reserve from being exhausted. Legitimate care costs, correlated losses, pricing error, payment delays, expenses and distributions also determine financial risk.

D.3. Participation and reserve-first distributions

The proposed Pool has three economic roles: members contribute for covered care; depositors provide capital and participate in surplus; and producers receive payment for verified work. Providing care does not itself require a capital contribution. The same person or organisation may have separately documented roles. Participation terms, loss allocation, entry, withdrawal, seniority and the legal form remain design work [Specified].

A bounded accounting model

To evaluate the distribution rule without inventing those terms, we assume a fixed cohort and opening capital committed to absorb losses over a finite horizon. There are no new subscriptions, withdrawals, guaranteed principal, investment returns or transfers from stablecoin backing. These are experiment assumptions, not depositor rights. Opening net assets do not exceed the fixed reserve target R^* . Let A_t be available Pool assets and Q_t incurred but unpaid covered-care obligations at the start of epoch t . Incurred cost L_t is recognised even when documentation or payment is delayed. Verified care payments P_t satisfy $0 \leq P_t \leq Q_t + L_t$; $E_t \geq 0$ represents other operating costs, excluding the fee already

deducted from contributions. The illustrative additional payment loss is $U_t = uL_t$, and C_t denotes member contributions in epoch t . Then

$$\begin{aligned} Q_{t+1} &= Q_t + L_t - P_t, \\ A_{t+1}^- &= A_t + C_t(1 - f) - P_t - U_t - E_t, \\ B_{t+1} &= A_{t+1}^- - Q_{t+1}. \end{aligned} \tag{6}$$

The process stops at the first $B_{t+1} < 0$, records the shortfall, and makes no subsequent distributions. Otherwise, for a fixed payout fraction $\chi \in [0, 1]$,

$$\begin{aligned} \Sigma_t &= \chi [B_{t+1} - R^*]^+, \\ A_{t+1} &= A_{t+1}^- - \Sigma_t, \\ R_{t+1} &= A_{t+1} - Q_{t+1} = B_{t+1} - \Sigma_t. \end{aligned} \tag{7}$$

The reserve target is a buffer above recognised unpaid obligations. For this experiment $P_t = Q_t + L_t$, $Q_0 = 0$, $E_t = 0$ and $A_0 = R^*$; $\chi = 0$ retains operating gains, while $\chi = 1$ distributes all assets above the target after obligations are recognised. The target is expressed in months of gross contributions, not an actuarial capital requirement. Real deployment also needs claim estimation, liquidity, asset restrictions and an explicit capital contract.

Proposition D.4 (Reserve-first accounting). *On an active epoch, the policy distributes nothing when $B_{t+1} \leq R^*$. When $B_{t+1} > R^*$, post-distribution net assets are at least R^* . The policy creates neither income nor a guarantee against future shortfalls.*

Proof. For $B_{t+1} > R^*$, $R_{t+1} = R^* + (1 - \chi)(B_{t+1} - R^*) \geq R^*$; otherwise $\Sigma_t = 0$. □

A proposed sharing parameter ω assigns $\omega\Sigma_t$ to members and $(1 - \omega)\Sigma_t$ to depositors. The experiment treats both as cash outflows in that epoch; a future contribution credit would instead remain a liability until used. It does not estimate a depositor return. In particular, $\mathbb{E}[X^+] \geq (\mathbb{E}X)^+$: positive distributions in some epochs do not establish positive drift or sustainable long-run yield.

In words. Pay covered obligations, keep the reserve, then share an eligible remainder. Verification can reduce some losses; it cannot make uncertain care costs or capital losses disappear.

E. Restricted Integer Threshold Certificate

The lemma concerns one restricted posterior family. Its arithmetic predicate is distinct from evidence authentication and clinical validity. The broader verification boundary is discussed in Section 10.

Statement and profile

The research circuit proves a threshold certificate for the unadjusted, non-decayed posterior $Q \sim \text{Beta}(\alpha_0 + s, \beta_0 + f)$, where s, f are committed integer counts. It does not prove the full risk-adjusted posterior of Definition B.1 or the correctness of the counts' derivation from clinical events. Public inputs are the tally commitment, its Merkle root, the epoch, and a fixed prior/threshold profile; the private witness contains the counts, salt and inclusion path.

Lemma E.1 (Integer certificate for an exact posterior quantile). *Let $s, f \in \mathbb{Z}_{\geq 0}$, $\alpha_0, \beta_0 \in \mathbb{Z}_{>0}$, $a = \alpha_0 + s$, $b = \beta_0 + f$ and $n = a + b$. Let $\theta = h/d$ with integers $0 \leq h < d$, $d > 0$, and $\delta = u/v$ with integers $0 < u < v$. Set $A = ad - hn$. If*

$$A > 0, \quad u(n+1)A^2 \geq (v-u)ab d^2,$$

then $F_Q^{-1}(\delta) \geq \theta$.

Proof. Write $\mu = a/n$ and $\varsigma^2 = ab/[n^2(n+1)]$. The positive gap is $x = \mu - \theta = A/(nd) > 0$. The one-sided variance inequality gives

$$\begin{aligned} \Pr(Q \leq \theta) &\leq \frac{\varsigma^2}{\varsigma^2 + x^2} \\ &= \frac{ab d^2}{ab d^2 + (n+1)A^2} \leq \frac{u}{v}. \end{aligned}$$

For completeness, let $Z = \mu - Q$, so $\mathbb{E}Z = 0$. For any $c \geq 0$, the event $Z \geq x$ implies $(Z + c)^2 \geq (x + c)^2$; Markov's inequality yields $(\varsigma^2 + c^2)/(x + c)^2$. Choosing $c = \varsigma^2/x$ gives the stated bound. The Beta CDF is continuous and strictly increasing on $(0, 1)$, so $F_Q(\theta) \leq \delta$ implies $F_Q^{-1}(\delta) \geq \theta$. $\square$

Bounded research implementation

The demonstration fixes $(\alpha_0, \beta_0, h, d, u, v) = (2, 3, 75, 100, 1, 20)$ and $0 \leq s, f < 2^{20}$. Thus its predicate is $A = 100(2 + s) - 75(5 + s + f) > 0$ and $(n+1)A^2 \geq 19ab 100^2$. The circuit constrains the count domains, the fixed public parameters, the epoch encoding and the nonnegative integer representation of A ; there is no freely chosen margin. Intermediate bounds keep the integer comparisons below the scalar-field modulus. The certificate can reject some eligible posteriors; the measured acceptance and recall describe that conservatism, not an alternative definition of PTI.

Authentication and production limits

Inclusion is proved relative to the supplied root. The verifier still needs an authenticated root, an accepted profile and epoch, participant/service binding, accountable tally construction, and replay policy. A self-selected tree does not establish clinical evidence. The development

ceremony has a single local contributor, and its fixture tallies, salt and witness are retained for inspection; it provides no production setup assurance or private patient-data demonstration.

Groth16 is classical, not post-quantum. Transparent hash-based proof systems remain a research direction, requiring a separate implementation, security assumptions and performance study. They are not a drop-in verified backend. Proving the full risk-adjusted, time-decayed estimator and its derivation from authenticated events remains open.

F. Conditional Fraud Calculation

The undetected-attempt probability u used below is local to this calculation. It is distinct from the Pool's additional-loss multiplier and from the rational certificate numerator denoted u above.

Consider a fixed integer number $z \geq 1$ of attempts. Assume attempts are independent with the same undetected-settlement probability u . Let $v > 0$ be the attacker's net gain on an undetected attempt and $c > 0$ its realized net loss on a detected attempt, after receipts, recovery, penalties and costs. Then $U \sim \text{Bin}(z, u)$ and $G = vU - c(z - U)$. The identity $u = \prod_j (1 - d_j)$ additionally assumes mutually independent detection channels; it does not follow merely from the independent-source condition (A2). The loss c must be collectible/realized. Returning the original payment alone cancels that receipt and must not also be counted as an additional net loss.

Theorem F.1 (Conditional bound on profit after fixed attempts). *Let $\beta^* = c/(v + c)$ and $0 < u < \beta^*$. Under the preceding assumptions,*

$$\begin{aligned} \Pr(G > 0) &= \sum_{j=\lfloor z\beta^* \rfloor + 1}^z \binom{z}{j} u^j (1 - u)^{z-j} \\ &\leq \exp[-z \text{KL}(\beta^* \| u)]. \end{aligned}$$

Here $\text{KL}(x \| y) = x \ln(x/y) + (1 - x) \ln[(1 - x)/(1 - y)]$ is the Bernoulli relative entropy. For $u = 0$ the probability is zero.

Proof. $G > 0$ exactly when $U > z\beta^*$. The sum is the binomial upper tail; the Chernoff bound applies because $\beta^* > u$. $\square$

In words. Under this fixed-attempt model, an exponentially decreasing bound limits the chance of a net gain. It does not say that every additional case reduces the exact probability by a fixed factor.

Correlated channels or attempts, adaptive stopping, non-collectible losses, identity changes and collusion require different models. Neither this bound nor a negative expected gain proves an equilibrium or production deterrence. (A2) alone is insufficient for the calculation's independence assumptions.

G. Source Reconciliation and Preservation

This revision preserves the full-system scope of the approved outline while retaining the corrected mathematical core of v1.2. Source specifications establish architectural intent; their descriptions of readiness do not establish implementation, clinical validity or deployment. Dated decisions govern current direction, while the claims below retain their own evidence requirements.

G.1. Source and version decisions

The source audit covered 728 core PDF artifact pages, including both overlapping architecture editions and the unchanged 75-page Gray Paper previously read in full, plus 148 supporting PDF pages and two guides. Those counts measure reading coverage, not independent scientific content. The two long Valentin Scale volumes were not fully read in that audit; four remote canonical documents were inaccessible and two historical PDFs could not be parsed. These limits remain unresolved.

Excluded-source disposition

The inaccessible remote documents are the Master Index, Product Constitution, Canonical Architecture and Initial Product Specification. The partly read rubric volumes are `valentin_scale_codex_v2.pdf` (175 pages) and `valentin_scale_codexv1.pdf` (192 pages). Parsing failed for `valentin_scale_codex_academic_masterpiece.pdf` and `vantheon_graypaper2.pdf`. This is a precise limit on reconciliation, not a finding that these sources conflict with the present direction. No claim in this revision depends on uninspected contents.

Table 9. Source/version reconciliation. Page counts identify the reviewed artifacts; source hashes are retained in the audit manifest.

Source	Treatment in this revision	Pages
Gray Paper v1.81	Preserve governance, attribution, Trust-Time, sovereignty and heterogeneous participation. Retain corrected v1.2 proof scope; do not restore unsupported historical guarantees.	75
SAS v2.0 and v1.0	Preserve distinct source identities. The cover versions and closing metadata conflict; use the governing L0–L6 decision rather than infer chronology from file dates.	90; 65
Codex v1.0	Preserve constitutional hierarchy, earned expertise, exceptions, independent appeal and compensation. Score-use and voting contracts require validation and adoption.	57

Table 9. Source/version reconciliation (continued).

Source	Treatment in this revision	Pages
Data & Integration v1.0	Preserve source/transform lineage, consent purposes and clinical/financial interoperability, including FHIR/HL7 and X12 interfaces. No adapter completeness is inferred.	73
Treasury and Subnet v1.0	Preserve entitlements, financial separation and network roles. Repair replay/recovery omissions; resolve the custom-VM versus EVM execution boundary.	105; 68
Operational Playbook; MVPilot v1.0	Preserve incident, escalation, deployment and recovery obligations. Targets, example pipelines and simulated status displays remain distinct from measurements.	101; 65
ASCLEPION design v2.0	Source for causal, uncertainty, ontology and coordination aims within THEON foundation-model and research-procedure studies.	16
SAS Addendum 1	Preserve temporal profiles and governed calibration. Repair rate units; neither normalization nor a sampling analogy proves clinical comparability.	13
Supporting corpus	Economic Architecture, engineering briefing, ENO/A2A, TIF, governance, simulator and security documents plus two integration guides inform the cross-component contracts.	148*

*Supporting PDF pages; the two guides are additional. Full source titles, identifiers and hashes remain in the source-audit records. The core sources are cited in the corresponding main sections [12, 13, 14, 17, 15, 30, 37, 27, 25].

The direction follows DEC-002 (layer responsibilities), DEC-005 (unified development), DEC-015 (the healthcare ecosystem and scoped machine approval), DEC-016 (Pool and token intentions), and DEC-017 (THEON and its research program) [20, 21, 18, 19, 16]. Original source artifacts and decision provenance remain in the source package. Research tasks, resource requirements and capital terms require their own current authorization and evidence.

Score scope and accountable change

PTI concerns provider performance for a named service. Its bounded posterior illustrates one endpoint under the stated assumptions; it does not validate every proposed incentive or voting use. CCS describes structured clinical assessment, CXS task execution, and CGS proposal-specific influence. Section 5 keeps these evidence objects separate. A changed estimator or use contract must identify affected decisions, recompute eligible projections and provide recourse. No universal conversion between objects is assumed.

G.2. Preservation ledger

Tables 10 and 11 account for every approved preservation row. The main narrative follows three connected cycles: accountable care and recognized contribution; participant-led revision of collective rules; and authorized model and research development. The ledger preserves scope while those cycles supply the reading order. A retained concept can remain a research obligation; preservation does not promote its evidence status.

Table 10. All fifteen preservation pillars and their manuscript locations.

ID	Preserved commitment and boundary	Sections
P-01	Sovereignty and ownership. Care, data-use and learning rights, portable identity/reputation, recovery and exit require operational contracts.	2, 9
P-02	Codex institution. Constitutional policy, proposals, participation, adoption, exceptions, appeal and remedies remain substantive governance responsibilities.	4
P-03	Learning purpose. Experience can motivate coordination, policy and model candidates; each needs independent evaluation and authorized adoption.	3, 9, 11
P-04	Trust with uncertainty. Context, outcomes, transparency and mission inform distinct constructs; service-specific calibration and cold-start opportunity remain necessary.	5
P-05	Causal attribution. Identify effects under explicit assumptions, then distinguish effect estimation from the agreement about contributor credit.	5, 11
P-06	Heterogeneous producers. Humans, institutions, agents and robots share accountable attribution, with task-specific authority and evidence.	2, 5, 6
P-07	Patient/provider agents. Preferences, negotiation, team formation, handoffs and escalation extend beyond conversational interfaces.	2, 6
P-08	Trust-Time. Reserve and allocate trustworthy service capacity; urgency/access constraints and auction design remain explicit research.	7
P-09	Pool and Treasury. Members fund care, depositors supply capital, producers earn; obligations and reserves precede eligible surplus.	8
P-10	Dual-token roles. Proposed VVAL backing/payment and VHTH network/staking roles remain separate from clinical competence and operating permission.	8, 10
P-11	Semantic interoperability. Clinical/financial records, devices, patient reports and social context need units, purpose and derivation lineage.	3, 5, 10
P-12	Verification and settlement. Preserve binding, replay and accounting scope; computational validity is distinct from clinical truth.	3, 8, 10
P-13	THEON. Healthcare foundation models and multi-agent/recursive research retain the ASCLEPION scientific lineage without claiming trained capabilities.	9

Table 10. Preservation pillars (continued).

ID	Preserved commitment and boundary	Sections
P-14	Operational accountability. Observe consequential behavior, support people, investigate failures and reconcile external effects during recovery.	6, 10, 11
P-15	Open applications. Patient/provider/institution workspaces, developer interfaces and third-party applications use common governed contracts.	2, 3, 6

Table 11. Eight additional obligations preserved explicitly.

ID	Obligation	Sections
X-01	Advances/TIF. Financing is distinct from earned outcome compensation; specify retained reserves, exposure and recovery.	8
X-02	Exceptions/dissent. Contextual clinical judgment, justification and appeal must survive beyond adherence scoring.	4, 5
X-03	Authority domains. Clinical, network, economic and individual powers need distinct eligible actors, scope and recourse.	2, 4, 8
X-04	Temporal addendum. Govern frequency, interval, half-life and calibration; test dependence and frequency inflation.	5
X-05	Broad contributor value. Patient, community, public-health and research contributions require defined rights and compensation contracts.	2, 8, 9
X-06	Correction/rollback. Preserve negative results and original history; determine affected scores, finance, datasets, policies and model releases.	3, 4, 9
X-07	Versioned ecosystem. Shared contracts support role workspaces, APIs and interoperable third-party execution.	3, 6
X-08	Formal corrections. Preserve all ten corrected result/proof blocks and their restrictions; do not extend them through prose.	10, 11; App. A–F

H. Shared Interface Contracts

The ten original audit boundaries remain the integration checklist in Table 12. These are minimum contract requirements, not deployed schemas or proved cross-system transitions. Each implementation must add types, validation, ownership, failure states and version migration.

Table 12. Shared interface boundaries and the records each transition requires.

ID	Boundary and required record	Sections
B-01	Care/data → evidence . Source/revision, units, permitted purpose, occurred/received times, completeness and transformation lineage.	2, 3, 10
B-02	Evidence → clinical claim . Endpoint, population, follow-up, reference/comparator, identifying assumptions, uncertainty and abstention.	5, 11
B-03	Claim → trust . Service/context, entity/model versions, risk model, effective evidence, decay, uncertainty and missingness.	5
B-04	Trust/evidence → governance . Proposal domain, competence/authority, frozen weights, representation, conflicts, quorum, challenge and effective epoch.	2, 4, 5
B-05	Policy → action . Principal and effective grant, allowed tools/physical effects, limits, delegation, escalation and responsible owner.	2, 4, 6
B-06	Action/outcome → Treasury . Obligation identity, evidence/policy/score snapshot, maturity, recipient allocation, asset, amount, custody and idempotency.	3, 8
B-07	Pool → distribution . Incurred liabilities, reserve method, capital claims, fees, liquidity and authorized distribution conditions.	8
B-08	Correction/appeal → projections . Superseded versions, remedy owner/deadline, score recomputation, payment compensation and policy/model re-evaluation.	3, 4, 8, 9
B-09	Experience → THEON candidate . Purpose-authorized data, artifact/model/tool versions, observed/predicted/causal labels, provenance and budget.	2, 5, 9
B-10	Candidate → deployment . Independent frozen evaluation, current-policy conformance, scope/promotion authority, staged integration, monitoring and fallback.	4, 9, 10, 11

Corrections and financial identity

An ordinary event identifier is not a complete economic-obligation identifier. The retained event-keyed conservation/replay model excludes corrections and contests; it does not model split releases, challenge windows or recovery. Those extensions require their own transition specification. A correction record should bind the predecessor identifier/digest, adjudicated change, effective authority, dependent versions, remedy owners and unresolved effects. Recomputation must preserve the historical decision context. An ethics or constitutional veto likewise needs an explicit holder, limited scope, expiry where applicable and review path; CP-01 illustrates a contract but does not adopt it.

Model and research records

A THEON candidate record binds dataset/model/tool versions, authorized purposes, task and population, budget, proposal history and the change under study. A research-run record adds independent evaluation versions, retained procedural changes, rejected candidates, human assistance and results. A promotion record identifies permitted use, current acceptance policy, accountable authority, monitoring and fallback. Dataset correction triggers assessment of dependent releases; it does not promise automatic unlearning. These records make the RSI comparison in Section 9 inspectable without claiming it has been executed.

I. Retained Evidence and Methods

The v1.2 source and outputs are frozen reference evidence. This revision develops architectural exposition, candidate contracts and source-derived visualizations. Its additional paired numerical and duration/severity summaries reanalyze the retained synthetic records (Appendix O); it adds no clinical study or THEON model run. Full file identities are recorded in `retained-evidence-records.json`; the copied formal input and ten exact result/proof blocks are checked against `proof-preservation.json`. Source paths below refer to the v1.2 research artifact, not a production service.

Table 1 is the manuscript’s contribution-to-evidence map. The methods below identify the retained generators and restrictions; the current RP-01, TP-01 and EP-01 profiles and parallel THEON charters have separate appendices. Candidate institutional contracts in Appendix N do not enlarge the proved or executed scope.

I.1. Trust experiment

The frozen calibration plot uses 500 synthetic participants for 24 monthly epochs, seed 20261012, and five probabilistically assigned classes with weights (0.20, 0.45, 0.15, 0.10, 0.10). Latent qualities follow Beta(18, 2), Beta(12, 5), Beta(5, 6), Beta(5, 6) and Beta(12, 5), respectively. Monthly event counts are Poisson with participant means drawn as $2 + \text{Gamma}(4, 3)$ (shape, scale). Risks are drawn from Beta(2, 5) and clipped to [0.01, 0.99]; outcomes follow equation (4) with slope 3 and reference risk 2/7. Verification of those outcomes is assumed, not measured.

The estimator uses 199 grid points from 0.005 to 0.995, prior Beta(2, 3), monthly decay 0.04 and lower quantile 0.05. The comparison is the implemented linear index $(0.4 \text{ OA} + 0.3 \text{ DT} + 0.2 \text{ MF})/0.9$ under the declared reporting strategy, not every possible historical PTI variant. The reported correlation is Pearson correlation across the simulated participants at the final epoch; mean absolute error uses their known latent q . Interval coverage is the fraction whose grid interval contains that q . The observed values are 0.988, 0.022 and 0.936,

respectively. This finite, decayed experiment is not the continuous, non-decayed asymptotic theorem and does not establish coverage under model error. Source: `sim/figures.py`, `run_population/fig3`; `sim/vantheon_sim.py`; `figures/results.json`.

I.2. Pool experiment

The frozen reserve plot uses 1,500 members, 2,000 independent paths and 60 monthly epochs. Incurred costs are compound Poisson at frequency 0.08 per member per month, with log-normal severity of mean 1 and log-scale standard deviation 1.1. Gross contributions are 1.10 times expected legitimate cost; the fee is 0.02. The assumed extra payment loss is uL_t , with $u \in \{0, 0.10\}$; no verification mechanism is executed or calibrated by that parameter. Retention and full reserve-first distribution use the same generated costs. Paths stop at first negative net assets. The named scenarios start at a three-month contribution-equivalent reserve/target; the sensitivity plot varies both from 0.25 to 4 months.

Both policies have zero observed shortfalls in the named $u = 0$ scenario, out of 2,000 paths over 60 months. The one-sided exact 95% upper limit is 0.150%; this is conditional Monte Carlo uncertainty, not a bound on an actual Pool. The plotted intervals are two-sided exact binomial intervals. There is no investment income, withdrawal, additional expense or payment delay in this experiment. Recognized unpaid obligations remain part of the general accounting model in Appendix D.3. Source: `sim/figures.py`, `fig8`, seed 20261020; `compound_poisson_claims` and `reserve_first_paths` in `sim/vantheon_sim.py`; `figures/pool_detail.json` and `figures/results.json`.

Other retained artifacts and corrections

The original CMS analyses remain descriptive companion material. Their periods, cohorts and filters do not calibrate the trust, fraud or Pool models. The restricted certificate's prior/threshold, count bounds and input-truth limitations remain in Appendix E; fresh proof timing is not a clinical authorization benchmark. Historical claims of logarithmic learning convergence, an inherited universal consensus threshold, blanket post-quantum privacy, instantaneous matured outcomes and sustainable surplus are not restored. In Economic Architecture Table 7b (pp. 16–17), the lower early-stage fee case is $500 \times \$50 \times 365 \times 0.0025 = \$22,812.50$ per year, not approximately \$2.3 million [28]. Correct arithmetic is not a revenue forecast. Current evaluation must include operating costs, liquidity, valid delayed liabilities and participant behavior.

J. RP-01: A Bounded Institutional Reference Profile

RP-01/0.2 makes selected interfaces in the running eye-examination episode executable. Its source, synthetic fixture, replay, independent workflow comparator and adverse-transition tests are packaged in `reference-profile/`. It is a research specification, not an adopted clinical policy or production subsystem. It does not modify the retained v1.2 implementation or ten mathematical statements and proofs. It records no patient data, makes no model or external tool calls and moves no funds.

J.1. Objects, domains and authority

Records concern episode E-01 under POL-07, with indivisible fictional accounting units and integer logical ticks. Ticks express ordering, not elapsed clinical time. Actor names, source independence and custodian attestations are fixture assumptions, not authenticated credentials or verified facts. The caller's in-memory state is trusted; submitted commands and capacity offers are validated as untrusted envelopes.

Table 13. RP-01/0.2 record contract. Modeled records do not establish external performance.

Object	Required fields and projection	Responsible role
Episode	Policy, performed flag, admitted evidence revision, obligations, contests and impacts	Coordinator; closure computed from registered records
Grant	Purpose, expiry and revocation; separate care and Treasury grants	Patient/clinical authority; release authority
Observation	Registered source, endpoint, strict positive revision, admission and linkage value; superseded records retained	Observer; correction by adjudicator
Obligation	Revision, due amount, outgoing/return receipts and submitted or unknown instructions	Adjudicator, Treasury, custodian fixture
Contest	Permanent identity, target revision, open or terminal status, reason and resolution command	Patient opens/withdraws; adjudicator upholds, denies or classifies a duplicate
Dependency	Explicit parent edges, affected revision, owner and disposition	Score, research, release or policy reviewer
Message	Exact closed schema, content digest, actor and nonnegative integer tick	Identical replay inert; conflicting identity rejected

Schema validation precedes hashing and mutation. Unknown fields, null or unregistered sources, unsupported actors and malformed identifiers are rejected. Integer fields exclude Booleans and floating-point coercion and are bounded by $2^{63} - 1$. Every transition operates on a copy, so rejection preserves both prior state and input. A permanent message iden-

tity binds its complete content. Care execution checks current care authority; instruction submission checks a distinct Treasury grant. Revocation of future care authority does not erase an incurred obligation. These local checks establish no cryptographic authentication, distributed locking or network finality.

J.2. Contest resolution, correction and closure

Contests can be opened only against an admitted obligation. An adjudicator can uphold or deny; the patient can withdraw; the adjudicator can classify a duplicate of a substantively upheld or denied contest. Each resolution records a reason. Duplicate classification is an authorized input, not an inferred equivalence of complaints. A closed contest identity cannot be reopened as a new complaint. A distinct complaint receives a distinct identity.

An upheld contest appends the next observation revision and a new obligation revision, preserving prior observations, instructions, receipts and projections. Multiple contests can receive successive corrections or distinct terminal dispositions. This supplies a completion path when responsible actors furnish authorized resolutions; it is not a deadline or unconditional liveness guarantee. Denial here concerns a contest disposition, not an automatic denial of care.

Let d be current due units, p confirmed outgoing units, r confirmed returned units and q submitted or unknown units. A new positive integer instruction requires no open contest, current admitted evidence, effective release authority, $q = 0$, and amount at most $d - (p - r)$. A matching simulated custodian receipt discharges an instruction once. Unknown results require reconciliation; a correction cannot pretend to cancel an external effect.

Closure requires performed service, admitted current evidence and obligation, no open contest, $q = 0$, $p - r = d$, and a completed authorized disposition of every affected registered descendant. A restricted release pending assessment blocks closure. Dependency traversal discovers consequences *within the registered graph*; it cannot discover an omitted care duty or missing graph edge. A later contest reopens the projection without erasing earlier closure.

The fixture initially determines $d = 100$: 80 service units plus 20 contingent report-linkage units. An admitted correction changes linkage qualification to false and d to 80. With $p = 100$, a separate 20-unit return receipt is required to reconcile the excess. No actual recovery, contractual write-off or solvency follows from this arithmetic.

Table 14. Executable late-correction trace, with completed and incomplete states preserved.

Tick	Authorized command	Modeled consequence
1–5	Perform, observe, determine, submit and confirm	Revision 1; 100 due and discharged; registered-set closure
6–8	Register research lineage; contest; uphold correction	Revision 2; 80 due, 100 paid; 20 recovery claim; registered descendants pending
9–10	Recompute score; quarantine dataset row	Toy posterior mean changes $1/2 \rightarrow 1/3$; affected research row ineligible
11–13	Restrict reference, record policy proposal, confirm return receipt	Net discharge is 80; assessment remains pending; closure false
14	Reassess hypothetical research reference	Zero eligible rows rejects reference; declared assessment completes; closure true
15–18	Open two contests; deny one; classify the other as its duplicate	Intermediate hold remains until both resolve; closure true; historical states retained

Score recomputation is now executed, but narrowly: an exact $\text{Beta}(2, 3)$ posterior *mean* for one fictional binary observation. It is neither clinical PTI nor a lower-quantile authorization, causal or calibration result. The research reassessment executes a deterministic data-lineage eligibility gate: quarantining its sole registered row yields zero eligible rows and rejection of a hypothetical reference. It does not evaluate trained THEON weights or implement unlearning. Recording a policy candidate does not activate POL-08. The ledger remains under POL-07. Clinical evaluation, model evaluation and governance adoption remain separate obligations.

J.3. Capacity validation

The single-request capacity profile accepts only unique bundle identifiers, nonnegative strict integer ticks, completion at or after visit, and strict Boolean authorization, patient acceptance, trust eligibility and slot-availability predicates. It validates the entire submitted snapshot, filters feasible bundles, and orders them by completion, visit, then identifier. Reservation repeats validation and feasibility checks before an in-memory update. The supplied instance rejects B-01, whose interpretation slot is unavailable, and selects B-02. Concurrent overlapping resources, urgency assignment, fairness, pricing and interrupted handoffs remain outside this profile.

J.4. Finite comparative evidence and its limit

The package contains 73 tests and a deterministic comparison of 12 named workloads with 116 paired command attempts. An independently written conventional append-only workflow ledger reconstructs its projection by replay. It implements its own domain reducer, graph traversal, closure and toy calculations, while sharing the command schema and initial state format. Both receive the same admitted obligations, dependency graph, authority assumptions, receipts and corrections. Both assume authenticated envelopes; neither implements signatures. Shared schema bugs can escape a differential comparison.

Workloads cover ordinary closure, completed and incomplete corrections, multiple contest dispositions, repeated upheld revisions, unknown external results, duplicate identity, malformed inputs, revoked authority and deliberate dependency omission. Admissions and projections agree at every step; the specified admission and known-unresolved-closure oracles have zero violations. These are finite authored cases, not statistical estimates, clinical evidence or a universal correctness proof.

The omission control removes the policy-candidate dependency. Both systems close relative to the remaining graph and reject the unregistered policy disposition. This negative control exposes the shared completeness assumption. The result establishes parity on the supplied contract, with *no evidence of architectural superiority or necessity of decentralization*. Elapsed latency, human effort, operating cost, crash durability and distributed failures were not measured. A future discriminating experiment must specify which additional property the coupling supplies, give a competent baseline the same information, and test the resulting tradeoff rather than assuming that integration is novel or better.

K. TP-01: Exact bounded gate and synthetic trust comparisons

TP-01 separates a repaired numerical decision from a synthetic estimator comparison. The former implements an exact predicate for a restricted conjugate posterior; the latter measures selected estimators under an authored generator. Neither establishes clinical trust, source independence, consent, effective authority or a valid healthcare authorization.

K.1. Exact integer-count decision

Admit nonnegative integer counts s, f with $s + f \leq 1000$, reference risk $r = 2/7$, no temporal decay, and the stipulated $\text{Beta}(2, 3)$ prior. Booleans, floating-point counts and out-of-domain values are rejected. For

$$a = 2 + s, \quad b = 3 + f, \quad n = a + b - 1,$$

the posterior is $\text{Beta}(a, b)$. The incomplete-beta/binomial identity in NIST DLMF 8.17.5 gives [53]

$$F_{a,b}(3/4) = \frac{\sum_{k=a}^n \binom{n}{k} 3^k}{4^n}. \quad (8)$$

Its density is positive on $(0, 1)$, so its lower 5th percentile is at least $3/4$ exactly when $F_{a,b}(3/4) \leq 1/20$. The implementation therefore evaluates

$$20 \sum_{k=a}^n \binom{n}{k} 3^k \leq 4^n \quad (9)$$

with arbitrary-precision integers. No floating quantile, finite grid or tolerance determines this result. The response **APPROVE** means only that this numerical predicate holds; its scope field explicitly excludes clinical sufficiency, source authenticity and grant validation. An unsupported posterior profile, including general risk conditioning or decay, returns **REVIEW**. This path does not approximate or authorize such a posterior. It is separate from the retained proof-circuit relation and from any operational API score.

Table 15. Five audited decision-boundary cases. Continuous quantiles are displayed numerically for explanation; the repaired decision uses Equation 9. Each retained-grid value would cross the 0.75 threshold.

Successes	Failures	Continuous lower 5%	199-point grid	Exact gate
32	3	0.749147	0.750000	Review
36	4	0.748492	0.750000	Review
40	5	0.748357	0.750000	Review
44	6	0.748534	0.750000	Review
48	7	0.748903	0.750000	Review

The supplied 16 tests exercise these cases, invalid count domains, unsupported profiles, cold start, a qualifying success case, all count pairs with total at most 120, and selected boundaries at totals 200, 500 and 1000. Agreement with a numerical library on these cases is a regression check; exactness follows from the identity and integer predicate within the declared domain. Numerical correspondence for the full risk-conditioned, decayed posterior remains an independent obligation.

K.2. Synthetic design and comparators

The estimator experiment uses four regimes, each with 24 independently seeded populations of 128 synthetic participants: 12,288 participant instances and 49,152 participant–method records. Every participant has 40 recorded observations. The duplicate regime contains 20 independently generated episodes copied twice with the same explicit identifiers. The other

regimes contain 40 independent episodes. There is no clinical dataset, fitted healthcare parameter, common population shock or prospective registration.

Latent reference-risk quality is drawn from $\text{Beta}(4, 2)$. Binary outcomes follow

$$\Pr(Y = 1 \mid q, r) = \sigma(\text{logit}(q) - \beta_r(r - 2/7)).$$

The matched regime uses risk $\text{Beta}(2, 5)$ and $\beta_r = 3$; the risk-shift regime uses risk $\text{Beta}(5, 2)$; the slope-misspecified regime uses true $\beta_r = 1.5$ while the estimators retain 3. The duplicate regime uses the matched likelihood with repeated records. “Matched” refers to the likelihood, not the prior: every estimator uses $\text{Beta}(2, 3)$ while generated quality follows $\text{Beta}(4, 2)$. Nominal Bayesian interval mass is therefore not a guarantee of frequentist coverage under this generator.

Four methods receive the same permitted observation projection. The retained 199-point grid uses all recorded rows. Its deduplicated variant selects one row per explicit episode identifier. A conventional risk-adjusted comparator uses the same Bayesian prior and likelihood on a finer 1001-point grid and the same deduplicated rows. An unadjusted conjugate Beta estimator uses the deduplicated outcomes but omits risk conditioning. Deduplication is executed from the exported identifiers, not privileged pre-duplication arrays. It does not solve unknown-identity record linkage or correlated distinct episodes. The conventional comparator is a strong established alternative, not evidence of a new estimator; its finer grid is not an exact-integration certificate.

Metrics are posterior-mean absolute error against generated q , coverage of the nominal central 90% interval, the numerical threshold-pass rate, and false passes among participants with $q < 0.75$. The latter two are simulation metrics, not clinical approvals. Mean absolute error and coverage use two-sided 95% t intervals across 24 population-level measurements. Pass rates use pooled counts and denominators with 95% Wilson intervals under the independent-participant generator. These sampling intervals are distinct from each participant’s posterior interval. Zero observed false passes retain a positive Wilson upper bound. No interval accounts for clinical transport, model selection across experiments or unmodeled dependence.

K.3. Observed parity and adverse results

The matched retained grid has mean absolute error approximately 0.0580, effectively the same as the conventional risk-adjusted comparator in this experiment. Its nominal 90% interval covers 85.03% of generated qualities; the conventional grid covers 85.25%. This is a calibration shortfall, not a successful uncertainty-validation claim. Risk adjustment materially changes the comparison with the unadjusted estimator after a risk shift, but supplies no advantage over the conventional estimator using the same likelihood.

In the duplicate regime, the retained grid’s coverage falls to 73.01%. Known-identifier deduplication raises it to 82.42%, still below 90%, while mean absolute error increases from 0.0771 to 0.0852. Deduplication therefore does not improve every reported metric. False passes change from 18 of 1,949 below-threshold participants to 1 of 1,949; their Wilson intervals remain in the machine-readable results. These finite synthetic observations do not establish a clinical error guarantee or general correction for dependence.

The package retains participant results, population summaries, compressed raw observations, seeds, source hashes, all four regimes and all comparators. The exact gate and estimator study have different domains: the risk-conditioned grid estimates in this experiment do not invoke Equation 9. Delayed labels, selective missingness, changing quality, generalized dependence, causal attribution and prospective clinical calibration remain untested. No superiority, novelty or full-model authorization claim follows from these artifacts.

L. Economic reference profile and paired stress study

EP-01 is a reproducible, nonbinding research profile. Its integer-cent ledger and synthetic Pool study are separate artifacts: the ledger includes an issuer and Treasury, whereas the study compares cash-only Pool policies. Neither operates custody or funds. The source owner is `economic-profile/`; `results/profile.json` declares all simulation inputs and `results/worked-ledger.json` contains the complete thirteen-state ledger.

L.1. Accounting, admission and failure fixtures

The ledger admits minting, redemption, member income, capital subscription, incurred care, partial care payment, external expense, network fee, TIF origination and resolution, gated distribution and gated capital exit. All amounts must be positive integer cents; Boolean and string amounts are rejected. An issuer records a cash freeze without deleting the token liability. An asset haircut reduces issuer equity and suspends further minting; balanced books may therefore report a negative capital position. No capital or clinical claim follows from bookkeeping consistency alone.

Seventeen tests cover these transitions, post-distribution capital protection, principal-versus-income classification, cross-entity conservation, failed-operation atomicity, emergency restrictions and 5,000 seeded operation attempts. A consolidated balance removes the Pool’s VVAL holding and the issuer’s matching liability before adding the three entities; the issuer’s backing is never simultaneously recorded as an independently available Pool asset. Unavailable redemption cash and exhausted capital are represented as failure conditions rather than erased obligations. The model assumes received cash and submitted payment amounts are

genuine; it does not authenticate bank receipts, establish segregated custody or choose enforceable insolvency priorities.

L.2. Selected illustrative comparison

Each of four scenarios uses 1,000 paired paths over 36 monthly periods, seed 20261015. The same random claim history is supplied to both policies. Each month has a Poisson claim count with mean 20 and independent lognormal severities with mean 5 and coefficient of variation 1. A monthly shock with probability 0.03 multiplies that month's aggregate by 1.6. Inputs are chosen sensitivity parameters, not healthcare estimates. Opening Pool capital is 300, monthly member income is 115 and monthly administrative expense is 8. All incurred care is recognized immediately; 70% is scheduled for payment that month, 20% one month later and 10% three months later. Overdue amounts remain liabilities and mature payments receive available cash before discretionary withdrawals or distributions.

Both policies advance 40 at month 1, with principal due at month 12 and no interest. Both receive the same aggregate capital-exit request of 90 at month 18; unfulfilled requests are carried forward without a promise of immediate par redemption. A withdrawal requires the remaining net capital to be at least 300 and cash to cover all unpaid care plus a 50 buffer. The conventional policy retains all gains. The shared-surplus policy applies the same accounting and gates, then distributes 50% of eligible surplus quarterly. This comparator already contains competent reserve and incurred-liability accounting. The experiment tests the distribution policy, not the necessity of tokens, decentralized governance or clinical verification.

The four scenarios are the base inputs; a 35% increase in incurred care during months 13–24; temporary unavailability of 80% of cash from month 12 until release at month 16; and complete impairment of the advance at month 12. The last two shocks distinguish unavailable assets from asset losses. Runs continue through all 36 months after a shortfall, with no external recapitalization. End-of-horizon unmatured care remains in the liability balance.

Table 16. Synthetic EP-01 endpoint counts. Each entry is conventional/shared-surplus out of the same 1,000 paired paths; events are counted once per path. Figure 10 includes uncertainty.

Scenario	Any negative equity	Any due-care shortfall
Base	37/47	19/22
Care cost +35%	622/770	488/628
Temporary cash unavailability	37/44	588/706
Advance impairment	49/65	25/30

Negative equity means assets less recognized unpaid care falls below zero at any recorded month. A due-care payment shortfall means an amount scheduled by that month remains unpaid, even when total assets exceed liabilities. The analysis reports 95% Wilson intervals for each event frequency and percentile intervals from 2,000 paired path-bootstrap samples for policy differences. Paths, not monthly observations, are the independent sampling units. These intervals concern Monte Carlo variation conditional on this model; uncertainty in healthcare pricing, reserve targets, liability estimation and behavior is unmeasured.

L.3. Artifacts and limits

The release retains 288,000 monthly state rows in `results/raw-paths.csv.gz`, exogenous claim draws, all eight policy/scenario summaries, paired differences, trajectory quantiles and source/output hashes. An independent parser checks six accounting identities on each retained row (1,728,000 identity checks), recalculates event counts and checks paired binary differences. Reproduction uses the governed figure environment's NumPy; the environment version is recorded. Generation does not overwrite the website or retained v1.2 artifacts.

The example removes several accounting ambiguities but leaves material economic questions open. It does not estimate incurred-but-unreported claims, adverse selection, member churn, individual ownership units or entry pricing; it excludes interest, taxes, token-market prices, strategic reporting, real recovery enforceability and governance adaptation. Individual capital loss and exit rights must be selected before an offer. Member rebates, depositor distributions and producer rewards require a separately declared allocation rule; the experiment measures their aggregate available amount only. A future calibrated model must compare useful care delivered, affordability, denied or delayed valid payments, concentration and continuity of care alongside financial endpoints. A distribution can advance shared value while reducing resilience; the chosen policy must make that tradeoff accountable.

M. THEON model and procedure study charters

M.1. TH-FM-01: first model-learning charter

The first weight-learning object is an adapter for structured evidence grounding. Inputs are synthetic FHIR R4 episode resources with stable source identifiers, time, service context and explicit missingness. Outputs are typed evidence-to-service links, contradicted or missing dependencies, correction descendants, source citations and calibrated abstention. It makes no diagnostic, treatment-effect or robot-control claim. Synthetic records must come from a versioned permitted generator such as Synthea, with the generator configuration and source-to-target mapping independently reviewed before a model study [54]. Generator truth is a task reference, not clinical truth.

The proposed first split has 6,000 training, 1,000 development and 2,000 protected assessment episodes, separated by synthetic patient and template family before generation. Numerical counts are feasibility targets. Counterfactual sibling episodes remain in the same partition; held-out correction patterns and service templates test transfer beyond alias memorization. Data provenance records the generator/license, seed, schema, transformations, objective and permitted research purpose. No real patient data, proprietary clinical corpus or implied training consent is introduced by this charter.

The first route is supervised low-rank adaptation of one licensed open-weight text checkpoint, preserving the frozen base and tokenizer [55]. The objective is token likelihood of the canonical typed output, with constrained decoding and separate validation of references. Retrieval content, prompt, tools and decoding are held fixed for the base-versus-adapter comparison. Rank, learning rate, batch, epochs and stopping rule must be frozen after a development-only feasibility rehearsal; checkpoint identifier/digest, license and exact environment are mandatory pre-run fields. No claim is made that a new pretraining architecture or training from scratch is justified. That decision requires an explicit capability, data-rights and compute comparison.

Baselines include deterministic extraction and graph linking, the identical frozen base model, and conventional full fine-tuning under a matched training-compute limit where feasible, reporting trainable parameters and memory separately. TH-FM-01 first establishes an engineering baseline for standard adaptation; no new adaptation method is claimed. If full fine-tuning is infeasible, omit that comparison explicitly rather than representing a second identical adapter as an independent method. A stronger model can be an additional costed comparator, not a substitute for the controlled ablation. Assessment reports macro edge F1, incorrect source attribution, missed required dependencies, unsupported closure recommendations, abstention/coverage, calibration, subgroup counts by service/template and full training/inference/review resources. Patient/template families, rather than individual edges, are the resampling units. Every run, failure and selected/rejected checkpoint remains in the result ledger. No task result automatically increases PTI or grants governance weight.

Before execution, freeze model and corpus rights, generator/mapping review, split and evaluator hashes, feasible resource limits and an independent assessment custodian. Before claiming improvement, define the primary metric, material-error constraint, repetition unit and analysis plan from development-only evidence. The model registry records identity, objective, training configuration, data versions, evaluated scope, failures and dependencies; unregistered weights cannot become a release. This is an executable-study specification to complete, not a trained-model artifact. Later multimodal, causal, Bayesian and robotics studies need corresponding labels, mechanisms, baselines and authority.

M.2. TH-RP-01: exploratory research-procedure study

TH-RP-01 makes an initial coordination-research task inspectable without claiming a trained healthcare model. Given a synthetic care episode, the system must link evidence to its service, identify missing or contradictory dependencies, and return the objects requiring reassessment after a correction. The proposed output is a typed dependency graph with source identifiers, an uncertainty or abstention flag, and a reason for each proposed link. It cannot close an event, deny care, release funds or modify a clinical rule. The profile connects to the institutional reference episode RP-01: a correction to observation OBS-01 in event E-01 can affect obligation O-01, a decision under POL-07, and research release DS-01. Reassessment is a distinct action from reversing a payment or removing learned information.

The first study would use authored synthetic episode graphs, not patient records. Inputs contain observations and incomplete or provisional source-reported links; they do not contain the completed gold dependency graph. Every baseline and experimental arm receives the same permitted input projection. Gold labels are the generator’s service membership, required-dependency edges, correction ancestry and unresolved-state annotations. Those labels are withheld from task inference and controlled by the evaluator; development feedback follows the declared split rules. Independent checks must test the labels against the written profile; agreement with a generator alone establishes no clinical validity. A proposed initial design uses 1,200 episodes: 600 for development, 200 for admission validation and 400 for final assessment. Template families, actor aliases and correction patterns are partitioned before generation to reduce direct template leakage. The final partition and its labels remain under a separate evaluator’s control; hashes of the generator, splits and evaluator are frozen before candidate search. The supplied small replay fixture exposes its complete graph to test bookkeeping and is not a task-evaluation input or a protected test set.

The proposed model route adapts the prompts, retrieval context and workflow around one licensed, frozen text checkpoint, denoted M_0 . A study registration must select its exact identifier, weight digest, license and inference configuration before any model run. The first profile does not train new foundation-model weights. It compares exact-identifier linking, a stronger deterministic dependency-graph baseline, a single agent using M_0 , and fixed versus adaptive research procedures using the same M_0 and tools. Model-weight learning remains a distinct THEON workstream. This bounded starting point does not replace the program’s multimodal, causal, probabilistic or robotic ambitions.

The initial exploratory experiment has five independently seeded research histories and three cycles per history. Each arm receives the same development and admission partitions and an aggregate limit of 60 model calls, 120,000 input-plus-output tokens and 120 execution minutes per history. Calls to critics, selectors, retries and unsuccessful candidates all count. A fixed executor records cumulative spend separately from learner-produced artifacts. Every

attempted candidate is charged, including rejected or malformed artifacts; unknown accounting halts the history rather than being treated as zero. The replay checks post-execution accounting; an experimental executor additionally needs pre-dispatch budget reservation and timeouts. At most three agent roles may operate in an arm; no unrecorded human repair is allowed. Infrastructure, model settings, cached context and permitted tools are matched, and actual model calls, tokens, tool time, accelerator time when observable, elapsed time and human interventions are reported separately. These are author-review parameters, not established sufficient sample sizes or approved institutional rules.

Only the adaptive arm may retain an edit to its candidate-generation or candidate-selection procedure: prompt sequencing, role allocation within the three-role limit, retrieval ordering or the ranking of development-tested candidates. Gold labels, split membership, evaluator code, admission thresholds, resource limits, tool permissions and Codex policy are outside the editable set. A proposed procedure P_{t+1} is retained only if a fixed admission check finds a valid artifact, no additional critical error, budget compliance and an improvement of at least 0.01 in macro edge F1 over P_t on the admission partition. Otherwise the arm retains P_t and records the rejection. Admission results are development feedback and do not establish generalization. Research admission authorizes another sandbox cycle; it does not constitute clinical deployment approval.

The frozen final evaluator compares both arms only after all histories end. It can then assess the frozen task artifacts from every cycle, enabling a prespecified later-cycle comparison without returning feedback to search. Its primary endpoint is macro edge F1 for the final-cycle task artifacts; secondary endpoints are correction-descendant recall, abstention and coverage, critical-error counts, per-cycle differences and complete resource use. A critical error includes treating an unresolved dependency as satisfied or an unauthorized action as permitted. The five-history stage is exploratory. It reports every paired history, final and per-cycle gains, critical errors, abstentions and resources; it has no confirmatory success declaration. A gain of 0.03 is a planning target rather than a decision threshold. A percentile bootstrap from five identical gains can have zero width, so nominal 95% coverage is not asserted. Before a confirmatory study, development-only pilot variation and simulated null/alternative histories must determine independent-history sample size, interval coverage, power and the exact analysis rule. The benchmark case count is not the number of independent research histories. Case-family uncertainty and research-history uncertainty are reported separately. Final scores and cases are unavailable to the search process; reuse for a later experiment requires a fresh protected partition. These comparisons can test task and procedure improvement, not diagnosis, patient benefit or healthcare RSI in general.

The package contains `research-profile/profile.json`, an executable deterministic replay and a recorded check report. The replay exercises identifier lineage, corrected-versus-

uncorrected dependencies, fixed admission rules, rejected edits, cumulative accounting and protected-evaluator access. Admission requires finite F1 in $[0, 1]$, nonnegative integer counts, finite nonnegative time and correctly typed fields; malformed records fail closed. It uses authored candidate records and exact graph operations. It neither invokes M_0 nor searches for a better procedure; its passing checks are harness evidence, not an RSI result. A model study remains blocked on checkpoint registration, independently reviewed synthetic labels, protected assessment custody and a pre-run protocol freeze.

N. Candidate institutional contracts and discriminating evaluations

The following profiles make unresolved design choices inspectable. They are *nonbinding candidates for research*, not adopted constitutional rules, participation terms, clinical authorizations or deployment descriptions. They preserve the confirmed direction while separating decisions that evidence can inform from rights that require legitimate adoption. The companion `evidence/design-profiles.json` records actors, objects, authority, preconditions, effects, recovery, alternatives and decision gates. CP-01, RP-01, TP-01 and EP-01 retain their existing scope; these profiles do not extend their tests or proofs.

N.1. DP-SCORE-01: meanings precede permitted uses

The candidate dictionary selects distinct objects rather than one universal reputation number. A record identifies its subject, service or task, population, evidence window, endpoint, estimator or rubric, uncertainty, source and policy versions, adjudication status and permitted use. Missing evidence is unknown, not zero quality. Supersession retains the old object, records a reason and schedules reassessment of its registered consumers.

Table 17. Candidate score/use dictionary. Each object has a defined purpose; an institutional use requires validation and adoption.

Object	Selected meaning for testing	Allowed use and boundary
PTI	Selected estimator profile: posterior over latent performance for a named service, risk model and endpoint; mean and lower quantile remain distinct.	Present profile: research eligibility/review. Proposed assignment, compensation/financing and voting uses require separate adopted contracts (Sections 4, 5, 8); this estimator grants no entitlement or patient-rights weight.
CCS	Structured clinical assessment: endpoint value, horizon, missingness, patient context and independently qualified assessment. No scalar composite is selected.	Clinical review and endpoint research. A causal effect is a separate object with an identification argument; neither directly determines contributor credit.
CXS	Task execution assessment: required steps, justified deviations, handoff completion and adjudicated failures, with counts and exposure.	Operational review for that task and entity type. Fast telemetry and self-reported confidence do not become independent clinical successes.
CGS	Proposal-specific influence record: domain, eligible principal, frozen service evidence, conflict status, chamber and declared weight.	Only the named proposal under its charter. It confers no clinical authority, patient representation or independent Treasury grant.
Credit allocation	Versioned normative rule allocating a recognized contribution among named producers.	Calculate an entitlement under declared terms. Allocation is not causal effect; ownership is not inferred from performance.

Table 18 distinguishes the evidence objects and the decisions they can support. A shared numerical range does not make their units, uncertainty or authority interchangeable.

Table 18. Score-use boundaries. Evidence informs a named decision under its own contract; it does not create a universal entitlement.

Object	Evidence and interpretation	Decision boundary
PTI	Service-specific provider performance under a declared endpoint, risk model, evidence window and uncertainty estimate. Transparency and mission alignment remain separately assessed attributes.	Assignment, compensation and relevant expertise weight each require adopted terms. Patient rights and representation remain independent.
CCS	Structured assessment of clinical evidence: endpoint, horizon, missingness, patient context and qualified review. A causal effect additionally needs an estimand and identifying assumptions.	Clinical review does not by itself establish causal benefit or divide contributor credit. A composite rubric needs its own validation.
CXS	Required task steps, adherence, justified deviations, timing, coordination, handoff completion and adjudicated failures, with exposure counts.	Operational performance is task-specific. Frequent telemetry or procedural compliance alone does not establish successful clinical outcomes.
CGS	Proposal domain, eligible principal, frozen service evidence, conflicts, chamber and declared influence rule.	A score alone grants no constitutional authority. Representation, eligibility, quorum, appeal and activation remain separate gates.

The clinical-methods owner and patient representatives should test this dictionary through blinded interpretation tasks, source-to-score traceability and misapplication cases. Advancement requires correct distinction of uncertainty, authority, credit and voting power, plus service validation and appeals. Proper scoring rules evaluate forecasts; they do not establish the clinical utility of a chosen threshold [48, 49].

N.2. DP-GOV-01: expertise with an independent patient boundary

The primary research case remains CP-01: three patient principals, three service experts, separate chamber gates and authored weights (3, 2, 1). These weights apply only to its

stated fixture values; no continuous political-entitlement function is adopted. Higher PTI can inform expertise in the relevant service while patient consent, constitutional rights, clinical authorization and financial ownership remain outside its influence. Robots and agents can supply proposals, evidence or delegated analysis through accountable registered principals. Section 4’s separate human–machine illustration connects a machine’s service record to its operator’s represented entitlement, while keeping compensation and voting terms distinct. Spawning another instance creates no additional principal or vote.

The comparator gives each qualified expert one vote while retaining the same patient chamber, information, conflicts and activation controls. A third candidate tests capped expertise bands, with boundaries frozen before voting and justified separately from the patient decision. All candidates require conflict screening, a fixed electorate, quorum, independent appeal and a separate activation grant. Suspension expiry requires review and cannot activate policy. A conflict discovered after the snapshot invalidates the tally and requires a newly disclosed snapshot and vote.

The governance research owner, patient representatives and domain professionals should compare proposal quality under blinded review, participation burden, chamber disagreement, influence concentration, strategic case selection and coalition capture. Sensitivity tests vary evidence uncertainty and work-allocation feedback. Passing a tally does not establish legitimacy. Adoption requires an accountable charter process and reasons for the chosen tradeoff. Snapshot voting and delayed execution are established techniques [47]; the VANTHEON question is how service evidence informs collective care improvement without overruling affected people.

N.3. DP-RIGHTS-01: participation, losses and exit

This candidate extends EP-01 only as a proposed individual-rights model. Members fund an agreed care arrangement and retain its coverage, representation and appeal rights. Depositors supply explicitly loss-bearing Pool capital and hold separately recorded participation units. Producers hold claims for eligible performed work; humans and accountable principals behind agents or robots receive compensation under their contracts. These capacities may coexist in one person, but their claims do not merge.

The candidate prices subscriptions and exits at assessed net asset value per active unit, after recognized care liabilities and an explicit estimate for unreported liabilities. A subscription purchases units; it creates neither operating income nor distributable surplus. Bootstrap pricing is declared before the first subscription. If assessed net resources are nonpositive, ordinary issuance and exit stop pending a separately authorized recapitalization or wind-down profile; issuing more units cannot conceal a deficit.

For testing, gains first rebuild the declared protection target. Losses consume retained operating gains and then erode contributed capital pro rata within the same unit class. Unpaid care and producer liabilities survive capital exhaustion. Payment-token holders retain a distinct issuer claim; issuer backing and restricted development resources are excluded from Pool loss absorption unless a separate permitted contract supplies them. This is a proposed accounting and allocation order, not a statement of enforceable insolvency seniority. EP-01 has not implemented individual units or an unreported-liability estimator.

Eligible surplus is divided by a predeclared fraction between member benefit and capital participation. The fraction is an experimental parameter, not an adopted rate. The candidate allocates the member portion by eligible contributions during the observation period and the capital portion by time-weighted units; producer compensation remains an expense or liability under its own rule. Sensitivity analysis must test disadvantage to high-need members and entry-timing incentives. Prospective benefits can lower member contributions; they do not buy authority over another person's care.

Exit requests receive one disclosed valuation date and a common cohort price. Only the amount compatible with recognized liabilities, the unreported-liability estimate, a liquidity buffer and the capital floor can be released. If liquidity is insufficient, eligible requests in the same cohort receive proportional fills and the remainder stays visible for review. Falling valuation can reduce eventual recovery; immediate par redemption is not promised. Care exit and portability are distinct from capital withdrawal. TIF is *Trust-Indexed Funding*: advances remain receivables with repayment, impairment and loss terms, never already earned care compensation.

The economic design owner must compare this candidate with fixed-term capital and a retained-gains policy under identical claims, costs and valuation information. Endpoints include care-payment delay, affordability, capital loss, cohort dilution, exit fulfillment, run incentives and care continuity. Valuation, surplus fractions, custody, enforceability and wind-down require adoption before an offer. The historical Dai paper makes debt, collateral and emergency unwind explicit [31]; it does not validate a reserve-backed VVAL issuer or healthcare Pool.

N.4. DP-CAP-01: a scheduling profile before a market

The first research profile is the RP-01 single-request bundle scheduler. Its objects are visit, imaging and interpretation bundles with identifiers and completion times. Patient acceptance and current authority, eligibility and availability filter offers; feasible offers are ordered by completion tick, visit tick and immutable identifier. Reservation rechecks the selected offer and marks its resources unavailable in one in-memory snapshot. These choices are implemented locally; clinical urgency and cross-institution concurrency are not.

A proposed extension records reservation expiry, cancellation reason and reassignment without labeling a reservation as performed care. Failed bundles return outstanding tasks to a human-supervised queue. Urgency, accessibility and patient preference are constraints supplied by adopted care policy. New participants receive a reviewed evidence-building path rather than permanent exclusion by a cold-start score.

The comparator is a dependency-aware earliest-feasible scheduler with identical offers and constraints. Evaluate completed episodes, waiting time, cancellations, burden, subgroup access and stale-offer sensitivity. A future Trust-Time market must choose bid language, private information, prices, payment and allocation rules, and test withholding and collusion; the local ranked example supplies none of those properties. Market design requires a mechanism and adversarial evaluation, not merely an auction label [50].

N.5. DP-DEPLOY-01: the first integration boundary

The candidate first deployment is a laboratory integration with synthetic records, named test principals, one authoritative event service and separately permissioned care, adjudication, Treasury and research adapters. It is a target to implement; RP-01 remains one local deterministic process. Record bodies stay in the laboratory store, with versioned references at interfaces. The custodian adapter accepts fictional or sandbox receipts. This profile admits no participant funds, PHI, issued token, treatment instruction or external research release.

The deployment owner must bind commands to authenticated principals, grants, object versions and idempotency keys; retain dependency lineage; distinguish instruction, unknown external effect and confirmed discharge; and reconcile outstanding effects before retry. Recovery drills cover stale grants, unavailable adjudicators, missing dependents, duplicate receipts, key loss and interrupted writes. Recovery restores consistent state and reopens unresolved obligations instead of asserting closure from an empty queue.

A conventional event service is the first comparator. A second candidate may add commitment anchoring only when an experiment names the independently verifiable property gained and its latency, privacy and operating cost. Neither Avalanche, an EVM deployment nor a custom VM is selected here. Each network candidate needs finality, serialization, fault assumptions, verifier compatibility and upgrade custody before comparison. A subsequent shadow pilot requires a named institution, accountable clinicians, patient/data authority, real custody boundaries where relevant, clinical stop criteria and an independently approved protocol. Those inputs have not been supplied by this paper.

N.6. DP-EVAL-01: what would distinguish the composition

The first discriminating study should give an integrated implementation and a competent conventional workflow identical episode records, dependency inventories, grants, correction

notices, cost budgets and human operators. Three paired task families test the full proposition: coordinate a patient-accepted care bundle through completion and compensation; turn a recurring care problem into a reviewed policy proposal with adoption or reasoned rejection; and produce an authorized, traceable research artifact under a locked evaluation. A correction perturbation then tests affected care, trust, entitlement, policy-candidate and learning dispositions. Additional workloads introduce late outcomes, newly discovered dependents, conflicting grants, failed handoffs and unknown custody effects. Blinded adjudication determines correctness; an empty dependency list is not the reference truth.

Report completed care bundles, access waiting time, participant burden, earned compensation discharged, blinded proposal quality, effort to adopt safely or reject, and research-artifact eligibility and reproducibility. Pair those benefits with false closure, time to correct every affected artifact, unresolved obligations, human actions, compute and operating cost, and adverse consequences. Give denominators for each task family, distributions, uncertainty and exclusions. The comparator may use provenance, replay and compensation; withholding these established techniques would create a weak baseline [10, 44, 45]. Repeated seeds test software behavior; patient outcomes require a separate clinical study.

THEON retains parallel weight-learning and research-procedure tracks. Locked model evaluation measures the former; a fixed-procedure comparator and held-out transfer test examine whether a changed research operator produces better subsequent candidates at matched budget. Model cards provide an established reporting precedent [51]; independent evaluator custody, patient/data authority, rollback and promotion remain explicit VANTHEON obligations. No integration, model-learning or RSI advantage is claimed before those comparisons are executed.

O. Derived views of retained synthetic results

The v2.3 visual revision adds exploratory analyses of the retained TP-01 and EP-01 data. It does not regenerate outcomes or introduce a new clinical dataset. The canonical transformation is `figures/build_empirical.py`; full-precision records, source hashes and definitions are packaged beside each figure.

O.1. Paired numerical refinement

Figure 9C pairs the mean absolute error of the deduplicated 199-point grid with that of the conventional 1001-point grid by regime and population. Both use the same known-identity deduplicated records, likelihood and prior. All 96 population pairs are shown; the largest absolute difference is 9.36×10^{-7} . The diamonds and intervals summarize 24 paired

differences per regime using a two-sided 95% t interval. This is observed numerical agreement in these instances, not an integration-error certificate or a universal tolerance guarantee.

O.2. Shortfall timing, duration and severity

A payment-shortfall event is `due_arrears` greater than 10^{-8} in a recorded end-of-month state. Cumulative incidence is the fraction of all 1,000 paths that have experienced this event by a month, with pointwise 95% Wilson intervals. No path is removed after its first event; all 36 monthly observations remain available. A step curve connects recorded months and makes no claim about within-month timing. Shaded windows in Figure 10A denote the selected cost and cash-lock shocks; a dotted line denotes the advance impairment.

Duration is the number of recorded months with an event, from zero to 36. The paired difference subtracts the retained-surplus path from the corresponding distributed-surplus path before averaging. A new 2,000-replicate percentile bootstrap resamples the 1,000 path pairs; NumPy seed 20263318 initializes one continuous PCG64 generator, traversing base, care-cost, cash-lock and advance-impairment scenarios in that order; percentile endpoints use linear interpolation. Both zero-event and nonzero-event paths remain in the estimand. This is total months in arrears, not the length of one uninterrupted episode. Peak severity is the maximum recorded due arrears on each path. Table 19 reports the unconditional 99th percentile of that maximum in arbitrary monetary units; it is not a confidence limit or a forecast of actual liabilities.

Table 19. Derived EP-01 duration and tail severity. R/D denotes retained/distributed surplus on the same 1,000 paths. Intervals concern paired mean duration differences; peak values are empirical path quantiles.

Scenario	Mean months R/D	Difference [95% interval]	Peak arrears, p99 R/D
Base	0.122/0.143	0.021 [0.010, 0.035]	44.5/59.4
Care cost	5.779/7.793	2.014 [1.811, 2.230]	513.3/545.8
Cash lock	1.758/2.135	0.377 [0.333, 0.421]	174.8/193.4
TIF impairment	0.183/0.225	0.042 [0.019, 0.069]	82.8/96.5

The largest additional mean arrears duration occurs under the care-cost shock (Table 19). The intervals quantify Monte Carlo variation conditional on the authored model, excluding uncertainty about healthcare prices, delayed claims, behavior, policy selection and real-world transport. No multiplicity-adjusted confirmatory claim is made. Distribution remains an institutional objective with a modeled resilience cost; these scenarios do not determine the preferred policy for an actual Pool.