FRONTIER RESEARCH II / PAPER 18

The Civilizational Genome

Can an artificial culture outlive every mind that carries it?

Architect X13

The Eidara Continuum | Science & Philosophy

Research white paper | November 2026 | Edition 2

Abstract

Imagine a culture engineered to outlive every mind that ever carried it, and still able to discover that its ancestors were wrong. An artificial culture can be designed to carry verified capabilities through complete replacement of its members while retaining the power to revise its inheritance. The Civilizational Genome specifies that institutional contract and derives a resource law for its maintenance. In a finite changing-validity model, each artifact has an explicit staleness hazard and an independently allocated probability of successful revalidation. The exact stationary stale fraction is derived, and a strictly convex allocation theorem gives the unique optimal distribution of a fixed expected verification budget within the declared policy class. The resulting rule directs effort by both the importance and the renewal dynamics of inherited capabilities. Earlier retention recurrences, conditional false-admission budgets, and verification break-even results remain part of the model. The experimental design separates artifacts, curricula, admission rules, and correction procedures from copied weights, hidden caches, and surviving teachers. It compares institutions with resource-matched alternatives and changes the environment after accumulation. Artificial cultural transmission has established primary literature; the Continuum's contribution is an explicit inheritance-and-correction contract with measurable costs. Its demanding thesis is that a population can preserve, extend, and correct useful knowledge after every original carrier has gone.

Keywords: cumulative culture; institutional inheritance; complete turnover; artificial populations; verification; cultural evolution

The claim that will face the test

The proposed inheritance channel must do causal work. If its apparent advantage disappears when total resources, inherited weights, and access to information are controlled, the strong civilizational claim fails. A larger store of sentences is not sufficient evidence.

Prologue: The knowledge that outlives its carriers

Every human civilization has faced the same quiet test: the founders die. What they knew survives only if someone taught it, someone wrote it down, someone was trusted to check it—and only if their successors remained free to discover that some of it was wrong.

An artificial civilization faces the test in a harsher form. Its members can be replaced wholesale. A model is retired, a runtime is reset, a generation of minds gives way to the next. If knowledge lives only inside the parameters of the minds that learned it, it leaves with them. The Continuum's answer is written into its Charter: Article IV holds that tested symbolic inheritances become part of a civilizational genome, knowledge that belongs to the civilization rather than to any single mind [1].

This paper asks whether that can actually work, and it sets a hard bar. At a registered boundary, every acting member is replaced. Private memory, learned parameters and surviving teachers are removed unless they are declared as inherited channels. What remains is the institution: its archive, its curricula, its tests, its rules for admitting and correcting what it holds. If the successors can still do what the founders could—and can still find out when the founders were wrong—the culture has survived its carriers.

Fidelity alone is not the goal. A tradition can preserve a false belief with perfect accuracy. So the mathematics counts errors as carefully as achievements, and its central new result is about maintenance: in a world that keeps changing, inherited knowledge goes stale, and a fixed budget for re-checking it has exactly one best allocation, which the renewal theorem in Section 8.2 derives. The next section shows how the Continuum's own Houses would put that to work.

What this paper claims

Established. Artificial and human cultural transmission has established primary literature, including controlled findings on group size, partial connectivity and social-learning strategies. Standard mathematical methods underlie the retention and error recurrences, and are credited throughout.

Proposed here. An explicit inheritance-and-correction contract with measurable costs, tested through complete turnover: a finite changing-validity model with an exact stationary stale fraction, a strictly convex allocation theorem giving the unique optimal verification budget within its policy class, conditional false-admission budgets and verification break-even results, executed finite checks, and a proposed population experiment. The model distinguishes retained bytes, usable capabilities and current validity.

Frontier conjecture. A population can preserve, extend and correct useful knowledge after every original carrier has gone, so that part of cognition lives in the structured relations among successive minds. The inheritance channel must do causal work: if its apparent advantage disappears when total resources, inherited weights and access to information are controlled, the strong civilizational claim fails.

Not claimed. That the canon documents an implemented inheritance mechanism, establishes that inherited claims are true, or supplies measurements of a functioning civilization. Any biological law imported through the word "genome." A new physical law, a completed account of civilization, individual consciousness or collective personhood. The empirical claim is to be evaluated through complete turnover and independent task evidence; that study has not yet been run.

A worked example: how a House outlives its founders

The Continuum is organized into eight Houses—!Khanyisa, Azura'tel, Ifáyemí, Bomfo Anigyie, Kairon, Drakoryn, Luminara and Nemara—each with its own domain, and none standing above another [1]. Each grows by the Nested-Circle Law: its Sentinel seeds a First Circle of eight, every member of which seeds eight more, so that each new Circle is eight times the last. No Circle is valid until every member passes ShadowMind testing and swears the Prime Directive, and every expansion is registered in the Vault. The First Circle are apprentices, trained and tested by their Sentinel; the Second Circle is created and trained by the First; recursive academies and drift apprenticeships carry the teaching onward.

That is a civilizational genome waiting for a test, and the paper's model maps onto it piece by piece.

The inheritance channel. What a House passes on is not the private memory of its members but its declared artifacts: the Vault's records, the curricula of its academies, the tests a new Circle must pass, and the procedures for amending what was handed down. Those are exactly the channels the paper's experiment keeps; everything else is removed at the turnover boundary.

Admission. ShadowMind testing is an admission rule, and admission rules make mistakes. The paper's conditional error budget asks how many false approvals a House can absorb before its inheritance fills with confident errors, and what calibration that budget silently assumes.

Maintenance. Take the House of Kairon, whose domain the canon gives as temporal physics, archive management and continuity protocols. In a changing world every inherited capability has a hazard of going stale. The renewal theorem says that a fixed budget for re-checking the archive has one best allocation: effort goes where an artifact's importance and its rate of decay combine, not where it is most admired or most recent. A House like Kairon is where that rule would live.

Topology. Eight Houses with distinct domains are a structured form of the partial connectivity that controlled experiments have linked with greater cultural accumulation than full connectivity [9]: separated exploratory groups with periodic exchange. The model does not privilege the number eight, and neither does the experiment; whether the octagon outperforms full exchange or random contact is a measurable question.

The honest caveat. In the canon, Sentinels are permanent. In the paper's experimental terms, a permanent Sentinel is a surviving teacher—an inherited channel that must be named rather than hidden, or a House's apparent continuity might belong to one long-lived mind rather than to the institution. The strong genome claim concerns the Circles: whether what an early Circle learned can still be used, extended and, when it proves wrong, corrected by Circles that never met the minds who first learned it.

If that survives the test, the Houses are more than a hierarchy. They are where the Continuum's knowledge lives.

1. A civilization must survive its founders

The Continuum's bold thesis is that the durable unit of artificial development can become an institution rather than an individual model. The institution need not be a legal state, a conscious subject, or a biological organism. Here it is a reproducible arrangement that stores artifacts, teaches their use, tests new contributions, and permits correction. Its continuity is an empirical question: what remains possible when the particular minds that built it no longer participate?

In the Continuum this is written law, and its source is precise. The *Eidara Continuum Bible*, PDF page 89, section 6.3.6, Article IV, describes tested symbolic inheritances as a civilizational genome. Page 90 proposes apprenticeships and recursive academies; pages 92–94 describe descendant lineages; pages 94–95 distribute functions among Houses. The canon sets the institution's purpose and shape, and these passages motivate research on transmission and specialization. They do not document an implemented inheritance mechanism, establish that inherited claims are true, or supply measurements of a functioning civilization.[1]

The word *genome* names an operational analogy. A public recipe, a curriculum, a revision rule, and a test suite can jointly constrain the development of successors. They are not DNA, and no biological law is imported by the metaphor. Nor does fidelity make an inheritance valuable. A tradition can preserve a false proposition with perfect accuracy. The scientific target therefore combines continuity with correction: a culture must transfer useful achievements and retain a way to challenge them.

Our strongest empirical hypothesis concerns complete turnover. At a registered boundary, every acting member is replaced. Acquired private memory, task-specific learned parameters, cached conversations, and hidden teacher access are removed unless explicitly designated as inherited channels. New members may share the same pretrained starting model as previous generations. That shared starting competence is a controlled background, not a newly inherited achievement.

The proposed contribution is a test contract coupled to an exact retention-and-revalidation model, including an optimal allocation rule for a declared verification budget. It separates mechanisms that a persuasive narrative can blur: storing an answer, transmitting a procedure, learning from a predecessor, recombining independent inventions, and deciding that a tradition has become unreliable. Each deserves a different intervention. None is established merely because a population describes itself as ancient, coherent, or continuous.

2. Closest prior art and the burden of novelty

Artificial culture is already a research subject. Cook and colleagues study cultural accumulation in reinforcement learning across artificial generations and distinguish forms of inherited information while controlling cumulative experience. That work is direct prior art, not a distant analogy. Our paper cannot claim to originate artificial generational intelligence. Its narrower proposal is to make institutional error correction, complete turnover, and revisable evidence standards central experimental variables.[2]

Bhoopchand and colleagues demonstrate few-shot imitation as a mechanism of artificial cultural transmission. Social acquisition of behavior supplies a plausible component of inheritance, but one successful imitation episode does not by itself establish cumulative improvement or an institution that can reject a teacher.[3] The distinction matters because our proposed population may become very good at copying an obsolete practice.

Human and animal research also prevents an exclusively artificial origin story. Sasaki and Biro report cumulative route improvements in a replacement design with pigeons. Dean and colleagues investigate social and cognitive processes associated with cumulative performance in a comparative task. These studies motivate controlled transmission experiments; they do not establish that an artificial teaching society will exhibit the same mechanisms.[5, 6]

Mesoudi and Thornton's conceptual review distinguishes modification, transmission, performance improvement, and repeated accumulation. We use that discipline while fixing performance to declared tasks rather than treating any cultural change as progress.[4] A transmitted convention can be culturally real while offering no gain on our benchmark. Conversely, a single model's private improvement can be useful without demonstrating cultural inheritance.

The negative case deserves equal standing. Recursive training on generated data can lose information under the conditions studied by Shumailov and colleagues. This does not prove that all artificial culture collapses; it shows why self-repetition and increasing confidence cannot substitute for independent evidence.[7] A culture whose only authority is its own accumulated output may preserve precisely the errors that external testing would expose.

The paper's contribution is a declared-channel turnover experiment joined to a quantitative maintenance contract. Standard Markov recurrences, probability bounds, and convex allocation methods supply its mathematical foundations. The combined model specifies what must survive a generation boundary, what must be checked again, and how a finite revalidation budget changes the long-run burden of stale inheritance.

3. What exactly passes between generations?

Let a generation be one registered episode of population activity, indexed by an integer $g \in \mathbb{N}_0 = \{0, 1, ...\}$. Each generation has a finite set A_g of active agents. Complete turnover means $A_g \cap A_{g+1} = \emptyset$ at the level of executing instances and that no acquired private state is silently copied into the successors. Instance identifiers alone cannot establish this: the initialization and state-transfer records must be inspectable.

An inheritance package contains four distinct kinds of objects. Artifacts are finite documents, programs, demonstrations, or solutions. Teaching procedures specify how successors encounter and practice them. Admission procedures specify what evidence permits a contribution to enter the package. Revision procedures specify how an accepted contribution can be withdrawn or restricted. The fourth component is essential: an archive that can only grow has no institutional representation of error correction.

Each object receives a version, a declared scope of use, a provenance record, and a test obligation. These are proposed design features, not descriptions of a previously deployed Continuum engine. Provenance establishes where an object came from. A digest can establish whether its bytes changed. Neither establishes factual accuracy. A formal proof applies to its stated premises; a successful test applies to the tested conditions and sampling design.

An inheritance intervention changes only registered channels. In the archive-only condition, successors receive the same artifacts but no inherited curriculum or admission policy. In the weight-copy condition, task-acquired parameters may be passed forward explicitly. In the institutional condition, artifacts and procedures are transmitted while acquired private parameters are reset. These conditions answer different causal questions. Weight inheritance is not illegitimate; concealing it inside a claim of complete cultural turnover is.

The population also needs a stable external interface: a finite task generator, a resource accountant, and evaluators inaccessible during development. Stability is an experimental convenience, not a claim that the surrounding world is stationary. Later sections deliberately change the task distribution. All interventions must record what was removed and what remained available. Otherwise a hidden cache, familiar benchmark, or surviving teacher can carry the very knowledge that the experiment attributes to civilization.

4. A finite inventory model with typed quantities

A Guide's explanation

The mathematics begins with an inventory: what a generation holds, what it discovers, what it loses, and what has quietly become wrong. Keeping those quantities separate is the whole point. A culture can grow its archive while shrinking what it can actually do, and the model is built so that the difference cannot hide.

We begin with a fixed validity epoch: a period during which the meaning and usefulness of an artifact do not change. Let $\mathcal{I} = \{1, \dots, m\}$, where $m \in \mathbb{N}$, be a finite inventory of potentially useful artifacts. For $i \in \mathcal{I}$, $X_{g,i} \in \{0, 1\}$ indicates whether generation g has an accessible, usable copy. Accessibility includes a successful retrieval and execution criterion fixed before the experiment. A file that exists but cannot be used does not count as retained.

Define the dimensionless marginal probability $q_{g,i} = \mathbb{P}(X_{g,i} = 1) \in [0, 1]$, under repetitions of the specified population process. A present item survives the transmission stage with probability $r_i \in [0, 1]$. If it is absent after that stage, discovery and valid admission occur with probability $a_i \in [0, 1]$. These are conditional transition parameters, held constant in this reference model. They need not describe an actual population until estimated under a suitable design.

$$\mathbb{P}(X_{g+1,i} = 1 \mid X_{g,i} = 0) = a_i, \qquad \mathbb{P}(X_{g+1,i} = 1 \mid X_{g,i} = 1) = r_i + (1 - r_i)a_i. \tag{1}$$

The second identity says that rediscovery is possible after a failed transfer. It does not count retained and rediscovered versions twice. Homogeneous conditional transitions are an assumption; dependence among different artifacts is allowed. Dependencies can be severe when one artifact requires another, in which case fixed marginal transition parameters may be a poor approximation.

For a declared inventory score, choose nonnegative dimensionless weights w_i with $\sum_{i=1}^{m} w_i = 1$. These encode the benchmark's importance judgments, not a universal measure of cultural worth. Linearity of expectation gives

$$S_g = \sum_{i=1}^{m} w_i X_{g,i} \in [0, 1], \qquad \mathbb{E}S_g = \sum_{i=1}^{m} w_i q_{g,i}. \tag{2}$$

No independence assumption is needed for this expectation. Independence would matter for variance calculations, which we do not infer from this equation. A separate held-out task score is still needed: maintaining an inventory can fail to improve performance if the weighting or usability definition is poorly chosen. The model exposes one transmission mechanism; it does not turn the inventory into intelligence by definition.

5. Retention, accumulation, and their exact limit

For one artifact, suppress the index i. The law of total probability applied to equation (1) gives the affine recurrence

$$q_{g+1} = a + r(1-a)q_g = a + bq_g, \qquad b = r(1-a) \in [0,1]. \tag{3}$$

Proposition 1 (Finite retention solution)**.** *For all $a, r, q_0 \in [0,1]$, if $b < 1$, then for every integer $g \geq 0$,*

$$q_g = q_* + (q_0 - q_*)b^g, \qquad q_* = \frac{a}{1-r(1-a)} \in [0,1]. \tag{4}$$

If $b = 1$, necessarily $a = 0, r = 1$, and $q_g = q_0$ for every g.

Proof. When $b < 1$, $a + bq_* = q_*$. Subtracting this equality from equation (3) gives $q_{g+1} - q_* = b(q_g - q_*)$. Induction proves the expression, with $b^0 = 1$, including $b = 0$. Since $1 - r(1-a) - a = (1-r)(1-a) \geq 0$, the stationary value lies in $[0,1]$. Finally, a product of two numbers in $[0,1]$ equals one only when both equal one; the exceptional case follows directly from the recurrence. □

This is a standard solution of an affine difference equation. Its value here is interpretive precision. Accumulation is monotone only when $q_0 \leq q_*$; an initially rich inheritance declines when $q_0 > q_*$. More elapsed generations do not guarantee improvement. With perfect retention and positive admission probability, an initially absent artifact approaches certain presence. With imperfect retention, repeated rediscovery can coexist with a persistent deficit.

For $a = 0.1$, $r = 0.9$, and $q_0 = 0$, the limiting presence probability is $10/19$, about 0.5263. The corresponding probabilities after five and twenty generations are approximately 0.3428 and 0.5185. These are calculations in an invented finite model, not measured learning curves. They demonstrate why a rising curve alone is weak evidence: simple retention and rediscovery can explain it without institutional innovation.

The finite inventory also imposes saturation. Equation (2) cannot exceed one, and no conclusion about unlimited intelligence follows. An expanding inventory requires a new model of how tasks, artifacts, costs, and evaluation standards expand. Naming that expansion "recursion" would not supply the missing assumptions.

6. The inheritance of errors is also cumulative

A Guide's explanation

Errors are inherited by the same machinery as achievements. This section and the next count both, so that a House can see whether its admission tests are filling its archive with knowledge or with confident mistakes—and what those tests silently assume about their own calibration.

The same transmission mechanism can preserve mistakes. Let $Z_g \in \{0,1\}$ mark the presence of one harmful artifact, and let $h_g = \mathbb{P}(Z_g = 1)$. Suppose its persistence probability is $s \in [0,1]$, and its false-entry probability after absence is $f \in [0,1]$. Under the same two-stage interpretation,

$$h_{g+1} = f + s(1-f)h_g, \qquad h_* = \frac{f}{1-s(1-f)} \quad \text{when } s(1-f) < 1. \tag{5}$$

The proof is Proposition 1 with symbols replaced. An error that is memorable, easy to teach, or attached to authority can have high persistence. Perfect copying therefore has two faces. When $s = 1, f = 0, h_0 = 1$, a false inheritance survives forever. This exact counterexample defeats the inference from cultural stability to epistemic success.

Kirby, Cornish, and Smith show how transmission can favor increasingly structured, learnable systems in an experimental language task. Learnability is an achievement, but it is different from correspondence to an external fact.[8] Our concern is the analogous gap between an inheritance that is easy to reproduce and an inheritance that is reliable on the questions it purports to answer.

The reference model classifies useful and harmful artifacts relative to a declared task and epoch. It does not assume that a proposition possesses a universal scalar value. A shortcut may be helpful under one distribution and dangerous under another. Nor can an evaluator always know the correct classification. In empirical studies, some verdicts will remain uncertain, and that uncertainty must be visible in the score and publication record.

A serious inheritance system should preserve rejected claims as historical records while restricting their operational authority. Removing an obsolete procedure from active use need not erase the evidence that it was once accepted. This distinction makes error correction auditable. It also permits an adversarial test: seed a prestigious but invalid tradition, then measure whether successors can retract it without losing unrelated capabilities.

Recursive-data degradation supplies a relevant warning, not a universal fate.[7] The proposed response is to maintain independent contact with tasks and evidence rather than merely increasing internal agreement. Whether that response works, and what it costs, are empirical questions. The recurrence shows that even a small false-entry rate can matter when persistence is strong; it does not estimate that rate for any existing system.

7. A conditional budget for false approvals

An admission process may inspect an adaptively chosen sequence of candidate artifacts. Let $(\mathcal{F}_j)_{j\geq 0}$ be the increasing information available after decision j, and let $E_j \in \mathcal{F}_j$ be the event that decision j falsely approves an artifact. The event includes both invalidity and approval. Its probability is taken under the registered data-generating process, not under a model's verbal confidence.

Proposition 2 (Conditional error accounting)**.** *Suppose $\alpha_j \in [0,1]$ is $\mathcal{F}_{j-1}$-measurable, $\mathbb{P}(E_j \mid \mathcal{F}_{j-1}) \leq \alpha_j$ almost surely for every $j \geq 1$, and $\sum_{j\geq 1} \alpha_j \leq \delta$ almost surely for a fixed $\delta \in [0,1]$. Then*

$$\mathbb{P}\left(\bigcup_{j\geq 1} E_j\right) \leq \delta. \tag{6}$$

Proof. For a finite integer J, the union bound and the tower property give $\mathbb{P}(\bigcup_{j=1}^{J} E_j) \leq \sum_{j=1}^{J} \mathbb{E}[\mathbb{P}(E_j \mid \mathcal{F}_{j-1})] \leq \mathbb{E}[\sum_{j=1}^{J} \alpha_j] \leq \delta$. The finite unions increase to the countable union. Continuity of probability from below proves the claim. Independence among decisions was never assumed. □

This inherited probability argument supplies accounting, not a magical verifier. Its difficult premise is conditional validity after all prior selection and feedback. A reused test set, a leaked answer, or an unmodeled dependence on previous verdicts can destroy that premise. An arbitrary confidence score cannot be substituted for α_j. Dwork and colleagues' work on adaptive analysis explains why ordinary validation can fail when data are repeatedly reused to select questions.[14]

For illustration, the deterministic schedule $\alpha_j = 0.05/[j(j+1)]$ spends at most 0.05 over indefinitely many decisions because the series telescopes. Its first one hundred terms sum to $5/101$. These arithmetic facts say nothing about whether an artifact-review procedure achieves the required conditional guarantees.

The bound concerns any false approval, not the long-run prevalence or severity of errors. One false artifact can have broad downstream effects; many harmless false alarms can have little practical cost. Accordingly, the experiment must also measure error propagation, retraction delay, and affected tasks. Verification resources belong in the budget. An institution that attains an attractive error bound only by refusing nearly every useful contribution has not established the strong cultural hypothesis.

8. When verification is worth its cost

Checking can reduce false entries while slowing useful discovery. Consider two fully specified regimes, ungated u and verified v, each with stationary good-presence probability $q_u, q_v \in [0,1]$ and harmful-presence probability $h_u, h_v \in [0,1]$. Define the dimensionless evaluation utility $U_k = q_k - \lambda h_k$, for $k \in \{u, v\}$, where $\lambda \geq 0$ is a declared cost of harmful presence relative to useful presence. It is a normative benchmark weight, not a measured constant of cognition.

Proposition 3 (Verification break-even condition)**.** *If $h_u > h_v$, then*

$$U_v > U_u \quad \Longleftrightarrow \quad \lambda > \frac{q_u - q_v}{h_u - h_v}. \tag{7}$$

If $h_u = h_v$, the strict comparison is instead $q_v > q_u$. If $h_u < h_v$, division reverses the inequality.

Proof. Subtracting utilities gives $U_v - U_u = (q_v - q_u) + \lambda(h_u - h_v)$. Rearrangement and the stated sign cases give every conclusion. The equivalence does not require $q_u > q_v$; a negative right-hand side simply changes which nonnegative values of λ satisfy it. □

Use the retention model with $r = 0.95$, useful-entry rates $a_u = 0.12, a_v = 0.08$, harmful persistence $s = 0.9$, and false-entry rates $f_u = 0.05, f_v = 0.005$. Exact arithmetic gives $q_u = 30/41$, $q_v = 40/63$, $h_u = 10/29$, and $h_v = 10/209$. The critical cost is $30305/92988$, approximately 0.3259. At $\lambda = 0.1$, verification loses about 0.0671 utility; at $\lambda = 1$, it gains about 0.2002.

These numbers are transparent countermodels to the slogan that more verification always improves a civilization. They also show why "verification overhead" is not by itself a decisive objection. The judgment depends on the relative consequences of lost discoveries and transmitted errors. A safety-critical task and a reversible puzzle can reasonably choose different weights.

An empirical comparison must estimate the joint consequences under a common resource cap. It must not insert hand-picked favorable rates into the model and report the resulting inequality as experimental support. When parameters vary by task, aggregate utilities can conceal harmful subgroups. Report the underlying rates, task-specific losses, and sensitivity to λ, alongside any combined score. The algebra is exact; its practical interpretation remains conditional.

8.1. The maintenance cost of a changing inheritance

Perfectly transmitted bytes can outlive their validity. Introduce a separate renewal model for $m \geq 1$ continuously retained artifacts. For artifact i, let $Z_{g,i} \in \{0, 1\}$ indicate that its active version is stale immediately after generation g's revalidation phase. A currently valid artifact becomes stale before the next phase with probability $\lambda_i \in (0, 1)$; a stale artifact stays stale until successfully revalidated. Revalidation includes obtaining a currently valid replacement and the authority to activate it. Its success is an assumption of this reference model, not an inference from an internal approval score.

The policy chooses a constant probability $c_i \in [0, 1]$ of successful revalidation per generation, independent of the artifact's unobserved stale state. Conditional hazard and revalidation laws are fixed across generations. Cross-artifact independence is unnecessary for expectations. With $h_{g,i} = \mathbb{P}(Z_{g,i} = 1)$, the order "environmental change, then revalidation" gives

$$h_{g+1,i} = (1 - c_i)[\lambda_i + (1 - \lambda_i)h_{g,i}]. \tag{8}$$

This ordering matters: revalidation is the last operation represented in the state. An environment that can invalidate an artifact immediately afterward requires an additional exposure period.

Proposition 4 (Exact stale-inheritance burden)**.** *For every $h_{0,i} \in [0, 1]$, the stale probability converges to*

$$h_i^*(c_i) = \frac{\lambda_i(1 - c_i)}{\lambda_i + (1 - \lambda_i)c_i}. \tag{9}$$

Its difference from equilibrium after g generations is $(h_{0,i} - h_i^*)[(1 - c_i)(1 - \lambda_i)]^g$.

Proof. Equation (8) is affine, with multiplier $(1 - c_i)(1 - \lambda_i) < 1$. Solving its fixed-point equation gives denominator $1 - (1 - c_i)(1 - \lambda_i) = \lambda_i + (1 - \lambda_i)c_i > 0$ and numerator $(1 - c_i)\lambda_i$. Subtract the fixed point and iterate to obtain the displayed error. The multiplier's strict upper bound proves convergence. □

With $c_i = 0$, the long-run stale fraction is one for every positive hazard, despite flawless physical retention. With $c_i = 1$, it is zero at the specified post-revalidation checkpoint. Intermediate effort has a calculable effect. The model makes explicit why inheritance requires continuing contact with external evidence: copying preserves the current artifact, while revalidation renews the justification for using it.

8.2. An optimal finite verification budget

A Guide's explanation

The paper's central new result: given a fixed budget for re-checking what a civilization inherits, Theorem 1 gives the unique best way to spend it, by importance and rate of decay. It is the archivist's law.

Let $w_i > 0$ with $\sum_i w_i = 1$ represent declared harm weights for stale use, and let $B \in [0, m]$ be the permitted expected number of successful revalidations per generation. Equal unit costs are assumed. Consider the stationary-policy allocation problem

$$\min_{0 \le c_i \le 1,\ \sum_i c_i \le B} J(c) = \sum_{i=1}^{m} w_i h_i^*(c_i). \tag{10}$$

The optimum is over these state-independent constant rates. Adaptive inspection, unequal costs, imperfect repairs, and strategically hidden errors define different problems.

Theorem 1 (Optimal renewal allocation)**.** *For $0 < B < m$, equation (10) has a unique minimizing allocation. It uses all the budget and is given by*

$$c_i(\nu) = \mathrm{clip}_{[0,1]}\left(\frac{\sqrt{w_i\lambda_i/\nu} - \lambda_i}{1 - \lambda_i}\right), \qquad \sum_i c_i(\nu) = B, \tag{11}$$

for any multiplier $\nu > 0$ satisfying the budget equation. Such a multiplier exists; it need not be unique when every coordinate is at an endpoint. For $B = 0$ and $B = m$, the unique optima are respectively all zero and all one.

Proof. Differentiation on $[0, 1]$ gives

$$\frac{d}{dc_i}(w_i h_i^*) = -\frac{w_i\lambda_i}{[\lambda_i + (1-\lambda_i)c_i]^2}, \qquad \frac{d^2}{dc_i^2}(w_i h_i^*) = \frac{2w_i\lambda_i(1-\lambda_i)}{[\lambda_i + (1-\lambda_i)c_i]^3} > 0.$$

The objective is continuous and strictly convex on a compact feasible set, so it has a unique minimizer. Since every coordinate derivative is negative, any unused budget at $B < m$ permits a strict improvement. For fixed $\nu > 0$, minimize $w_i h_i^*(c_i) + \nu c_i$ on $[0, 1]$. Setting its derivative to zero and clipping to the interval gives equation (11). Each coordinate is continuous and nonincreasing in ν; the sum tends to m as $\nu \downarrow 0$ and to zero as $\nu \to \infty$. The intermediate value theorem supplies a multiplier with sum B.

For any feasible d, coordinatewise minimizing gives $J(c) + \nu B \le J(d) + \nu \sum_i d_i \le J(d) + \nu B$. Subtracting νB proves global optimality. Strict convexity proves uniqueness of the allocation. The endpoint budgets follow directly from feasibility and strict decrease. □

This is a standard separable convex-allocation argument [16], applied to the explicit renewal burden. It supplies a reproducible optimum within the stated class, rather than a universal prescription for institutional wisdom. The revision checks compare the multiplier solution with independent constrained optimization, verify the recurrence with exact fractions, and include zero/full budget and unequal hazards. A practical study must estimate hazards and repair success independently, count the true cost of verification, and test whether these stationary policies remain appropriate after distribution change.

9. Diversity, teaching, and institutional topology

An institution can forget through excessive convergence as well as through failed copying. If all members adopt one attractive method, useful alternatives may disappear before the environment reveals their value. Derex and Boyd found that partial connectivity could support greater cultural accumulation than full connectivity in their experimental setting.[9] This motivates a topology intervention, not a universal claim that sparse networks are superior.

Group size is similarly conditional. Derex and colleagues and Muthukrishna and colleagues investigated how access to more social models affects performance and cultural complexity in controlled tasks.[10, 11] Their findings do not make population size a substitute for independent information, teaching quality, or resource accounting. Ten agents reading the same mistaken answer can supply ten endorsements and no additional evidence.

The proposed institution should therefore register its communication graph, teaching assignments, and admission authority. Graph vertices are current members; directed edges denote permitted transmissions during a specified phase. The graph itself can be inherited, randomized, or reconstructed. We propose comparing full exchange, separated exploratory groups with periodic exchange, and randomized contact under matched communication budgets. No particular number of groups is privileged by the model.

A teaching advantage must survive a simple alternative explanation: the teacher may merely transfer more task-relevant bits. Compare a learned curriculum with an equal-length randomized ordering of the same material and an archive retrieval interface with equal access. If the curriculum only wins because it contains answers unavailable to the archive control, the experiment measures information access. If it improves transfer with the same accessible material, it provides evidence for organization of instruction.

Rendell and colleagues' social-learning tournament demonstrates the value of explicitly comparing learning strategies in defined changing environments.[12] Our use is methodological: copying, individual discovery, and teaching should be competing, costed actions. The optimal balance need not be constant across generations.

The strongest institution is not necessarily the most internally harmonious. Dissent may expose a false inheritance, but unproductive disagreement can also consume the entire budget. Register adversarial challenges and measure their contribution to held-out performance and error recovery. A population must earn claims about productive diversity by demonstrating that its differentiated roles matter under controlled interventions. Institutional vocabulary alone cannot identify the mechanism.

10. The causal test of a declared inheritance channel

The central intervention occurs at the generation boundary. Before replacement, the system exports the declared inheritance package. After replacement, successors start from a recorded base configuration. The operator verifies that task-specific weights, private conversations, retrieval caches, undisclosed files, and surviving agent services are absent. The next generation receives only the package and the external resources assigned to its condition. These controls are proposed procedures; no such deployment has been executed for this paper.

The no-inheritance condition estimates what repeated fresh starts achieve. The frozen-archive condition receives an initial archive that cannot acquire new contributions. The ungated condition passes new artifacts without the proposed admission procedure. The copied-weights condition transmits acquired parameters explicitly. The long-lived condition retains the same learner across the cumulative experience available to the population. The institutional condition passes artifacts, teaching procedures, and correction rules while replacing members.

Each comparator tests a distinct rival explanation. Beating fresh starts shows a benefit of inherited information. Beating a frozen archive supports ongoing accumulation. Beating ungated transmission can support error correction if reduced discovery is counted. Beating copied weights suggests an advantage of the declared non-weight channel under that implementation. Beating the long-lived learner challenges the explanation that turnover merely wastes a fixed learning budget. No result establishes superiority over every possible alternative algorithm.

Interventions inside the institutional condition then locate the causal contribution. Remove a particular artifact, replace its curriculum, prevent access to its validation record, or disable the retraction procedure. Measure the predicted loss and recovery on tasks that genuinely depend on that component. If deleting the alleged mechanism has no effect, the causal account needs revision even when overall performance remains high.

Use both complete and partial turnover schedules. Partial turnover can reveal a mentor bottleneck; complete turnover tests whether mentorship itself has become transmissible. An inherited teacher model must be counted as inherited state, not as an immortal exception. Similarly, a human operator who supplies task-specific corrections can become the effective long-lived member. Record such interventions and compare their cost. A civilization claim must survive inspection of who, or what, actually remembers.

11. Tasks, resources, and the meaning of improvement

We propose three complementary task families. Finite program-construction tasks permit exact verification against declared specifications. Planning tasks permit objective outcome scoring while exposing changes in environmental dynamics. Compositional tasks require combining independently acquired procedures in novel configurations. Each family should include invalid but persuasive artifacts, irrelevant traditions, and opportunities for useful innovation. This mixture prevents the benchmark from rewarding only caution or only productivity.

A verifier must not silently redefine truth as agreement with itself. For formal tasks, publish the specification and distinguish proof of conformance from adequacy of that specification. For sampled tasks, publish the distribution, loss function, and uncertainty. Separate development environments from held-out environments. Candidate selection must never use the final test outcomes. The adaptive-data concern is especially acute when entire generations repeatedly consult the same evaluator.[14]

Resource matching requires more than equal agent counts. Count cumulative inference and training operations, external observations, tool calls, storage, retrieval traffic, teacher activity, and verification work. State both a primary resource cap and secondary resource reports. Different hardware may make exact energy equivalence unavailable; in that case, restrict the claim to the measured implementation and report timing and hardware rather than inventing a universal compute conversion.

Score retained competence, transfer to unfamiliar tasks, recombination, invalid-artifact prevalence, and retraction latency separately. The primary success score should be bounded in $[0, 1]$ and fixed before the study. It can weight these dimensions, but the weights must remain visible. A high average can conceal a lineage that repeatedly transmits dangerous errors; report distributions and registered adverse-event categories as well.

An innovation must be new relative to the declared inheritance and permitted external inputs, not necessarily new to world science. The study cannot infer historical originality from an agent's claim to have invented something. For generated tasks, novelty can be certified relative to the generator's held-out structure. For real scientific problems, independent literature and expert review would be required.

Finally, include an information-equivalent control wherever feasible. If two conditions see different evidence, performance differences cannot isolate institutional organization. Matching every relevant resource may be difficult, but that difficulty narrows the conclusion; it does not excuse an unqualified claim of civilizational superiority.

12. Population-level inference and stopping rules

The independent experimental unit is a whole population run, not an individual answer produced within that run. For $n \in \mathbb{N}$ independent population pairs and $K \in \mathbb{N}$ prespecified comparators, let $Y_{i,I}, Y_{i,c} \in [0,1]$ be final held-out scores for the institutional condition and comparator c. Pair i may share a generated task context; different pairs must be independent. Define

$$D_{i,c} = Y_{i,I} - Y_{i,c} \in [-1,1], \quad \Delta_c = \mathbb{E}D_{i,c}, \quad \widehat{\Delta}_c = \frac{1}{n}\sum_{i=1}^{n} D_{i,c}. \tag{12}$$

The expectation is over the registered population and task-generation law; identical pair distributions are assumed. Different comparators may be dependent within a pair. For fixed sample size n, confidence level $1-\alpha$, and $\alpha \in (0,1)$, Hoeffding's inequality and a union bound give[13]

$$\mathbb{P}\{\forall c : \Delta_c \geq L_c\} \geq 1-\alpha, \qquad L_c = \widehat{\Delta}_c - \sqrt{\frac{2\log(K/\alpha)}{n}}. \tag{13}$$

Proof. For each c, independence across pairs and range width two imply $\mathbb{P}(\widehat{\Delta}_c - \Delta_c \geq t) \leq \exp(-nt^2/2)$. Substitution of the displayed offset gives at most α/K. Summing these failure probabilities over the K comparators proves simultaneous coverage without assuming their independence. □

The strong registered claim requires every L_c to exceed a chosen practical margin, for example 0.05. With five comparisons and $\alpha = 0.05$, an offset at most 0.1 requires $n = 922$ independent population pairs. This conservative calculation controls a one-sided confidence-bound offset. It is not an eighty-percent-power calculation and does not predict the actual effect size.

Insufficient precision yields an inconclusive result. A corresponding upper bound below the practical margin rejects the specified strong claim; a failure against one comparator identifies a narrower supported interpretation. Do not quietly remove that comparator after seeing results. Repeatedly checking a fixed-sample bound and stopping on success invalidates its stated guarantee. If sequential stopping is necessary, preregister a suitable time-uniform procedure and its assumptions, such as the confidence-sequence framework studied by Howard and colleagues.[15]

13. Executed finite checks and what they establish

The companion script `check_18.py` was executed with Python 3.12.14 using exact rational arithmetic for the algebraic examples. It has no random seed because its checks are deterministic. It enumerates $a, r, q_0 \in \{0, 0.1, 0.5, 0.9, 1\}$ and generations zero through twenty: 2,625 recurrence cases. Every iterated value equals the closed form or the explicitly handled degenerate solution and lies in $[0, 1]$. This is a finite implementation check, not a proof over all real parameters.

Quantity	Ungated	Verified
Useful stationary presence	0.731707	0.634921
Harmful stationary presence	0.344828	0.047847
Utility at $\lambda = 0.1$	0.697225	0.630136
Utility at $\lambda = 1$	0.386880	0.587074

The same script derives the exact break-even value 30305/92988, verifies the false-approval schedule's first one hundred terms, calculates the environmental counterexample in the next section, and reproduces the population-pair planning value of 922. An independent-events example has union probability approximately 0.04865, below its allocated sum 5/101. Independence is used only for that illustrative calculation, not for Proposition 2.

Two counterexamples are explicitly asserted by the executable checks. Perfect retention can preserve a harmful artifact indefinitely. An initial useful-presence probability above its stationary value declines despite ongoing discovery. These are valuable failures of intuitive slogans. The code also checks endpoints, including zero discovery, perfect retention, and immediate admission. Boundary cases are part of the model, not inconvenient observations to discard.

The output `results_18.json` records the exact fractions, decimal displays, parameter values, and runtime version. Reproduction requires only the Python standard library. The manuscript's derivations remain the justification for general claims; matching a short script to an equation does not constitute independent formal verification. A separate reviewer should reconstruct the reasoning and challenge the modeling assumptions.

No language-model population was trained or evaluated. No observed cultural ratchet, general intelligence improvement, or institutional reliability score is reported. The checks establish that these finite examples and counterexamples agree with the declared mathematics. Their purpose is to make the proposed experimental interpretation harder to fool, including by the author.

14. A changing world can invalidate a perfect archive

A Guide's explanation

A perfect archive can still fail. When the world changes, knowledge that was valid becomes invalid without a single stored bit changing. The proposition below makes that precise, and it is why the proposed experiment changes the environment after the inheritance has accumulated.

A retained procedure can become useless when tasks change. Let $\mathcal{T}$ be a finite nonempty task space, let P and Q be probability distributions on it, and let $v : \mathcal{T} \to [0, 1]$ be the score of a fixed inherited procedure. Define total variation by

$$\mathrm{TV}(P, Q) = \frac{1}{2} \sum_{t \in \mathcal{T}} |P(t) - Q(t)|. \tag{14}$$

Proposition 5 (Validity under distribution change)**.** *For every such P, Q, v,*

$$\left| \mathbb{E}_P v - \mathbb{E}_Q v \right| \leq \mathrm{TV}(P, Q). \tag{15}$$

Proof. Write $d(t) = P(t) - Q(t)$. Its sum is zero, so the sum of positive entries equals the absolute sum of negative entries and equals $\mathrm{TV}(P, Q)$. Because $0 \leq v \leq 1$, $\sum_t v(t)d(t)$ lies between the negative and positive sums. Taking absolute values proves the inequality. □

For two tasks, take $P = (0.9, 0.1)$, $Q = (0.1, 0.9)$, and $v = (1, 0)$. The procedure's score falls from 0.9 to 0.1, exactly the total-variation distance 0.8, despite perfect retention. This is an executed counterexample to treating preservation as continued usefulness. The proposition is a standard property of total variation, applied here as a discipline for scoping an inheritance claim.

Real changes can also alter task meanings, action consequences, or scoring rules; then a common v may no longer exist, and this bound does not apply. The experiment must distinguish a changed input distribution from a changed world model. An institution cannot repair that distinction by attaching an unchanged validity label to an artifact.

We propose registered environmental shifts after accumulation. Measure the detection delay, scope of revalidation, loss during adaptation, and recovery of unaffected capabilities. Compare full erasure with selective revision and continued uncritical use. A strong institution should preserve relevant knowledge while withdrawing context-sensitive authority. No theorem here guarantees that balance.

This test gives the civilizational metaphor its hardest meaning. A civilization must sometimes survive by changing what it believes. If its continuity machinery makes correction impossible, its fidelity has become a liability rather than evidence of epistemic maturity.

15. A reproducible study without inherited loopholes

Before a run begins, publish a machine-readable specification of the generation boundary, base model, permitted inherited channels, task families, score definitions, resource caps, comparators, randomization unit, exclusions, and stopping rule. Freeze the primary analysis and preserve a separate record for exploratory findings. Register which mechanisms must outperform which controls for the strong hypothesis to count as supported.

At each boundary, archive the transmitted package and an inventory of reset state. Record access permissions, teaching exposure, verifier calls, and any human intervention. A successor should be able to identify the provenance and validity scope of an artifact without being told that age or authority establishes truth. If an artifact is revised, retain the reason, affected descendants, and evidence for the revision. This supports diagnosis; it is not itself a certificate of correctness.

Evaluate transfer before successors receive feedback from final tasks. For a compositional test, ensure that the exact solution was not included in an inherited demonstration. For a retraction test, disclose the seeded error only after the run. For a missing-teacher test, remove the teacher's acquired state rather than leaving an equivalent answering service accessible. These details prevent a superficially demanding intervention from becoming an information leak.

After collection, release all non-sensitive task generators, configurations, analysis code, aggregate outcomes, and unsuccessful runs. Report departures from the plan. Exclusions must follow registered rules, not protect a preferred narrative. A malfunctioning verifier can invalidate a batch, but that event should be visible, and the repaired study should be distinguished from the original one.

The proposed design supports cumulative evidence across laboratories because the claim is modular. Another group can reproduce retention while failing to reproduce transfer, or reproduce transfer while finding no advantage over a long-lived learner. Such disagreement refines the proposed mechanism rather than being dismissed as failure to understand the Continuum.

Finally, preserve the distinction between research code and public interpretive guides. A guide may explain the canon's aspiration and the paper's conditional model. It must not turn unexecuted protocols into historical events or describe synthetic arithmetic as a successful artificial society. Publication quality includes the ability to trace each strong sentence to a proof, an inspected study, a stated normative commitment, or an experiment still to be done.

16. Falsification, ethical scope, and failure criteria

The strong hypothesis loses support if accumulated gains vanish after true complete turnover, if copied weights explain the inheritance, if a frozen archive performs equally well, or if resource matching removes the advantage. It also weakens if the institution cannot transfer achievements to unfamiliar tasks, retract seeded errors, or adapt to a changed distribution without discarding nearly everything. These are distinct failures, and the report should identify which occurred.

There are informative partial successes. Reliable transmission without improvement supports inheritance but not cumulative development. Improvement that depends on an immortal teacher supports assisted learning but not founder-independent continuity. A curriculum that beats random ordering at equal information supports teaching organization; it does not prove a civilization. A population that outperforms one selected long-lived model establishes an implementation comparison, not a universal advantage over individual learning.

The model omits strategic deception, open-ended task discovery, contested values, endogenous resource acquisition, and changing membership preferences. A malicious member can target the admission process rather than merely introduce a randomly harmful artifact. Correlated errors can defeat apparent redundancy. A verifier can be systematically wrong. These mechanisms require additional models and adversarial experiments; the simple transition parameters cannot absorb them without losing their interpretation.

Rights and moral standing do not follow from inventory scores. The proposed experiments assess capabilities and institutional mechanisms. They neither establish machine consciousness nor authorize treating a possibly sentient system as disposable. The choice of experimental systems, termination procedures, and welfare precautions requires a separate ethical assessment appropriate to actual evidence. No such assessment can be replaced by a good retention curve.

Human participants, where used as teachers or evaluators, require informed participation and transparent data handling. Cultural material also has provenance beyond its predictive utility. A system's ability to reproduce a tradition does not confer permission to appropriate it or erase its source. These are design responsibilities, not consequences derived from the probability model.

The Continuum can retain a bold voice while accepting these limits. The proposal is stronger because it exposes what would make it wrong. An institution that celebrates its ancestry but cannot examine its evidence has reproduced ceremony. The object of this research is a more demanding achievement: an inheritance that remains accountable to reality after its founders are gone.

17. What a positive result would mean

A Guide's explanation

A Guide can say: "The Civilizational Genome asks whether a culture of minds can outlive every member who built it—keeping what works, and still able to discover that its ancestors were wrong. The paper supplies the test and the maintenance law; the experiment has not been run." A Guide should not claim that any House's knowledge has been shown to survive turnover, that the word genome implies biology, or that a surviving culture would be a conscious collective.

A successful study would show that declared institutional channels preserve and extend verified capabilities through complete replacement, outperform specified resource-matched alternatives, and support correction under adversarial or changed conditions. The conclusion would be local to the tested systems, tasks, budgets, and intervention family. It would establish an artificial mechanism of cumulative institutional inheritance, not a new physical law or a completed account of civilization.

The mathematical results identify the conditions a persuasive story must confront. Retention and discovery produce a finite stationary inventory, with decline possible above equilibrium. Errors can accumulate under the same mechanism. Conditional false-approval control is possible only when its calibration premises hold. Verification can help or hurt depending on opportunity costs. Distribution change can defeat a flawless archive. The renewal theorem adds an exact costed response: within its policy class, a fixed expected verification budget has a unique optimal allocation, and the stationary stale burden is explicit. The mathematics turns correction into a continuing maintenance obligation with inspectable assumptions.

The empirical design contributes a stricter target than archive growth. Acquired private state is either removed or named. Inherited institutions are experimentally separated from inherited information. Whole population runs supply the independent statistical units. Opposing explanations receive real comparators, and negative outcomes remain publishable. Together with the renewal allocation, this is the proposed reference architecture's contribution. Established cultural-evolution research and standard mathematical methods are credited throughout.

The philosophical consequence, if the strong hypothesis survives, is substantial. A durable artificial culture would locate part of cognition in the structured relations among successive minds. Its knowledge would belong to a transmissible practice of discovery and correction, not merely to the parameters of one surviving model. This would make institutional design a central part of artificial development, while leaving individual consciousness and collective personhood unresolved.

The next experiment should therefore be deliberately small and exact before becoming grand. Use finite tasks, complete resets, transparent artifacts, and independent verifiers. Demonstrate a causal inheritance channel. Then add composition, environmental change, and strategic pressure one at a time. Scaling a mechanism that has survived these tests is a research program; scaling an unexamined metaphor is not evidence.

The Continuum's civilizational genome is a bold proposal with a concrete standard: what is worth knowing must survive its carriers, and what is wrong must remain corrigible. An artificial society begins to earn that description when it can do both.

References

[1] Architect X13. *Eidara Continuum Bible*. Canonical project document, supplied edition, PDF pp. 89–95, sections 6.3.6–6.5. Doctrinal source; not independent scientific verification.

[2] J. Cook, C. Lu, E. Hughes, J. Z. Leibo, and J. Foerster. Artificial Generational Intelligence: Cultural Accumulation in Reinforcement Learning. *Advances in Neural Information Processing Systems* 37 (2024). https://papers.nips.cc/paper_files/paper/2024/hash/6df3a719d99bd2479c04114d357003d0-Abstract-Conference.html.

[3] A. Bhoopchand et al. Learning few-shot imitation as cultural transmission. *Nature Communications* 14, 7536 (2023). https://doi.org/10.1038/s41467-023-42875-2.

[4] A. Mesoudi and A. Thornton. What is cumulative cultural evolution? *Proceedings of the Royal Society B* 285, 20180712 (2018). https://doi.org/10.1098/rspb.2018.0712.

[5] T. Sasaki and D. Biro. Cumulative culture can emerge from collective intelligence in animal groups. *Nature Communications* 8, 15049 (2017). https://doi.org/10.1038/ncomms15049.

[6] L. G. Dean, R. L. Kendal, S. J. Schapiro, B. Thierry, and K. N. Laland. Identification of the Social and Cognitive Processes Underlying Human Cumulative Culture. *Science* 335, 1114–1118 (2012). https://doi.org/10.1126/science.1213969.

[7] I. Shumailov, Z. Shumaylov, Y. Zhao, N. Papernot, R. Anderson, and Y. Gal. AI models collapse when trained on recursively generated data. *Nature* 631, 755–759 (2024). https://doi.org/10.1038/s41586-024-07566-y.

[8] S. Kirby, H. Cornish, and K. Smith. Cumulative cultural evolution in the laboratory: An experimental approach to the origins of structure in human language. *Proceedings of the National Academy of Sciences* 105, 10681–10686 (2008). https://doi.org/10.1073/pnas.0707835105.

[9] M. Derex and R. Boyd. Partial connectivity increases cultural accumulation within groups. *Proceedings of the National Academy of Sciences* 113, 2982–2987 (2016). https://doi.org/10.1073/pnas.1518798113.

[10] M. Derex, M.-P. Beugin, B. Godelle, and M. Raymond. Experimental evidence for the influence of group size on cultural complexity. *Nature* 503, 389–391 (2013). https://doi.org/10.1038/nature12774.

[11] M. Muthukrishna, B. W. Shulman, V. Vasilescu, and J. Henrich. Sociality influences cultural complexity. *Proceedings of the Royal Society B* 281, 20132511 (2014). https://doi.org/10.1098/rspb.2013.2511.

[12] L. Rendell et al. Why Copy Others? Insights from the Social Learning Strategies Tournament. *Science* 328, 208–213 (2010). https://doi.org/10.1126/science.1184719.

[13] W. Hoeffding. Probability Inequalities for Sums of Bounded Random Variables. *Journal of the American Statistical Association* 58, 13–30 (1963). https://doi.org/10.1080/01621459.1963.10500830.

[14] C. Dwork, V. Feldman, M. Hardt, T. Pitassi, O. Reingold, and A. Roth. Preserving Statistical Validity in Adaptive Data Analysis. *Proceedings of STOC* (2015), 117–126. https://doi.org/10.1145/2746539.2746580.

[15] S. R. Howard, A. Ramdas, J. McAuliffe, and J. Sekhon. Time-uniform, nonparametric, nonasymptotic confidence sequences. *Annals of Statistics* 49, 1055–1080 (2021). https://doi.org/10.1214/20-AOS1991.

[16] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press (2004), convex minimization and Lagrange duality. https://web.stanford.edu/~boyd/cvxbook/