FLOOD AND HARVEST: THE PROVABLE NECESSITY OF TRIVIA FOR GENERATING VALUABLE MATHEMATICS

Anonymous authors
Paper under double-blind review

ABSTRACT

A proof checker certifies that a statement is valid, not that it is worth keeping. We ask what a generator can learn about value from an incomplete record of valuable mathematics, starting from one elementary fact: a generator acts only on its information, so everything it can guarantee is decided by the worlds that its information cannot separate. First, a membership verifier *localizes*: a task judged in the limit is solvable with a verifier exactly when it is solvable inside each formal world, where the verifier is silent. Second, inside a formal world, a record that might be complete leaves value ambiguous. If the record lists K and a larger candidate K' is equally consistent with it, every output in $K' \setminus K$ is trivia under one reading and a discovery under the other, and the two trade at an exchange rate of exactly one. If m of the first N points of K' are recorded, the best coverage of them that can be guaranteed by round $T \geq 2N+2$ is $m/2+\min\{g(T), N-m\}$ up to $O(\sqrt{N})$, for every trivia budget $g(T) \leq \sqrt{T}$ that grows by at most one per round. Guaranteed coverage therefore takes only two values across all trivia regimes: $\alpha/2$ when trivia must stay bounded, and $1-\alpha/2$ when it may grow without bound, however slowly, where α is the relative density of K in K'. Recovering the unrecorded region takes the inverse budget at its mass, $(1-\alpha-\varepsilon)^2N^2$ rounds under a square-root budget. A description-length model of value realizes $\alpha = 0$: bounded trivia guarantees nothing, and every unbounded budget guarantees everything.

> *"Discovery consists precisely in not constructing useless combinations, but in constructing those that are useful, which are an infinitely small minority. Discovery is discernment, selection."*
> — Henri Poincaré, *Science and Method* (1914)

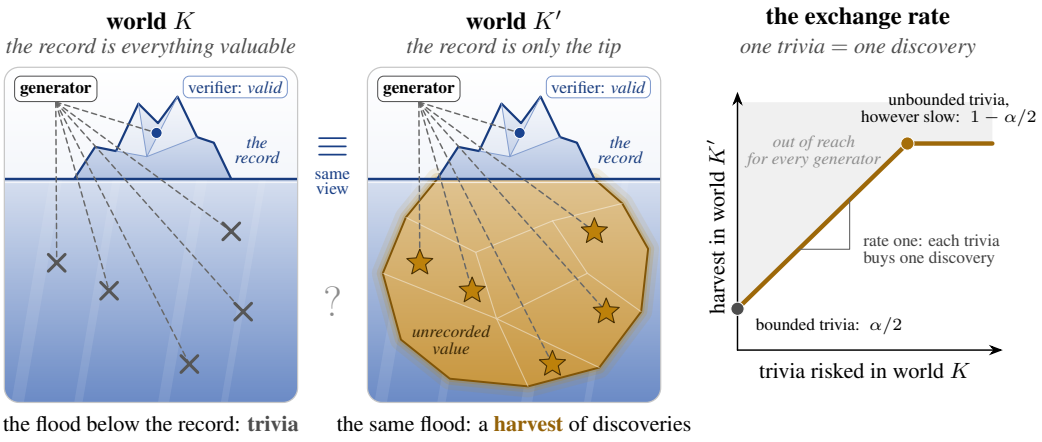


Figure 1: **The paper in one picture.** Above the waterline the two worlds are identical: the same record and the same verifier answers. Every output below the record, the *flood*, is therefore trivia in K, where the record is everything valuable, and a discovery in K', where it is only the tip. The two trade one for one (Theorem 4.1), and guaranteed coverage of K' is $\alpha/2$ under bounded trivia and $1-\alpha/2$ under unbounded trivia, however slow (Corollary 4.2), where α is the recorded share.

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

1 INTRODUCTION

In August 2026, OpenAI reported ten advances on longstanding problems in mathematics and theoretical computer science, together with Lean certificates (OpenAI, 2026b). In September, it released a proposed solution to the Navier–Stokes Millennium Prize problem and a Lean formalization (OpenAI, 2026a). Anthropic, meanwhile, reported a formalization of Fermat's Last Theorem comprising roughly thirteen million lines of Lean code (Anthropic, 2026). Finding a new argument and formalizing an existing one are different achievements. Both, however, make it worth asking how mathematics should develop when machines can produce and check substantial bodies of work.

A proof checker answers one question about a statement: is it valid? Mathematicians ask a second: is it worth keeping? A named conjecture answers the second question in advance, because the community has already decided that the statement is worth proving; a formalization project likewise starts from mathematics already judged worth preserving. A generator sent to explore beyond such assignments has to answer the second question for itself. It can be shown the results that mathematicians chose to keep. It cannot be shown the results nobody has found yet, and finding some of those is the reason to run it. Terence Tao and other mathematicians have warned that producing solutions does not by itself develop mathematical understanding (Avila et al., 2026). We isolate one narrower difficulty that can be made precise: what can an incomplete record of valuable mathematics tell a generator about statements that have never appeared in it?

We answer starting from a single fact. A generator acts only on its information, meaning the record it has observed and the answers of its verifier. Whatever it can guarantee is therefore decided by the *fibers* of the map that sends a possible world to the information it presents, that is, by the sets of worlds that the generator cannot tell apart (Figure 2). Every output is judged in all worlds of its fiber at once, and nothing about value can be learned beyond what separates them. The paper computes the consequences for the two fibers that matter in automated mathematics: the worlds that a verifier cannot separate, and the worlds that a record cannot separate.

Verification localizes. A membership verifier observes the formal world F, the set of valid statements, and nothing else. In the limit it identifies F, and inside a fixed F it gives the same answers for every candidate notion of value. Theorem 3.1 shows that nothing more is true. A generation task judged in the limit is solvable with a verifier exactly when it is solvable inside each formal world separately, and there a verifier is useless. Characterizations of generation with breadth transfer fiber by fiber, and every countable model admits eventual sound coverage. Verification thus reduces the question of value to a single formal world, but it does not settle the question there.

A record cannot certify its own completeness. Inside a formal world the generator learns about value only from the record, and a record never announces that it is complete. Let K' be the even numbers and $K = \{2, 6, 10, \ldots\}$, and suppose the record lists K. If K is the valuable set, the record is complete. If K' is, the record omits every multiple of four and always will. No observation and no verifier query distinguishes the two worlds, so a deterministic generator produces the same outputs in both. An output such as 4 is a discovery in one world and trivia in the other: every output in the unrecorded region is entered twice, in two ledgers that the generator can never compare (Figure 1).

Trivia and discovery trade one for one. The double entry yields an inequality, and our main theorem shows that it is exact. Suppose that m_N of the first N points of K' are recorded and $n_N = N - m_N$ are not. Consider a generator allowed $g(T)$ trivia by round T under the smaller reading. On its worst record it can have covered at most $m_N/2 + \min\{g(T), n_N\}$ of these points by round T under the larger reading, up to a logarithmic term. A simple generator attains this on every record, for every budget that grows by at most one per round: it races for recorded points and spends each unit of budget on the next unrecorded one (Theorem 4.1). Its error is $o(N)$ whenever the budget is sublinear, and at most $O(\sqrt{N})$ from round $2N + 2$ on when $g(T) \leq \sqrt{T}$. The first term is the known price of racing an adversarial record (Kleinberg & Wei, 2026), and the second is new. Each unit of trivia risked under one reading buys exactly one discovery under the other (Figure 1, right).

The exchange rate settles two questions at once. The first is how much trivia is necessary. Let α be the relative density of K in K'. Over every notion of "few trivia" that allows fewer whenever it allows more, the guaranteed coverage of K' takes one of two values: $\alpha/2$ if trivia must stay bounded, and $1 - \alpha/2$ if it may grow without bound, however slowly (Corollary 4.2). No bounded amount of trivia buys any of the unrecorded share $1 - \alpha$, and any unbounded amount buys all of

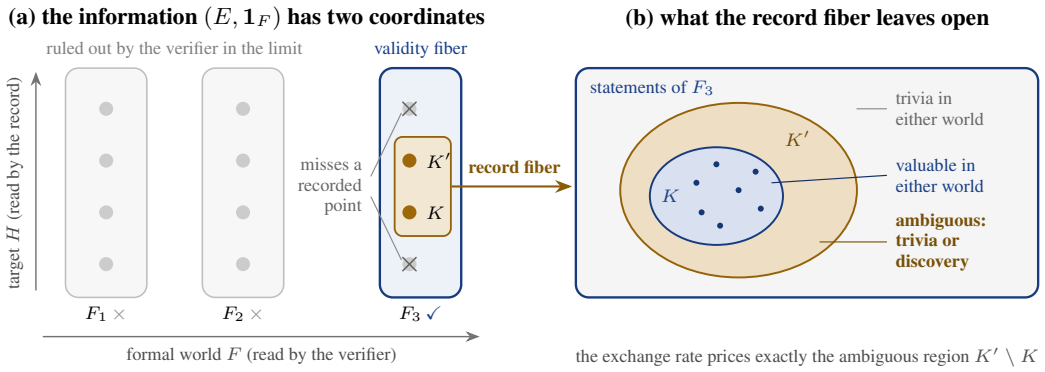


Figure 2: **The two fibers of the information map.** (a) The verifier separates formal worlds in the limit and is silent inside one, the validity fiber (Theorem 3.1). There the record rules out every target that misses a recorded point, leaving the record fiber $\{K, K'\}$. (b) Only the ambiguous region $K' \setminus K$, between the fiber's intersection and its union, can turn trivia into discovery (Theorem 4.1).

it; this is the precise sense in which trivia is provably necessary. The second question is how long discovery takes. Eventual coverage ignores time, but the frontier does not. For every sublinear budget, coming within εN of the larger frontier takes the inverse budget evaluated at the unrecorded mass, $g^{-1}\big((1-\alpha-\varepsilon)N \pm o(N)\big)$ rounds. Under a square-root budget this is $(1-\alpha-\varepsilon)^2 N^2$ rounds up to a factor $1 + O(N^{-1/2})$, which pins the leading constant (Corollary 4.3). Rare exploration preserves the harvest and postpones it.

The ambiguity is not an artifact of integer examples. In a description-length model of value in the spirit of Aksenov et al. (2026), a word is valuable when a dictionary describes it by logarithmically many blocks. Enlarging the dictionary with new tags yields nested targets in which the record of the smaller target has relative density zero in the larger (Proposition 5.1). There, bounded trivia guarantees no coverage at all and every unbounded budget guarantees full coverage. Almost everything that a richer vocabulary can compress is invisible from the record of a poorer one.

What is new. The protocol is the partial-enumeration model of Kleinberg & Wei (2026); their race supplies the term $m_N/2$ and the endpoint $\alpha/2$. The exact frontier with its converse and attainment, its two corollaries, and the density-zero compression construction are new. The localization theorem assembles classical ingredients, namely identification by enumeration (Gold, 1967) and the breadth characterizations of Kalavasis et al. (2026). What it adds is the statement that these ingredients exhaust what a verifier can do, together with separations that show where they stop. Sparse errors are known to help in timely generation (Ganju et al., 2026), and error rates have been organized into hierarchies (Panigrahi et al., 2026). Here the error is trivia inside a verified world, the constraint binds under an indistinguishable reading, and its price is exact. Appendix G compares the results with prior work theorem by theorem.

2 Worlds, Records, and Information

Worlds. Statements are elements of a countable universe $U = \{x_1 \prec x_2 \prec \cdots\}$, listed in a fixed order of type ω, and $L[1 \dots N]$ denotes the first N elements of an infinite set $L \subseteq U$. A *world* is a pair (F, H) with $H \subseteq F \subseteq U$. The formal world F contains the statements that pass the verifier; mathematically, its elements encode statements with certificates. The target H contains the valuable statements. An output is valuable if it lies in H, *trivia* if it lies in $F \setminus H$, and invalid otherwise. The generator knows a countable *pair model* $\mathcal{P}$ of candidate worlds but not which candidate is true.

Records. The adversary fixes a true world (F, H) and an infinite core $C \subseteq H$, and enumerates C without repetition so that every point of C eventually appears; this is the partial-enumeration protocol of Kleinberg & Wei (2026). The generator is not told C. In round t it sees an observation w_t, may query the membership oracle $\mathbf{1}_F$ finitely many times, and outputs some $a_t \notin S_t = \{w_1, \dots, w_t\}$.

We write $O_T = \{a_1, \ldots, a_T\}$ and $O = \bigcup_T O_T$. Generators are deterministic and computationally unrestricted, since our statements concern information rather than computation.

Coverage and trivia. An infinite output set can still occupy a vanishing fraction of the target, so we measure breadth along prefixes. Coverage counts distinct valuable outputs in a prefix of the target, $|O_T \cap H[1 \ldots N]|$, and its asymptotic form is the lower density

$$\underline{d}(O, H) = \liminf_{N \to \infty} \frac{|O \cap H[1 \ldots N]|}{N}. \tag{1}$$

Trivia are counted the same way: $b_H(T) = |O_T \cap (F \setminus H)|$ is the number of distinct trivial outputs by round T. Our generators never repeat an output, so counting trivial rounds would change nothing.

Information. An instance $\iota = ((F, H), C, E)$ presents to the generator exactly the information

$$\pi(\iota) = (E, \mathbf{1}_F),$$

its record and its oracle, and nothing else. A deterministic generator's run is a function of $\pi(\iota)$, so instances in one fiber of π receive the same outputs.

Lemma 2.1 (One transcript)**.** If two instances share the formal world F and the observation sequence, then a deterministic generator makes the same queries, receives the same answers and produces the same outputs in every round of both.

Proof. Induct over rounds and, within a round, over queries. The next observation is shared, and every query q receives the answer $\mathbf{1}_F(q)$ in both instances. Lemma A.1 records the details. □

The lemma is elementary, but its consequences are not: every guarantee becomes a statement about the fibers of π. Its two coordinates organize the paper. Fixing the oracle gives the *validity fiber* $\mathcal{P}_F = \{(F', H') \in \mathcal{P} : F' = F\}$, the candidates on which the verifier answers identically. Fixing a complete record of a set C gives the *record fiber*, the candidate targets that contain C; once the whole record is known, these are the targets that remain possible. Points in the intersection of the record fiber are valuable in every possible world; points outside its union, in none. Between them lies the *ambiguous region*: there the worlds disagree, and value is not a function of information.

The ledger. Let $H \subsetneq H'$ be targets in one validity fiber, and let the record enumerate H completely. Then (F, H) and (F, H') with this record are both legal instances, and by Lemma 2.1 a generator's outputs form one sequence read in two ways. The two readings differ exactly on the outputs in $H' \setminus H$: for every T,

$$b_H(T) - b_{H'}(T) \;=\; |O_T \cap (H' \setminus H)| \;=\; |O_T \cap H'| - |O_T \cap H|. \tag{2}$$

Every output in the unrecorded region $H' \setminus H$ is a discovery under H' and trivia under H, and nothing the generator observes can tell which entry is correct. Section 3 asks what a verifier adds to this picture, and Section 4 computes what the identity (2) costs.

3 Verification localizes

A verifier answers questions about F. We first ask whether it can answer anything else. Call a success criterion *judged in the limit* if changing finitely many outputs of a run never changes whether that run succeeds. Generation in the limit, generation with exact or approximate breadth (Kleinberg & Mullainathan, 2024; Kalavasis et al., 2026), lower-density coverage and eventual soundness are all judged in the limit. A trivia budget that must hold at every round is not.

Theorem 3.1 (Verification localizes)**.** Let $\mathcal{P}$ be a countable pair model, and fix a protocol and a success criterion.

(i) (Stripping.) For each formal world F, every verifier-assisted generator, run on instances with formal world F, coincides round by round with an ordinary generator. Hence, on a validity fiber, a verifier improves no guarantee of any kind.

(ii) (Gluing.) Suppose the criterion is judged in the limit. Then some verifier-assisted generator succeeds on every instance of $\mathcal{P}$ if and only if, for each formal world F occurring in $\mathcal{P}$, some ordinary generator succeeds on every instance of $\mathcal{P}_F$. The glued generator queries one new point per round.

Proof. (i) Answer every query with $\mathbf{1}_F$; by determinism, the simulation reproduces every round. (ii) Necessity follows from (i). For sufficiency, list the candidate formal worlds $F_1, F_2, \ldots$ and query x_t in round t. Let $\hat{\imath}_t$ be the least index whose world agrees with every answer received so far; this is identification by enumeration from an informant (Gold, 1967). The least index i^* naming the true F is never excluded. Each of the finitely many smaller indices names a different set and is excluded after finitely many rounds, so $\hat{\imath}_t = i^*$ from some round on, although the generator cannot tell when. Take ordinary generators Γ_F for the fibers; each may be taken to be a function of the observations alone, because a deterministic generator can recompute its own earlier outputs. In round t, output $\Gamma_{F_{\hat{\imath}_t}}(w_1, \ldots, w_t)$, replaced by a fixed legal output whenever it is illegal. From the stabilization round on, these are exactly the outputs of Γ_F on the same record. A criterion judged in the limit cannot tell the two runs apart. □

Both parts are simple, and their ingredients are classical. Together they say that a verifier contributes one coordinate of the world, eventually, and nothing else. The following corollary makes this concrete; Appendix B gives the proofs.

Corollary 3.2 (What verification resolves)**.** Let $\mathcal{P}$ be a countable pair model and consider complete records with set-valued outputs $G_t \subseteq U \setminus S_t$.

(a) Exact [approximate] breadth is achievable with a verifier if and only if every validity fiber satisfies Angluin's tell-tale condition [its weak form].

(b) With a verifier, every $\mathcal{P}$ admits eventual sound coverage, meaning that $H \setminus S_t \subseteq G_t \subseteq F$ from some round on.

(c) Both parts of Theorem 3.1 are sharp. Without a verifier, some $\mathcal{P}$ admits no eventual sound coverage. On some $\mathcal{P}$, no verifier-assisted generator is sound at every round while covering soundly in the limit.

Part (a) is Theorem 3.1 applied to the characterizations of Kalavasis et al. (2026), and part (b) is gluing applied to the fiber generator $G_t = F \setminus S_t$. The separations in (c) need genuine constructions. One uses a chain of formal worlds whose sound-coverage requirements conflict after any finite record. The other uses worlds that differ at a single point the generator has not queried.

The corollary splits the problem of value in two. Across formal worlds, verification does all the work: once F is known, outputting everything valid covers every valuable statement soundly. Within a formal world it does none. The exhaustive output of (b) contains every unrecorded valuable statement, but it also contains all of $F \setminus H$. Whether that trivia is necessary is a question about a single validity fiber, where the record is the only information about value. A trivia budget imposed at every round is not judged in the limit, and it does not glue across formal worlds (Appendix B.4): a candidate world may exclude exactly the points the generator has already output, and no finite number of queries rules that world out. We therefore analyze budgets inside one formal world. From now on we fix the formal world and set $F = U$, so that every output is valid and the verifier falls silent.

4 The Exchange Rate

Fix the formal world $F = U$ and two nested candidate targets $K \subsetneq K'$ whose difference $D = K' \setminus K$ is infinite. Among the first N points of K', let

$$m_N = |K \cap K'[1 \ldots N]|, \qquad n_N = N - m_N = |D \cap K'[1 \ldots N]|$$

count the recorded and the unrecorded points. In the running family K' is the even numbers and K keeps the points of rank $\lceil k/\alpha \rceil$ in K', so that $m_N = \lfloor \alpha N \rfloor$ for any real $\alpha \in (0, 1)$; at $\alpha = 1/2$ this gives $K = \{2, 6, 10, \ldots\}$. Let $\mathcal{E}(K)$ be the set of complete repetition-free enumerations of K. Each $E \in \mathcal{E}(K)$ is a legal record under both readings, and we measure coverage of K' on such records in the worst case:

$$Q_{G,E}(T, N) = |O_T(G, E) \cap K'[1 \ldots N]|, \qquad E \in \mathcal{E}(K).$$

Trivia are constrained under *both* readings; were they constrained only under K', exploration would be free. A *budget* is a nondecreasing $g : \mathbb{N} \to \mathbb{N}$ with $g(0) = 0$ and $g(t + 1) \leq g(t) + 1$. The last condition loses nothing: a trivia count grows by at most one per round, so it stays below a nondecreasing function exactly when it stays below the largest budget beneath that function (Appendix A). Let $\mathcal{G}_g$ be the class of generators whose output is trivial in at most $g(T)$ of the first T rounds, for every T and on every legal instance of either target. We use the convention $\min \varnothing = \infty$.

4.1 CONTESTED AND UNCONTESTED POINTS

Points of K and points of D behave differently. A point of K may be observed before the generator reaches it, and an observed point can never become a new output; the record contests it. A point of D is never observed, so it waits. Contested points reward speed, uncontested points reward persistence, and the two mechanisms do not interact.

For contested points the answer is known (Kleinberg & Wei, 2026). Greedy racing outputs the least point of K that has been neither observed nor output. Until a prefix of K is resolved, each race round claims a new point of it and each point left unclaimed costs an observation, so racing secures about half of every prefix. Conversely, a racing record observes the least point that neither party has used and occasionally recycles points that the generator has claimed, which holds every generator to about half. We use both directions in the following finite form (Lemma C.1). If racing is interrupted in $w(t)$ of the first t rounds, it has claimed at least $(M - 1 - w(t_M))/2$ points of $K[1 \dots M]$ by the round t_M at which that prefix is resolved. And for every generator G there is a record $E_G \in \mathcal{E}(K)$ with

$$|O(G, E_G) \cap K[1 \dots M]| \le M/2 + \log_2 M + 3 \qquad \text{for all } M \ge 2. \tag{3}$$

4.2 THE FRONTIER

Theorem 4.1 (Exchange rate one)**.** Let $\Phi_g(T, N) = m_N/2 + \min\{g(T), n_N\}$.

(a) (Converse.) For every deterministic generator G there is a record $E_G \in \mathcal{E}(K)$ such that, for all $T \ge 1$ and all N with $m_N \ge 2$,

$$Q_{G,E_G}(T, N) \;\le\; m_N/2 + \min\{b_K(T), n_N\} + \log_2 m_N + 3,$$

where $b_K(T)$ counts the distinct trivia that G has produced by round T when the same record is read under K. In particular, $Q_{G,E_G}(T, N) \le \Phi_g(T, N) + \log_2 m_N + 3$ for every $G \in \mathcal{G}_g$, whatever other candidates the model contains.

(b) (Attainment.) Let $\tau_g(M) = \min\{t : t - g(t) \ge M\}$. Some $G_g \in \mathcal{G}_g$ produces no trivia under K' and satisfies, for every $E \in \mathcal{E}(K)$, every N with $m_N \ge 1$ and $\tau_g(m_N) < \infty$, and every $T \ge \tau_g(m_N)$,

$$Q_{G_g,E}(T, N) \;\ge\; \Phi_g(T, N) - \big(1 + g(\tau_g(m_N))\big)/2.$$

If $g(T) = o(T)$, the error is $o(m_N)$. If $g(T) \le \sqrt{T}$, then $\tau_g(M) \le 2M + 2$ and the error is at most $\big(1 + \sqrt{2m_N + 2}\big)/2$.

Proof. (a) Let E_G be the racing record against G, and split $K'[1 \dots N]$ into $K[1 \dots m_N]$ and $D \cap K'[1 \dots N]$. By (3), the first part contributes at most $m_N/2 + \log_2 m_N + 3$, even to the limit set. Every output in D is trivia under K, so by the ledger (2) the second part contributes at most $b_K(T)$, and at most n_N.

(b) The generator G_g counts its sweeps. In round t it *sweeps* if it has made fewer than $g(t)$ sweeps so far, outputting the least point of D it has not yet output. Otherwise it *races*, outputting the least point of K that has been neither observed nor output. If an observation ever falls outside K, the smaller reading is refuted and G_g races on K' from then on. Under K its trivia are exactly its sweeps, at most $g(T)$ by round T, and under K' it produces none. On a record in $\mathcal{E}(K)$, the unit steps of g make the sweep count equal $g(t)$ exactly. Because points of D are never observed and only sweeps output them, the jth sweep outputs the jth point of D, so the second part of the prefix contributes exactly $\min\{g(T), n_N\}$. The sweeps interrupt the race in $w(t) = g(t)$ of the first t rounds, so the Mth race round is $\tau_g(M)$, and the finite race bound yields at least $(m_N - 1 - g(\tau_g(m_N)))/2$ points of $K[1 \dots m_N]$ by that round. Finally, $\tau_g(M) = M + g(\tau_g(M))$: a sublinear budget makes the error $o(M)$, and $g(t) \le \sqrt{t}$ gives $\tau_g(M) \le 2M + 2$. □

For a sense of scale, take $\alpha = 1/2$ and $g(T) = \lfloor \sqrt{T} \rfloor$. The first thousand points of K' contain 500 recorded and 500 unrecorded statements. After 10^4 rounds the best guarantee on them is about $250 + 100$, and it approaches 750 only after about 2.5×10^5 rounds. Each unit of budget buys exactly one unrecorded point, and part (a) shows that nothing can buy more: the exchange rate between trivia risked under K and discoveries made under K' is one. The only other loss is the half of the record

lost to the race, the price of contested points rather than of trivia. Computed runs meet both sides of the frontier (Figure 7), and both parts of the theorem are checked in Lean (Appendix F).

The theorem also says where to explore. Only outputs in the ambiguous region D are exchanged: an output outside K' is trivia under both readings and buys nothing. A generator that sweeps the whole formal world, in the spirit of the exhaustive generator of Corollary 3.2(b), trades at a lower rate. In the running family at $\alpha = 1/2$, with $F = \mathbb{N}$, two of every three trivial outputs of such a sweep are odd numbers, so it buys one discovery per three trivia. Exploration pays at the full rate only where the plausible accounts of value disagree.

4.3 Two values

Asymptotically the frontier collapses to two numbers. A *trivia regime* is a nonempty set $\mathcal{R}$ of nondecreasing functions $\mathbb{N} \to \mathbb{N}$ that contains, with each member, every smaller nondecreasing function. A generator obeys $\mathcal{R}$ if, on every legal instance of either target, its count of distinct trivia as a function of T lies in $\mathcal{R}$. Zero trivia, finitely many trivia, a prescribed budget and a vanishing trivia rate are all regimes.

Corollary 4.2 (The necessity of trivia)**.** Suppose that $m_N/N \to \alpha \in [0,1)$, and let $V(\mathcal{R}) = \sup_{G \text{ obeys } \mathcal{R}} \inf_{E \in \mathcal{E}(K)} \underline{d}\big(O(G,E), K'\big)$. Then

$$V(\mathcal{R}) = \alpha/2 \text{ if every member of } \mathcal{R} \text{ is bounded}, \qquad V(\mathcal{R}) = 1 - \alpha/2 \text{ otherwise},$$

and the second value is attained with no trivia under K'.

Proof. If every member is bounded, apply Theorem 4.1(a) with b_K bounded and let T and then N grow; pure racing ($g = 0$) attains $\alpha/2$ with no trivia at all. If $f \in \mathcal{R}$ is unbounded, then $g(t) = \min\{f(t), \lfloor\sqrt{t}\rfloor, g(t-1)+1\}$ is an unbounded budget lying below f, so G_g obeys $\mathcal{R}$. Letting $T \to \infty$ in Theorem 4.1(b) gives $|O \cap K'[1 \dots N]| \geq N - m_N/2 - O(\sqrt{N})$. The matching upper bound $1 - \alpha/2$ is Theorem 4.1(a) with $\min\{b_K, n_N\} \leq n_N$. □

The rate of trivia is irrelevant to the eventual harvest. What matters is whether exploration may continue forever: uncontested points wait, so infinitely many excursions reach all of them, however far apart the excursions are. The converse holds in every pair model that contains K and K', so whenever a record might be complete, breadth beyond it is taxed at this rate.

4.4 The clock

The eventual harvest ignores time, and the frontier restores it. For $0 < \varepsilon < 1 - \alpha$, let $\tau_N(\varepsilon)$ be the first round with $Q(T,N) \geq N - m_N/2 - \varepsilon N$, that is, within εN of the eventual optimum on the prefix. Write $g^{-1}(y) = \min\{T : g(T) \geq y\}$.

Corollary 4.3 (The clock is the inverse budget)**.** Let $\alpha \in (0,1)$, let g be an unbounded budget with $g(T) = o(T)$, and let $m_N = \lfloor \alpha N \rfloor$. Every $G \in \mathcal{G}_g$ has $\tau_N(\varepsilon) \geq g^{-1}\big((1-\alpha-\varepsilon)N - \log_2 N - 3\big)$ on E_G, for all N with $m_N \geq 2$. For all sufficiently large N, the generator G_g has $\tau_N(\varepsilon) \leq g^{-1}\big((1-\alpha-\varepsilon)N + o(N)\big)$ on every $E \in \mathcal{E}(K)$, and if $g(T) \leq \sqrt{T}$, more precisely $\tau_N(\varepsilon) \leq g^{-1}\big((1-\alpha-\varepsilon)N + \sqrt{2N+2}\big) + 2N + 2$. Under $g(T) = \lfloor\sqrt{T}\rfloor$, the minimax recovery time $\inf_{G \in \mathcal{G}_g} \sup_{E \in \mathcal{E}(K)} \tau_N(\varepsilon)$ is therefore $(1-\alpha-\varepsilon)^2 N^2 \big(1 + O(N^{-1/2})\big)$ for fixed α and ε.

Both bounds follow by solving Theorem 4.1 for $g(T)$. Under $g(T) = \lfloor \log_2 T \rfloor$, for example, the clock is exponential, $2^{(1-\alpha-\varepsilon)N \pm O(\log N)}$. Slowing the budget never lowers the eventual harvest, but it postpones the harvest without limit. On computed runs at $\alpha = 1/3$, the square-root recovery time at $N = 3200$ is within 1% of $(1-\alpha-\varepsilon)^2 N^2$ for $\varepsilon \in \{0.05, 0.1\}$ (Figure 3). The faster budget $g(T) = \lfloor T^{3/4} \rfloor$ lies outside the square-root regime, and there the general attainment bound applies: at $N = 1600$ and $\varepsilon = 0.05$ its computed recovery time is within 5% of $((1-\alpha-\varepsilon)N)^{4/3}$ (Appendix E).

Scope. The half is the price of one observation per output. Against r observations per round it becomes $1/(1+r)$, while the term $\min\{g, n_N\}$ is unchanged (Remark C.2). The shared record is

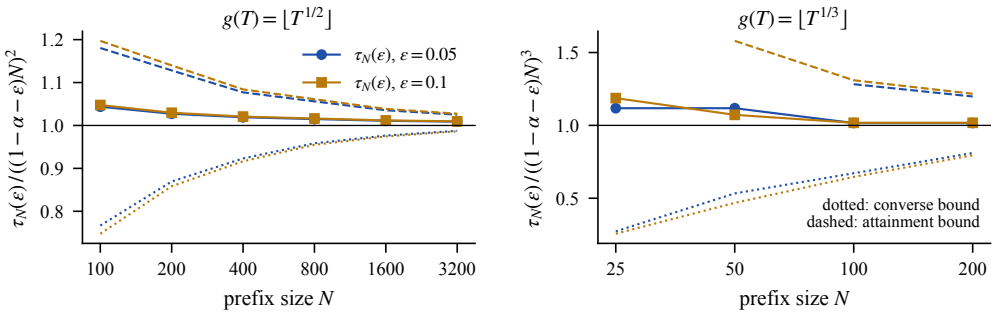


Figure 3: The clock of Corollary 4.3 on computed runs at $\alpha = 1/3$. Recovery times are divided by the leading term $((1-\alpha-\varepsilon)N)^2$ under $g(T) = \lfloor T^{1/2} \rfloor$ (left) and $((1-\alpha-\varepsilon)N)^3$ under $g(T) = \lfloor T^{1/3} \rfloor$ (right). Solid lines show the race-and-sweep generator on its racing record, dotted lines the converse bound that every budgeted generator obeys on its own racing record, and dashed lines the attainment bound. Both bounds are stated inequalities, not fits.

essential: once an observation leaves K, the smaller reading is refuted, the ledger closes, and G_g covers K' like any racer. In a general countable family, sweeping the union of consistent candidates at a vanishing rate still outputs every unrecorded valuable point (Proposition C.4); only the exact frontier needs a shared record.

5 A VALUE RULE THAT HIDES ALMOST EVERYTHING

The integer family fixes value by fiat. We now let value come from a description rule and ask whether the unrecorded share can be all of it. Following the macro framework of Aksenov et al. (2026), call a word valuable when a dictionary describes it with few factors. Words are finite strings over $\{a, x\}$, ordered by length and then lexicographically, and $|w|$ is the length. Fix integers $1 \le p < q$ and a real $\kappa > 0$. A *block* with scale $m \ge 1$ and tag $i \in \{1, \dots, q\}$ is

$$\widehat{\mu}_{m,i} = aa\, x^i a\, x^{q+1-i} a\, x^{2^m}, \tag{4}$$

and the dictionary $\widehat{D}_s$ consists of the blocks with tags $i \le s$. A word is valuable under $\widehat{D}_s$ if it is a concatenation of at most $\kappa \log_2(1+|w|)$ blocks of $\widehat{D}_s$:

$$\widehat{K}_s = \big\{ w = d_1 \cdots d_k : \; d_j \in \widehat{D}_s, \; k \le \kappa \log_2(1+|w|) \big\}, \qquad \widehat{K} = \widehat{K}_p \subsetneq \widehat{K}' = \widehat{K}_q.$$

The larger dictionary has $q-p$ additional tags at every scale (in the metaphor, new concepts), and nothing else changes.

Three features make the tags count. The marker aa occurs exactly at block starts, so every word has at most one block decomposition and its tags are well defined. The runs x^i and x^{q+1-i} have a fixed total length, so the length of a block depends on its scale but not on its tag, and whether a word passes the length test depends only on its scales. Hence among the k-block words of $\widehat{K}'$ with given scales, exactly a fraction $(p/q)^k$ lies in $\widehat{K}$. With $p = 1$ and $q = 2$, for instance, a word of ten blocks is recorded under only one of its 2^{10} tag patterns.

Proposition 5.1 (The record can have density zero)**.** The sets $\widehat{K}$ and $\widehat{K}' \setminus \widehat{K}$ are infinite, and the fraction of $\widehat{K}$ among the first N words of $\widehat{K}'$ is $O(1/\log L)$, where L is the length of the Nth word. Hence $m_N/N \to 0$.

The tag fraction alone does not prove this. One-block words keep the fixed fraction p/q at every scale, and a prefix may end inside a length layer that happens to be rich in $\widehat{K}$. The proof (Appendix D.3) controls both effects. Counting scale choices shows that the number $C(L)$ of words of $\widehat{K}'$ with length at most L is $\exp\big(\Theta(\log L \, \log\log L)\big)$. Consequently, words with $O(\log\log L)$ blocks are negligible, and every other word of $\widehat{K}'$ lies in $\widehat{K}$ for at most a $(\log L)^{-3}$ fraction of tag choices. Deleting a shortest block maps a word of length L into the ball of radius $L-1$ and keeps it compressible, which gives $C(L)/C(L-1) = O(\log^2 L)$ and controls prefixes that stop inside a layer.

By Theorem 4.1 and Corollary 4.2 with $\alpha = 0$, a generator whose trivia under $\widehat{K}$ must stay bounded guarantees coverage zero of $\widehat{K}'$. A generator allowed any unbounded budget guarantees coverage one. Under a budget $g(T) \leq \sqrt{T}$, the worst-case time to cover all but εN of an N-word prefix is $g^{-1}\big((1 - \varepsilon \pm o(1))N\big)$. If the richer dictionary is the right account of value, almost everything valuable is unrecorded, and the record of the poorer dictionary cannot tell any of it apart from trivia. Compression is a model of value here, not a thesis about mathematical importance: it shows that a natural value rule can hide the entire harvest from the record. Appendix D also realizes every rational α, and gives commutative dictionaries with identifiable targets, where breadth needs no trivia.

6 Related work

Generation in the limit. Kleinberg & Mullainathan (2024) separate generation from identification: every countable family admits a generator that eventually emits only unseen members of the target. Breadth asks for more, and Kalavasis et al. (2026) show that exact and approximate breadth revert to the identification conditions of Angluin (1980). Kleinberg & Wei (2025) measure breadth by limit-set density. Kleinberg & Wei (2026) study partial enumeration and prove that half of the core's density is attainable and, against a racing record, optimal. Cai et al. (2026) simplify dense generation and obtain the randomized optimum $1 - 1/e$ against oblivious enumeration. We keep the partial-enumeration protocol and read one record in two ways: the half is the price of their race, and the exchange term is the price of the second reading.

Errors, hallucination and time. Ganju et al. (2026) show that sparse hallucinations overcome coverage barriers in timely generation, and Panigrahi et al. (2026) organize hallucination rates into hierarchies. There, an error is an output outside the target. Here every output is valid: the error is trivia inside a verified world, it is priced under a reading that the generator cannot rule out, and its price is exact at every finite time.

Verifiers and value. Proof search and program search (Polu & Sutskever, 2020; Lample et al., 2022; Yang et al., 2023; Romera-Paredes et al., 2024) motivate the ambient verifier, and Botta et al. (2025) count the verifier queries that generation needs when the verifier certifies prefixes. Interestingness has been studied since automated theory formation (Colton et al., 2000) and is now learned from data (Tsoukalas et al., 2025). Those works propose measures of value. We ask what can be guaranteed before any such measure is available, and Aksenov et al. (2026) supply the description-length view of value that our compression construction instantiates.

Inductive inference. A finite portion of a text never rules out a larger language. This is the source of Gold's theorem that superfinite classes are not identifiable from text (Gold, 1967), and of Angluin's tell-tale characterization (Angluin, 1980). Partial records make the obstruction permanent: even the entire record of K never rules out K'. Identification asks whether such an obstruction can be overcome; generation with breadth asks what it costs. On a shared record, the cost is the unrecorded share of the larger target, paid in trivia at rate one.

7 Conclusion

A generator can guarantee only what the worlds its information cannot separate have in common, and this single fact organizes the paper. A verifier separates formal worlds in the limit and nothing more, so it localizes the question of value to one validity fiber (Theorem 3.1). Inside that fiber a record cannot certify its own completeness, and every output in the ambiguous region is entered twice, as trivia under one reading and as a discovery under the other. The two entries trade at exactly one (Theorem 4.1): guaranteed coverage is $\alpha/2$ when trivia must stay bounded and $1 - \alpha/2$ when it may grow without bound, however slowly (Corollary 4.2); the budget sets only the clock (Corollary 4.3); and a description-length model of value realizes $\alpha = 0$, where bounded trivia guarantees nothing and every unbounded budget guarantees everything (Proposition 5.1). Trivia is thus the price of reaching what the record does not yet contain. The harvest, however, arrives unlabeled, because the generator covers the unrecorded region without learning where it is. The discernment in Poincaré's epigraph is exactly the information that neither the record nor the verifier supplies; Appendix I discusses this gap, what the theorems advise within the model, and open directions.

Reproducibility Statement

The results are mathematical. Section 2 specifies the protocol, and Sections 3 to 5 state the results with their principal proofs. Appendices A to D give the complete definitions, statements and proofs. The supplementary code and result files reproduce the deterministic computations of Appendix E, including Figure 7, and Appendix F identifies the checked Lean version, its entry points and the scope of the formalized statements. No empirical benchmark results are claimed.

AI Use Statement

The research question came from the authors, and the initial draft was written by the authors. Generative AI tools (OpenAI Codex and Anthropic Claude Code) were used for literature formatting, for drafting, editing and restructuring the manuscript, for formulating and checking mathematical statements and proofs, for writing the Lean formalization of main results, and for writing and checking the simulation code. The mathematical uses included the organization around localization and the exchange-rate frontier, and the derivations and bounds. The authors checked all proofs, simulation outputs and formal correspondences, and take full responsibility for the content and correctness of the manuscript.

Ethics Statement

This work develops mathematical guarantees and counterexamples. It involves no human participants, no personal data and no empirical deployment. The word "valuable" denotes membership in a specified latent target, not a universal ranking of mathematical subjects, researchers or communities. The results concern an abstract exploration stream and do not justify flooding public mathematical repositories; the numerical illustrations use synthetic integer languages.

References

Vitaly Aksenov, Eve Bodnia, Michael H. Freedman, and Michael Mulligan. Compression is all you need: Modeling mathematics. *arXiv preprint arXiv:2603.20396*, 2026. URL https://arxiv.org/abs/2603.20396.

Dana Angluin. Inductive inference of formal languages from positive data. *Information and Control*, 45(2):117–135, 1980. doi: 10.1016/S0019-9958(80)90285-5. URL https://www.sciencedirect.com/science/article/pii/S0019995880902855.

Anthropic. Formalizing Fermat's Last Theorem, 2026. URL https://www.anthropic.com/research/formalizing-fermats-last-theorem. Accessed September 20, 2026.

Artur Avila, Manjul Bhargava, Caucher Birkar, et al. A severe misalignment of AI in mathematics. Open letter, September 11; also published on Terence Tao's blog, 2026. URL https://terrytao.wordpress.com/2026/09/11/a-severe-misalignment-of-ai-in-mathematics/.

Edoardo Botta, Yuchen Li, Aashay Mehta, Jordan T. Ash, Cyril Zhang, and Andrej Risteski. On the query complexity of verifier-assisted language generation. In *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pp. 5124–5155, 2025. URL https://proceedings.mlr.press/v267/botta25a.html.

Ziyi Cai, Shuangping Li, Yiheng Shen, Kangning Wang, and Peng Zhang. Dense language generation made simple: Deterministic, randomized, and multi-order algorithms. *arXiv preprint arXiv:2608.01320*, 2026. URL https://arxiv.org/abs/2608.01320.

Simon Colton, Alan Bundy, and Toby Walsh. On the notion of interestingness in automated mathematical discovery. *International Journal of Human-Computer Studies*, 53(3):351–375, 2000. doi: 10.1006/ijhc.2000.0394. URL https://www.sciencedirect.com/science/article/pii/S107158190090394X.

Atul Ganju, Travis McVoy, Shaddin Dughmi, and Shang-Hua Teng. A theory of time-sensitive language generation: Sparse hallucination beats mode collapse. *arXiv preprint arXiv:2605.11302*, 2026. URL https://arxiv.org/abs/2605.11302.

E. Mark Gold. Language identification in the limit. *Information and Control*, 10(5):447–474, 1967. doi: 10.1016/S0019-9958(67)91165-5. URL https://www.sciencedirect.com/science/article/pii/S0019995867911655.

Alkis Kalavasis, Anay Mehrotra, and Grigoris Velegkas. On characterizations for language generation: Interplay of hallucinations, breadth, and stability. In *Proceedings of the 37th International Conference on Algorithmic Learning Theory*, volume 313 of *Proceedings of Machine Learning Research*, pp. 1–49, 2026. URL https://proceedings.mlr.press/v313/kalavasis26a.html.

Jon Kleinberg and Sendhil Mullainathan. Language generation in the limit. In *Advances in Neural Information Processing Systems*, volume 37, 2024. doi: 10.52202/079017-2111. URL https://proceedings.neurips.cc/paper_files/paper/2024/hash/7988e9b3876ad689e921ce05d711442f-Abstract-Conference.html.

Jon Kleinberg and Fan Wei. Density measures for language generation. In *2025 IEEE 66th Annual Symposium on Foundations of Computer Science (FOCS)*, pp. 620–658, 2025. doi: 10.1109/FOCS63196.2025.00034. URL https://arxiv.org/abs/2504.14370.

Jon Kleinberg and Fan Wei. Language generation and identification from partial enumeration: Tight density bounds and topological characterizations. In *Proceedings of the 58th Annual ACM Symposium on Theory of Computing*, pp. 1684–1691, 2026. doi: 10.1145/3798129.3800876. URL https://arxiv.org/abs/2511.05295.

Guillaume Lample, Timothée Lacroix, Marie-Anne Lachaux, Aurélien Rodriguez, Amaury Hayat, Thibaut Lavril, Gabriel Ebner, and Xavier Martinet. HyperTree proof search for neural theorem proving. In *Advances in Neural Information Processing Systems*, volume 35, 2022. doi: 10.52202/068431-1910. URL https://proceedings.neurips.cc/paper_files/paper/2022/hash/a8901c5e85fb8e1823bbf0f755053672-Abstract-Conference.html.

Shuangping Li and Peng Zhang. Language generation in the limit: A Lean library and paper map. https://github.com/pengzhang91/generation-in-the-limit-lib, 2026. Lean library, Apache-2.0; source snapshot documented in the supplement.

OpenAI. On the Navier–Stokes Millennium Prize Problem. Research announcement, September 8, 2026a. URL https://openai.com/index/navier-stokes-solution/.

OpenAI. Ten advances in mathematics and theoretical computer science. Research announcement, August 1, 2026b. URL https://openai.com/index/ten-advances-in-mathematics/.

Debmalya Panigrahi, Fan Wei, and Ian Zhang. Hallucination rates in language generation. *arXiv preprint arXiv:2607.23361*, 2026. URL https://arxiv.org/abs/2607.23361.

Henri Poincaré. *Science and Method*. Thomas Nelson and Sons, London, Edinburgh, Dublin, and New York, 1914. English translation of *Science et méthode*.

Stanislas Polu and Ilya Sutskever. Generative language modeling for automated theorem proving. *arXiv preprint arXiv:2009.03393*, 2020. URL https://arxiv.org/abs/2009.03393.

Bernardino Romera-Paredes, Mohammadamin Barekatain, Alexander Novikov, Matej Balog, M. Pawan Kumar, Emilien Dupont, Francisco J. R. Ruiz, Jordan S. Ellenberg, Pengming Wang, Omar Fawzi, Pushmeet Kohli, and Alhussein Fawzi. Mathematical discoveries from program search with large language models. *Nature*, 625:468–475, 2024. doi: 10.1038/s41586-023-06924-6. URL https://www.nature.com/articles/s41586-023-06924-6.

George Tsoukalas, Rahul Saha, Amitayush Thakur, Sabrina Reguyal, and Swarat Chaudhuri. Learning interestingness in automated mathematical theory formation. In *Advances in Neural Information Processing Systems*, volume 38, 2025. doi: 10.52202/085713-1298. URL https://proceedings.neurips.cc/paper_files/paper/2025/hash/37678356aa009dd4205c813c02fa4985-Abstract-Conference.html.

Kaiyu Yang, Aidan M. Swope, Alex Gu, Rahul Chalamala, Peiyang Song, Shixing Yu, Saad Godil, Ryan J. Prenger, and Anima Anandkumar. LeanDojo: Theorem proving with retrieval-augmented language models. In *Advances in Neural Information Processing Systems*, volume 36, 2023. doi: 10.52202/075280-0944. URL https://proceedings.neurips.cc/paper_files/paper/2023/hash/4441469427094f8873d0fecb0c4e1cee-Abstract-Datasets_and_Benchmarks.html.

TABLE OF CONTENTS

1 Introduction **2**

2 Worlds, records, and information **3**

3 Verification localizes **4**

4 The exchange rate **5**

4.1 Contested and uncontested points 6

4.2 The frontier 6

4.3 Two values 7

4.4 The clock 7

5 A value rule that hides almost everything **8**

6 Related work **9**

7 Conclusion **9**

A Formal framework **15**

A.1 Notation and proof dependencies 15

A.2 Universe and interaction protocol 16

A.3 Candidate pairs and verifier queries 16

A.4 Partial records, coverage and trivia 16

A.5 Set-valued tasks and tell-tales 17

A.6 Transcript invariance and the ledger 17

A.7 External results used in the proofs 18

B Verification localizes: proofs **18**

B.1 Proof of the localization theorem 18

B.2 Breadth and sound coverage 19

B.3 Both parts of the localization theorem are sharp 19

B.4 Gluing fails for criteria not judged in the limit 21

C The exchange rate: proofs **22**

C.1 The race lemma 22

C.2 Proof of the exchange theorem 23

C.3 Proof of the two-value corollary 24

C.4 Proof of the clock 24

C.5 Boundary checks and the role of each assumption 25

C.6 Persistent exploration harvests in every family 25

D Compression constructions and counting **26**
D.1 Mathematical background . . . 26
D.2 Abelian refinement chains . . . 27
D.3 Free-monoid pairs . . . 29
D.4 Imprecision of exhaustive generation . . . 31

E Numerical illustration **32**

F Lean formalization **33**

G Extended context and related work **36**
G.1 Theorem-level comparison . . . 36
G.2 Generation in the limit and density . . . 36
G.3 Errors, hallucination rates and timely generation . . . 36
G.4 Verifiers, interestingness and automated discovery . . . 38
G.5 Compression as a model of mathematics . . . 38

H How to read the results **38**
H.1 The words of the title . . . 38
H.2 The quantifiers . . . 38
H.3 Counts, rates and clocks . . . 39
H.4 What is fixed by the model . . . 39

I Discussion **39**
I.1 Coverage is not selection . . . 39
I.2 Advice within the model . . . 39
I.3 Open directions . . . 40

GUIDE TO THE APPENDICES

Appendix A fixes the formal framework: notation and a map of proof dependencies, the protocol, trivia counts, budgets and regimes, the transcript lemma, and the two external results. The two organizing theorems are then proved in the order of the main text. Appendix B proves that verification localizes, together with the separations showing that both halves of the localization theorem are sharp. Appendix C proves the exchange-rate frontier, the two-value corollary and the clock, checks the role of each assumption, and shows that persistent exploration harvests in every countable family. Appendix D gives the compression constructions and their counting proofs.

The remaining appendices serve other purposes. Appendix E documents deterministic computations of the frontier and the clock, and Appendix F records the scope of the Lean development; neither substitutes for a proof. Appendix G compares the results with prior work theorem by theorem, Appendix H fixes the reading of the title and the quantifiers, and Appendix I discusses selection, advice within the model and open directions.

A FORMAL FRAMEWORK

A.1 NOTATION AND PROOF DEPENDENCIES

Table 1 collects the notation, and Figure 4 shows which results are used in the proofs of which others.

Table 1: Notation, with the place where each object is introduced.

Symbol	Meaning	Where
U; $x_1 \prec x_2 \prec \cdots$; $L[1\ldots M]$	countable universe; canonical order; first M elements of L	Appendix A
w_t, a_t; S_t, O_t; O	observation and output in round t; seen and output sets; limit set	Appendix A
$\mathcal{P}$; (F, H); $\mathcal{P}_F$, $\mathcal{P}_F$	pair model; formal world and target; validity fiber and its targets	Section 2, Appendix A
$\mathbf{1}_F$; G_t	verifier oracle; set-valued output	Appendix A
T_L; Angluin, weak Angluin	tell-tale of L; the two breadth conditions	Appendix A.5
C; E; $\mathcal{E}(K)$	observed core; its enumeration; complete repetition-free records of K	Section 4
$\underline{d}(A, L)$, $\overline{d}(A, L)$, $d(A, L)$	lower, upper and exact density of A in L	(5)
$b_H(T)$, $B_H(T)$	distinct trivia and trivial rounds by T at truth H	(6)
g; g^{-1}; $\mathcal{G}_g$; $\mathcal{R}$	budget; its inverse; budgeted class; trivia regime	Section 4, Appendix A
K, K', D; m_N, n_N; α	nested pair and difference; recorded and unrecorded prefix counts; limit of m_N/N	Section 4
$Q_{G,E}(T, N)$; $\Phi_g(T, N)$	prefix coverage of K' by round T; the exchange frontier	Theorem 4.1
E_G; t_M, τ_M, $w(t)$	racing record against G; resolution time, Mth race round, non-race rounds	Lemma C.1
G_g; $\tau_g(M)$; $\tau_N(\varepsilon)$	race-and-sweep generator; its Mth race round, $\min\{t : t - g(t) \geq M\}$; recovery time	Theorem 4.1, Corollary 4.3
$\mathcal{H}_t$, J_t, U_t; G_{sw}	consistent candidates, identified intersection and union; universal sweep	Appendix C.6
A_n, F_n; G, D; $\lvert w\rvert_G$, $\lvert w\rvert_{G\cup D}$, ℓ_s	monoids; generators and dictionary; raw, wrapped and dictionary-only length	Appendix D.1.2
H_D; θ_κ; D_c; $s_c(N)$; $B_G(r)$	compressible language; cut; geometric dictionary; digit sum; raw ball	Appendix D.1.2
$\mu_{m,i}$, $\widehat{\mu}_{m,i}$; $\widehat{K}$, $\widehat{K}'$	tag words; tag-restricted pair	Section 5, Proposition D.5
L^*, A, $R_{\geq j}$, H_j, F_j; $\mathcal{P}^\dagger$	separation construction	Appendix B.3

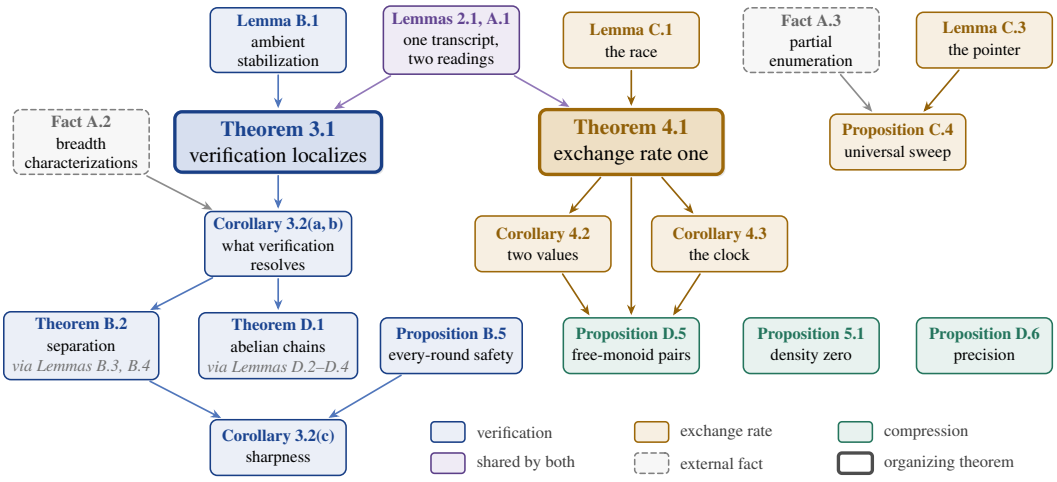


Figure 4: Dependency graph of the results. An arrow from one result to another means that the first is used in the proof of the second. Colors mark the strands, verification (blue), the exchange rate (gold) and compression (green); thick boxes are the two organizing theorems, the plum box is the transcript lemma they share, and dashed boxes are the external facts of Appendix A.7. Every numbered result appears, and a lemma used in a single appendix proof is listed inside the node it serves. Corollary 3.2 is drawn twice: parts (a) and (b) follow from Theorem 3.1, and the sharpness in part (c) comes from Theorem B.2 and Proposition B.5. Proposition 5.1 and Proposition D.6 are counting results that stand alone; their consequences for coverage are read off from Theorem 4.1 and Corollary 4.2.

A.2 Universe and interaction protocol

Fix a countable universe $U = \{x_1 \prec x_2 \prec \cdots\}$ in a canonical order of type ω. All target languages are infinite, and $L[1\ldots M]$ denotes the first M elements of an infinite set $L \subseteq U$; for $L \subseteq L'$ the induced order gives $L \cap L'[1\ldots N] = L[1\ldots m]$ with $m = |L \cap L'[1\ldots N]|$. Rounds are indexed by integers $t \geq 1$, with $S_0 = O_0 = \varnothing$. In round t the generator observes w_t and then emits an output before w_{t+1} arrives; write $S_t = \{w_1, \ldots, w_t\}$ and $O_t = \{a_1, \ldots, a_t\}$. A generator G is a deterministic strategy on the entire ordered history, which includes its own outputs and all verifier queries and answers. The candidate model is given to G; the true pair and the core are not. Distinct statements are counted under one fixed canonical encoding. No computability or efficiency constraint is imposed, and adversaries may be non-effective.

A.3 Candidate pairs and verifier queries

A pair model is a countable collection $\mathcal{P} = \{(F_i, H_i)\}_i$ with $H_i \subseteq F_i \subseteq U$. Its fiber over F is $\mathcal{P}_F = \{(F_i, H_i) : F_i = F\}$, and $\mathcal{P}_F = \{H_i : F_i = F\}$ is the family of targets in that fiber; ambients are compared as sets. The adversary fixes a true pair (F, H). The membership oracle $\mathbf{1}_F$ returns 1 on $q \in F$ and 0 on $q \notin F$. A verifier-assisted generator makes finitely many adaptive queries to $\mathbf{1}_F$ per round, with no fixed bound on their number; each round terminates after finitely many queries along every realized interaction. The verifier is an oracle for the formal world, not for value.

A.4 Partial records, coverage and trivia

For element generation the adversary enumerates an infinite core $C \subseteq H$ without repetition, every point of C appearing eventually; this is the partial-enumeration protocol of Kleinberg & Wei (2026). The generator emits one point $a_t \notin S_t$ per round. Freshness excludes observations, not earlier outputs: a generator may repeat its own unobserved outputs, although the generators we construct never do. Points already output must still appear in the record eventually if they belong to the core. Let $O = \bigcup_t O_t$ be the limit set, and for $A \subseteq U$ define

$$\underline{d}(A, L) = \liminf_{M\to\infty} \frac{|A \cap L[1\ldots M]|}{M}, \qquad \overline{d}(A, L) = \limsup_{M\to\infty} \frac{|A \cap L[1\ldots M]|}{M}, \tag{5}$$

writing $d(A, L)$ when the limit exists. Coverage at truth H is $c = \underline{d}(O, H)$, and its finite form is the prefix count $|O_T \cap H[1..N]|$. A legal instance specifies the truth, the core and a repetition-free enumeration of the core. The core is not given to the generator, and every theorem that restricts evaluation to a particular core says so.

An output lies in exactly one of H (valuable), $F \setminus H$ (trivial) or $U \setminus F$ (invalid). Two trivia counts are used:

$$b_H(T) = |O_T \cap (F \setminus H)|, \qquad B_H(T) = |\{t \le T : a_t \in F \setminus H\}|. \tag{6}$$

The first counts distinct trivial statements and the second counts trivial rounds; $b \le B$, and the two coincide for generators that never repeat an output. Both are nondecreasing, vanish at $T = 0$ and grow by at most one per round. A *budget* is a function $g : \mathbb{N} \to \mathbb{N}$ with these three properties. Every nondecreasing g' can be replaced by the largest budget below it, $g(T) = \min_{0 \le t \le T}\big(g'(t)\mathbf{1}\{t > 0\} + T - t\big)$, because a trivia count obeys g' at every round exactly when it obeys g. A *trivia regime* is a nonempty set $\mathcal{R}$ of nondecreasing functions $\mathbb{N} \to \mathbb{N}$ closed under passing to smaller nondecreasing functions. A generator obeys $\mathcal{R}$ if, on every legal instance of the model, its profile $T \mapsto b_H(T)$ lies in $\mathcal{R}$, where H is the truth of that instance. Zero trivia ($\{0\}$), finitely many trivia (the bounded functions), a prescribed budget ($\{f : f \le g\}$) and a vanishing rate ($\{f : f(T)/T \to 0\}$) are regimes. The same definitions with B in place of b give round-count regimes; Corollary 4.2 holds for both (Appendix C.3).

A.5 Set-valued tasks and tell-tales

Breadth and sound coverage. For the breadth task the stream enumerates all of H, repetitions allowed, and outputs are sets $G_t \subseteq U \setminus S_t$. Exact breadth means that $G_t = H \setminus S_t$ eventually. Approximate breadth means that eventually $G_t \subseteq H$ and $|H \setminus G_t| < \infty$. Sound coverage relaxes exact breadth to the sandwich

$$H \setminus S_t \;\subseteq\; G_t \;\subseteq\; F \qquad \text{for all sufficiently large } t. \tag{7}$$

These are set-valued tasks. They imply no density guarantee for a generator restricted to one output per round.

Tell-tales. A countable family $\mathcal{H}$ satisfies *Angluin's condition* if every $L \in \mathcal{H}$ has a finite $T_L \subseteq L$ such that no $L' \in \mathcal{H}$ satisfies $T_L \subseteq L' \subsetneq L$ (Angluin, 1980). The *weak* version requires only that $T_L \subseteq L' \subsetneq L$ imply $|L \setminus L'| < \infty$. Exact and approximate breadth are equivalent to these two conditions, respectively (Kalavasis et al., 2026, Theorems 3.3 and 3.8). The restriction to sublanguages of L is part of both definitions, and an ambient verifier is not a positive-and-negative presentation of the target.

A.6 Transcript invariance and the ledger

Lemma A.1 (One transcript, two readings)**.** Fix a candidate pair model containing (F, K) and (F, K'), where $K \subseteq K' \subseteq F \subseteq U$ and K is infinite. Let $E = (w_t)_{t \ge 1}$ list every element of K exactly once. For a fixed deterministic verifier-assisted element generator G, compare the instances

$$\mathcal{I}_0 = ((F, K), K, E), \qquad \mathcal{I}_1 = ((F, K'), K, E),$$

whose entries are the true pair, the observed core and the observation sequence. Both instances are legal, and in the two runs every round has the same ordered sequence of queries, oracle answers and output. In particular the two runs share the output sets $O_T(G, E)$ and the limit set $O(G, E) = \{a_t : t \ge 1\}$, and for every T

$$b_K(T) - b_{K'}(T) = |O_T \cap (K' \setminus K)| = |O_T \cap K'| - |O_T \cap K|.$$

The same identity holds for round counts, $B_K(T) - B_{K'}(T) = \sum_{t \le T} \mathbf{1}\{a_t \in K' \setminus K\}$. Equality of transcripts therefore does not assert equality of trivia counts.

Proof. The sequence E is a repetition-free enumeration of the infinite core K. Since K is contained in each target, both instances are legal. The candidate model and the initial state of G are the same in both runs, and both verifier oracles are the function $\mathbf{1}_F$.

We prove that the completed round histories agree, by induction on t. The histories before round 1 are empty. Suppose the histories through round $t-1$ agree. Both runs next receive the same observation w_t. For the interaction within round t, induct on the number of oracle queries already answered. Before any query the available histories agree. If they agree after j answers, determinism makes the next action identical. If that action queries a point $q \in U$, both oracles return $\mathbf{1}_F(q)$, which preserves agreement after $j+1$ answers. If the action emits an output, both runs emit the same a_t and end the round. The finite-query requirement ensures that the round terminates. Hence the histories through round t agree, which completes the outer induction.

Equality of every output gives the common sets O_T and $O(G, E)$. Finally, the disjoint union $F \setminus K = (F \setminus K') \dot\cup (K' \setminus K)$ gives both displayed identities, the first by counting distinct elements of O_T and the second by summing indicators over $t \le T$. □

A trivia constraint imposed on the truth-K instance therefore controls the same output set that is evaluated on the truth-K' instance. Every lower bound in Section 4 is a statement about one output set under two readings. None of them asserts that a run at truth K' contains trivia.

A.7 External results used in the proofs

We restate, in the form in which they are used, the two external results that the proofs invoke. Results cited only for context are not restated: the identification theory of Gold (1967) and Angluin (1980), the generation model of Kleinberg & Mullainathan (2024), the timely-generation and hallucination-rate results of Ganju et al. (2026); Panigrahi et al. (2026), and the compression framework of Aksenov et al. (2026). None of these is invoked in a proof.

Fact A.2 (Breadth characterizations; Kalavasis et al., 2026)**.** A countable family admits exact breadth in the limit if and only if it satisfies Angluin's condition. It admits approximate breadth in the limit if and only if it satisfies the weak condition. These are Theorems 3.3 and 3.8 of the cited paper, with the definitions of Appendix A.5.

This fact is used in Corollary 3.2(a), in both directions, and through Corollary 3.2(a) in Theorem D.1.

Fact A.3 (Generation from partial enumeration; Kleinberg & Wei, 2026)**.** For every countable target family and every partial-enumeration instance with infinite core $C \subseteq H$, there is a generator that is eventually valid, emitting elements of H from some finite round on, and whose limit set satisfies $\underline{d}(O, H) \ge \frac{1}{2}\underline{d}(C, H)$. Its construction races on identified intersections of consistent candidates, which become eventually valid. The same paper shows that no coverage guarantee above one half of the core is attainable against an adaptive racing adversary.

This fact is used in Proposition C.4, only through the identified-intersection principle, which is proved there in the form needed. The racing record and the two-sided race bounds of Lemma C.1 are proved in Appendix C.1 in the form needed here and are not imported.

B Verification localizes: proofs

This appendix proves Theorem 3.1 and Corollary 3.2. The pair model and the set-valued tasks are defined in Appendices A and A.5, and Fact A.2 is the only external input. The proofs separate the two roles of a verifier: it selects the correct formal world in the limit, and it supplies no further distinction between targets inside that world.

B.1 Proof of the localization theorem

Fix a protocol, element-valued with partial records or set-valued with complete records, and a success criterion S: a set of pairs (instance, output sequence), fixed in advance and used unchanged on $\mathcal{P}$ and on each of its fibers. Records are fixed sequences; an adaptive adversary against a deterministic generator is one such sequence, so this loses nothing. An output is *legal* at a history if it avoids the observed set, $a_t \notin S_t$ or $G_t \subseteq U \setminus S_t$. Legality therefore depends only on the observations, and a legal output always exists. A generator *succeeds on* a set of worlds if its run on every legal instance whose true world lies in the set produces only legal outputs and belongs to S. The criterion is *judged in the limit* if $(\iota, (y_t)_t) \in \mathsf{S}$ and $y_t = y'_t$ for all large t imply $(\iota, (y'_t)_t) \in \mathsf{S}$.

Stripping. Fix F and a verifier-assisted generator G. Define the ordinary generator G^F as follows: on an observed history, replay the deterministic interaction of G and answer every query q with $\mathbf{1}_F(q)$. Each realized round of this replay is a realized round of G on an instance with formal world F, so it ends after finitely many queries; no uniform bound over hypothetical oracle paths is needed. The induction in the proof of Lemma A.1 shows that G^F and G produce the same output in every round of every instance with formal world F. This proves Theorem 3.1(i). It also proves necessity in (ii): if G succeeds on $\mathcal{P}$, then G^F succeeds on $\mathcal{P}_F$.

Observation-only form. An ordinary deterministic generator Γ reads its own earlier outputs, but these are themselves functions of the observations. Define $\bar{\Gamma}(w_1, \ldots, w_t)$ recursively as the output that Γ produces when it is fed $w_1, \ldots, w_t$ together with the outputs $\bar{\Gamma}(w_1), \ldots, \bar{\Gamma}(w_1, \ldots, w_{t-1})$. On every run, Γ outputs exactly $\bar{\Gamma}(w_1, \ldots, w_t)$ in round t.

Lemma B.1 (Ambient stabilization)**.** List the candidate formal worlds as $F_1, F_2, \ldots$, with repetitions allowed. In round n, query the new point x_n and retain earlier answers, and let

$$\widehat{z}_n := \min\{i \geq 1 : F_i \cap \{x_1, \ldots, x_n\} = F \cap \{x_1, \ldots, x_n\}\}.$$

For $z_0 = \min\{i : F_i = F\}$ there is a finite N_1 with $\widehat{z}_n = z_0$ for every $n \geq N_1$.

Proof. Index z_0 is always consistent, so $\widehat{z}_n \leq z_0$. For each $i < z_0$ choose the least index $m(i)$ with $x_{m(i)} \in F_i \triangle F$; candidate i is inconsistent from round $m(i)$ on. Put $N_1 = \max(\{1\} \cup \{m(i) : 1 \leq i < z_0\})$. From round N_1 on, every smaller candidate has been eliminated and z_0 remains. Duplicate ambients cause no difficulty. The lemma asserts stabilization; it does not give the generator a test for whether the current guess is already final. □

Gluing. Suppose S is judged in the limit and that for every formal world F occurring in $\mathcal{P}$ an ordinary generator Γ_F succeeds on $\mathcal{P}_F$. There are countably many such worlds, and we fix one Γ_F for each, using countable choice; we do not claim that these functions are effective. Define the verifier-assisted generator G: in round t, query x_t, compute $\widehat{z}_t$ as in Lemma B.1, and let $y_t = \bar{\Gamma}_{F_{\widehat{z}_t}}(w_1, \ldots, w_t)$. Output y_t if it is legal at the current history, and otherwise output a fixed legal replacement (the least point of $U \setminus S_t$, or $y_t \setminus S_t$). Consider any legal instance with true world (F, H). By Lemma B.1, $F_{\widehat{z}_t} = F$ for all $t \geq N_1$. The history is a legal history of an instance of $\mathcal{P}_F$, on which Γ_F outputs only legal outputs, so from round N_1 on no replacement occurs and G outputs exactly what Γ_F outputs on the same instance. Since Γ_F succeeds and S is judged in the limit, G succeeds. The history is never restarted when the guess changes, and G makes one new query per round. This proves Theorem 3.1(ii). □

B.2 Breadth and sound coverage

Proof of Corollary 3.2(a). Exact breadth and approximate breadth depend only on the target and the outputs, and both are judged in the limit. An ordinary generator succeeds on $\mathcal{P}_F$ exactly when it achieves the corresponding breadth for the family $\mathcal{P}_F$ of targets in that fiber. By Theorem 3.1, a verifier-assisted generator achieves breadth on $\mathcal{P}$ if and only if each family $\mathcal{P}_F$ admits ordinary breadth, which by Fact A.2 happens exactly when $\mathcal{P}_F$ satisfies Angluin's condition, or its weak form for approximate breadth. For a model with a single ambient, these are the oracle-free characterizations; with several ambients, the verifier reduces the problem to them. □

Proof of Corollary 3.2(b). Sound coverage (7) is judged in the limit. On $\mathcal{P}_F$ the ordinary generator $G_t = F \setminus S_t$ satisfies $H \setminus S_t \subseteq G_t \subseteq F$ in every round, because $H \subseteq F$. Gluing gives a verifier-assisted generator that outputs $F \setminus S_t$ from its stabilization round on. Since the record is positive, $S_t \subseteq H$, so from that round on the trivia in the output are exactly $F \setminus H$, an infinite set whenever that difference is infinite. No claim is made about soundness of the outputs before stabilization. If the entire F is supplied as a parameter instead, the same sandwich holds from round 1. □

B.3 Both parts of the localization theorem are sharp

Gluing needs the verifier: without it, the formal world cannot be identified, and the following model has no eventual sound coverage at all. To separate assisted from unassisted coverage we need formal

worlds that demand incompatible output sets after the same finite record. An increasing chain of targets keeps every finite record of the union compatible with some member of the chain. Removing infinitely many points of the union from that member's ambient makes their coverage requirements incompatible even after finitely many observations. A second column supplies a common infinite trivia set.

Work on $U = \mathbb{N} \times \{0, 1\}$, taking $0 \in \mathbb{N}$. Put

$$\begin{aligned} L^* &= \mathbb{N} \times \{0\}, \qquad A = \mathbb{N} \times \{1\}, \\ R_{\geq j} &= \{(2m, 0) : m \geq j\}, \qquad H_j = L^* \setminus R_{\geq j} \quad (j \geq 0), \\ F_j &= H_j \cup A, \qquad F_\infty = L^* \cup A = U. \end{aligned}$$

Using the same infinite A in every pair is sufficient; no disjointness condition on auxiliary trivia sets is needed.

Theorem B.2 (No sound coverage without a verifier)**.** The countable pair model $\mathcal{P}^\dagger = \{(F_j, H_j) : j \geq 0\} \cup \{(F_\infty, L^*)\}$ has infinite targets and every difference $F \setminus H$ equal to the infinite set A, and no oracle-free set-valued generator achieves sound coverage in the limit on it. Consequently a verifier strictly enlarges the class of pair models admitting eventual sound coverage, to all countable pair models.

Lemma B.3 (Sandwich-exclusion)**.** For any $j \geq 0$, finite $S \subseteq U$, and set $B \subseteq U$, the following two sandwiches cannot both hold:

$$L^* \setminus S \subseteq B \subseteq F_\infty, \qquad H_j \setminus S \subseteq B \subseteq F_j.$$

Proof. The lower inclusion of the first sandwich and upper inclusion of the second give $L^* \setminus S \subseteq F_j$. Intersecting with column 0 gives $L^* \setminus S \subseteq H_j$, because A lies in column 1. Hence $R_{\geq j} = L^* \setminus H_j \subseteq S$, contradicting that $R_{\geq j}$ is infinite and S is finite. □

Lemma B.4 (Fresh-index availability)**.** For every finite $T \subseteq L^*$ and every integer J, there is $j > J$ with $T \subseteq H_j \subsetneq L^*$ and $|L^* \setminus H_j| = \infty$. The targets form a strictly increasing chain with union L^*.

Proof. Take $j > J$ larger than each m for which $(2m, 0) \in T$. Then $T \cap R_{\geq j} = \varnothing$, so $T \subseteq H_j$. Every $R_{\geq j}$ is infinite. Every H_j is infinite because it contains all odd points in column 0. Passing from j to $j + 1$ adds exactly $(2j, 0)$, and every even point is eventually added. Thus the union is L^* and the inclusions are strict. □

Proof of Theorem B.2. Let

$$\mathcal{P}^\dagger = \{(F_j, H_j) : j \geq 0\} \cup \{(F_\infty, L^*)\}$$

with the parameters above. All languages are infinite, and every formal minus valuable difference equals the infinite set A. All ambients are distinct, so the fibers are singletons.

Fix an arbitrary deterministic oracle-free set-valued generator G. Let $E^* = (e_1, e_2, \ldots)$ be a fixed repetition-free listing of L^*. We build an actual ordered history by phases. At the end of a completed phase ℓ, maintain that its seen-set is $\{e_1, \ldots, e_{n_\ell}\}$ and its length is n_ℓ. The history itself need not be in E^*-order. Start with the empty history, $n_0 = 0$, and $j_0 = -1$.

Suppose phases before ℓ have completed. Choose, by Lemma B.4, an index $j_\ell > \max(j_{\ell-1}, n_{\ell-1})$ for which the current finite seen-set lies in H_{j_ℓ}. Consider the infinite continuation obtained by appending the as-yet unseen elements of H_{j_ℓ} in E^*-order. Together with the current ordered history, this is a legal full enumeration of H_{j_ℓ}. Examine only rounds strictly after $n_{\ell-1}$ on this continuation.

If none of these rounds satisfies

$$H_{j_\ell} \setminus S_n \subseteq G_n(\sigma_n) \subseteq F_{j_\ell},$$

choose (F_{j_ℓ}, H_{j_ℓ}) as the target and use this entire continuation. The sandwich fails at every subsequent round, so G fails sound coverage in the limit. This completes the counterexample in this case.

Otherwise let $n'_\ell > n_{\ell-1}$ be the first round on that continuation satisfying the sandwich, and append exactly the finite continuation through that round. By Lemma B.3, at this same ordered history G fails the sound-coverage sandwich for (F_∞, L^*).

To finish the phase, let P_ℓ be the largest E^*-position of a point shown through round n'_ℓ. Choose an integer

$$M_\ell > \max(P_\ell, n'_\ell, \ell)$$

and append in E^*-order all points among $e_1, \dots, e_{M_\ell}$ not already shown. This retains the actual earlier order and introduces no repetitions. The resulting seen-set is exactly the first M_ℓ points of E^*, so its length is $n_\ell = M_\ell$ and the invariant holds. In particular $n_\ell > n'_\ell > n_{\ell-1}$.

If any phase took the infinite-failure continuation, we have already obtained a legal target and enumeration defeating G. Otherwise all phases finish, $n_\ell \to \infty$, and their nested ordered histories define a repetition-free full enumeration of L^*: the catch-up step eventually includes every e_k. At each of the infinitely many distinct rounds n'_ℓ, G fails sound coverage for (F_∞, L^*). Choose that pair and the constructed infinite enumeration. Again G fails in the limit. The construction is an existence argument; deciding whether a successful continuation round exists need not be effective. It constructs one fixed target and one complete enumeration as a function of G; these can be fixed before the resulting interaction, so the target never changes along the realized run. Since G was arbitrary, the separation follows.

Combined with Corollary 3.2(b), this gives the strict enlargement. The separating model varies the ambient and has singleton fibers; on any fixed fiber, stripping preserves all outputs and therefore preserves sound coverage as well. □

B.4 Gluing fails for criteria not judged in the limit

Proposition B.5 (Every-round safety is a stronger requirement)**.** There is a countable pair model with infinite trivia complements on which no finite-query set-valued generator is both formally sound at every round on every instance and sound-covering in the limit on every instance.

Proof. Let $U = \mathbb{N} \times \{0, 1\}$, $L^* = \mathbb{N} \times \{0\}$, and $A = \mathbb{N} \times \{1\}$. Use the pairs

$$(U, L^*) \quad \text{and} \quad (U \setminus \{x\}, L^* \setminus \{x\}) \quad (x \in L^*).$$

Every formal minus valuable difference equals the infinite set A. Suppose G is sound at every round on every instance of this model. Run it on any full enumeration of L^* with oracle $\mathbf{1}_U$. At a finite round n, let Q_n contain all points queried in rounds $1, \dots, n$. It is finite, as is S_n. Take any $x \in L^* \setminus (Q_n \cup S_n)$. The same ordered positive history is legal for $L^* \setminus \{x\}$ and extends to a full enumeration of that language. Replacing the oracle by $\mathbf{1}_{U \setminus \{x\}}$ changes none of the answers along the entire interaction through round n: the oracles differ only at x, which was not queried. By determinism the output set is unchanged. Soundness on this alternative instance forces x to be absent from that output. There are infinitely many such x, so the output on (U, L^*) does not cover $L^* \setminus S_n$. This holds at every round n, contradicting eventual sound coverage on (U, L^*). □

The criterion "sound at every round and eventually sound-covering" is not judged in the limit, and the proposition shows that gluing fails for it. On each fiber of this model the ordinary generator $G_t = F \setminus S_t$ is sound at every round and covers soundly, yet no verifier-assisted generator does both on the whole model. (Soundness at every round alone glues trivially, since $G_t = \varnothing$ is always sound.) The obstruction is exactly the finite information available before the formal world is identified.

The same obstruction defeats prescribed trivia budgets, which is why Section 4 fixes the formal world. On $U = \mathbb{N} \times \{0, 1\}$ with $L^* = \mathbb{N} \times \{0\}$, take the pair (U, L^*) together with the pairs $(U \setminus \{(\langle j, M\rangle, 1)\}, L^* \setminus M)$ for all $j \in \mathbb{N}$ and all finite $M \subseteq L^*$, where $\langle \cdot, \cdot \rangle$ is an injective code. Each fiber is a single pair, on which outputting the least unobserved point of the target produces no trivia at all. Now let G be any verifier-assisted element generator. Run it at truth (U, L^*) on the record that always reveals the least point of L^* that is neither observed nor output. Fix T, let M be the set of outputs in L^* by round T, and choose j so that $(\langle j, M\rangle, 1)$ was neither queried nor output. At truth $(U \setminus \{(\langle j, M\rangle, 1)\}, L^* \setminus M)$ the same observations form a legal prefix, because they avoid M, and every query receives the same answer. So the first T rounds coincide, and every one

of the T outputs is trivia there. Hence no verifier-assisted generator keeps $B_H(T) \le g(T)$ on every instance once $g(T) < T$ for some T, although every fiber admits zero trivia. Conversely, when every occurring formal world is isolated, meaning that finitely many queries separate it from all others (for instance, when only finitely many formal worlds occur), the generator can identify F with certainty during the first round, and gluing then holds for every criterion.

C THE EXCHANGE RATE: PROOFS

Throughout this appendix $F = U$, so every output is valid and the verifier is silent. The objects $K \subsetneq K'$, $D = K' \setminus K$, m_N, n_N, $\mathcal{E}(K)$, $Q_{G,E}(T, N)$, budgets and $\mathcal{G}_g$ are those of Section 4. Because $K \subseteq K'$ carry the induced order,

$$K'[1 \dots N] = K[1 \dots m_N] \dot\cup D[1 \dots n_N]. \tag{8}$$

The running family has $K' = \{0, 2, 4, \dots\}$ with $K'[j] = 2(j-1)$ and, for $0 < \alpha < 1$,

$$K = \{K'[\lceil k/\alpha \rceil] : k \ge 1\}, \qquad m_N = \lfloor \alpha N \rfloor, \tag{9}$$

because $\lceil k/\alpha \rceil \le N$ exactly when $k \le \alpha N$. The selected ranks are distinct for $\alpha < 1$, so D is infinite. No computability is asserted for irrational α.

C.1 THE RACE LEMMA

A point x is *resolved* by round t if $x \in S_t \cup O_t$. Following Kleinberg & Wei (2026), the *racing record* A_K usually reveals the least element of $K \setminus (S_{t-1} \cup O_{t-1})$. At rounds $t = 2^j$ it instead reveals the least element of $(O_{t-1} \cap K) \setminus S_{t-1}$, recycling a point that the generator has claimed; if no such point exists, it reveals the fresh point as usual. Recycling makes the record complete without materially changing the race.

Lemma C.1 (Race lemma; adapted from Kleinberg & Wei (2026))**.** Let K be infinite.

(R1) Suppose a generator outputs the least element of $K \setminus (S_t \cup O_{t-1})$ at every round of a set T_{race}, and let $w(t) = t - |T_{\mathrm{race}} \cap [1, t]|$. Consider any record whose observations lie in K, and any M such that T_{race} has an Mth element τ_M. Then $K[1 \dots M]$ is resolved by a round $t_M \le \tau_M$, and

$$|O_T \cap K[1 \dots M]| \;\ge\; \tfrac{1}{2}\big(M - 1 - w(t_M)\big) \qquad \text{for all } T \ge t_M.$$

If $w(t) = o(t)$, then $\underline{d}(O, K) \ge 1/2$.

(R2) For every generator G, the racing record run against G is a complete repetition-free enumeration $E_G \in \mathcal{E}(K)$, and $|O(G, E_G) \cap K[1 \dots M]| \le M/2 + \log_2 M + 3$ for every $M \ge 2$.

Proof. **(R1).** Let t_M be the first round by which $K[1 \dots M]$ is resolved. At a race round $t < t_M$ the prefix is still unresolved after the observation w_t, so the least element of $K \setminus (S_t \cup O_{t-1})$ lies in the prefix, and the generator claims a new prefix point. Hence M race rounds suffice to resolve the prefix, which gives $t_M \le \tau_M$, and

$$|O_{t_M} \cap K[1 \dots M]| \;\ge\; (t_M - 1) - w(t_M - 1) \;\ge\; t_M - 1 - w(t_M).$$

Every prefix point that has not been output by round t_M has been observed by then, and there are at most t_M observations, so $M - |O_{t_M} \cap K[1 \dots M]| \le t_M$. Adding the two bounds gives $2|O_{t_M} \cap K[1 \dots M]| \ge M - 1 - w(t_M)$, and output sets only grow. If $w(t) = o(t)$, then $\tau_M = M + w(\tau_M)$ gives $\tau_M \sim M$ and $w(t_M) \le w(\tau_M) = o(M)$; dividing by M proves the density bound.

(R2): legality. Every observation is in $K \setminus S_{t-1}$, so there are no repetitions. Fresh choices exist because K is infinite and the used sets are finite. To prove completeness, suppose x is the least point of K never observed. After every predecessor in K has been observed, if x is ever output, the next recycling round observes it. If it is never output, the next fresh round observes it. There are infinitely many rounds of both kinds. Either case contradicts the choice of x.

(R2): the cap. Fix $M \ge 2$, and let E_M be the first round by which $K[1 \dots M] \subseteq S_{E_M} \cup O_{E_M}$. While that prefix remains unresolved, each fresh observation resolves one of its points. By time $4M$

there are at least

$$4M - \lfloor \log_2(4M) \rfloor - 1 \geq M$$

non-recycling rounds. Thus $E_M \leq 4M$. Let F_M be the points in the prefix first claimed by a fresh observation, and $G_M = O \cap K[1 \ldots M]$. The sets are disjoint because outputs are fresh relative to all observations. Every point of G_M first enters the output by time E_M: later fresh output of a previously resolved point would either repeat an output or belong to the seen-set, and so cannot first enter G_M.

Pair the first-output round t of a point $y \in G_M$ with the observation at the *same* round whenever that observation is fresh. Just before the observation, y is unused; hence the adversary's least unused point is no larger than y and lies in the prefix. It belongs to F_M. This is injective, since the observation stream has no repetitions. Only scheduled recycling rounds can be exceptions, at most $\lfloor \log_2 E_M \rfloor + 1$. Therefore

$$|G_M| \leq |F_M| + \log_2 E_M + 1, \qquad |F_M| + |G_M| \leq M.$$

It follows that $|G_M| \leq M/2 + (\log_2 M + 3)/2 \leq M/2 + \log_2 M + 3$, proving the displayed cap. Here E_M is the resolution time, not the time at which every output-hoarded point has been recycled into the stream; the latter can be exponentially larger. The cap bounds the limit set O, so it bounds O_T for every finite T, which is the form used in Theorem 4.1(a). □

C.2 Proof of the exchange theorem

Proof of Theorem 4.1(a). Let E_G be the racing record of Lemma C.1(R2) run against G. It lies in $\mathcal{E}(K)$, so by Lemma A.1 the instances (K, K, E_G) and (K', K, E_G) share one output sequence. Split the prefix by (8). For the recorded part, the cap bounds the limit set and hence every O_T:

$$|O_T \cap K[1 \ldots m_N]| \leq |O \cap K[1 \ldots m_N]| \leq m_N/2 + \log_2 m_N + 3 \qquad (m_N \geq 2).$$

For the unrecorded part, $O_T \cap D \subseteq O_T \setminus K$, and the latter set is the set of distinct trivia by round T on the reading under K. Hence $|O_T \cap D[1 \ldots n_N]| \leq \min\{b_K(T), n_N\}$. Adding the two parts proves the first display. If $G \in \mathcal{G}_g$, then on the legal instance (K, K, E_G) we have $b_K(T) \leq B_K(T) \leq g(T)$. Only the two instances with record E_G were used, and both are legal in every model containing K and K'. □

Proof of Theorem 4.1(b). Let $s(t)$ be the number of sweeps made in rounds $1, \ldots, t$, with $s(0) = 0$. In round t, after observing w_t, the generator G_g outputs

$$a_t = \begin{cases} \min\big(K' \setminus (S_t \cup O_{t-1})\big), & S_t \not\subseteq K, \\ \min\big(D \setminus O_{t-1}\big) \ \text{(a sweep)}, & S_t \subseteq K \text{ and } s(t-1) < g(t), \\ \min\big(K \setminus (S_t \cup O_{t-1})\big) \ \text{(a race)}, & \text{otherwise.} \end{cases}$$

Legality. All three sets are infinite minus finite sets, so the minima exist. In the second case $S_t \subseteq K$ is disjoint from D. Every output therefore avoids S_t and all earlier outputs.

Trivia. At truth K every legal core lies in K, so the first case never occurs. Races output points of K and sweeps output points of D, so the trivial rounds are exactly the sweeps. A sweep occurs only when $s(t-1) < g(t)$, so $s(t) \leq g(t)$ for every t by induction, and $B_K(T) = s(T) \leq g(T)$ on every legal instance at truth K. At truth K' every output lies in K', so there are no trivia. Hence $G_g \in \mathcal{G}_g$.

The unrecorded part. Fix $E \in \mathcal{E}(K)$. The record never leaves K. We claim that $s(t) = g(t)$ for all t. This holds at $t = 0$. If $s(t-1) = g(t-1)$ and $g(t) = g(t-1) + 1$, then round t sweeps and $s(t) = g(t)$. If instead $g(t) = g(t-1)$, it does not sweep and again $s(t) = g(t)$. Points of D are never observed and only sweeps output them, so the jth sweep outputs the jth point of D. Therefore $O_T \cap D = D[1 \ldots g(T)]$, and $|O_T \cap D[1 \ldots n_N]| = \min\{g(T), n_N\}$.

The recorded part. The race rounds are the rounds without a sweep, so $w(t) = s(t) = g(t)$, and exactly $t - g(t)$ of the first t rounds are race rounds. Since g has unit steps, $t - g(t)$ is nondecreasing with unit steps, so the Mth race round is $\tau_M = \tau_g(M) = \min\{t : t - g(t) \geq M\}$ whenever this set is nonempty. By Lemma C.1(R1) with $M = m_N \geq 1$ we have $t_{m_N} \leq \tau_g(m_N)$, and for every $T \geq \tau_g(m_N)$,

$$|O_T \cap K[1 \ldots m_N]| \ \geq \ \tfrac{1}{2}\big(m_N - 1 - g(\tau_g(m_N))\big),$$

since w is nondecreasing. Adding the two disjoint parts proves the bound.

The two regimes of the error. From $\tau_g(M) = M + g(\tau_g(M))$: if $g(t) = o(t)$, then $g(t) \le t/2$ for large t, so $\tau_g(M) \le 2M$ for large M, and $g(\tau_g(M)) = o(\tau_g(M)) = o(M)$. If $g(t) \le \sqrt{t}$, then $\tau_M \le M + \sqrt{\tau_M}$; if $\tau_M \ge 4$ this gives $\sqrt{\tau_M} \le \tau_M/2$ and $\tau_M \le 2M$, and otherwise $\tau_M \le 3$. In both cases $\tau_g(M) \le 2M + 2$ and $g(\tau_g(M)) \le \sqrt{2M+2}$. □

C.3 Proof of the two-value corollary

Proof of Corollary 4.2. Upper bounds. Let G be any generator and E_G its racing record. Letting $T \to \infty$ in Theorem 4.1(a) gives

$$|O(G, E_G) \cap K'[1 \dots N]| \le m_N/2 + \log_2 m_N + 3 + \min\{|O \setminus K|,\, n_N\}.$$

If G obeys a regime all of whose members are bounded, its profile on the legal instance (K, K, E_G) is bounded, so $|O \setminus K| < \infty$. Dividing by N gives $\overline{d}(O(G, E_G), K') \le \alpha/2$, and hence $\inf_E \underline{d}(O(G, E), K') \le \alpha/2$. Without any constraint, $\min\{\cdot, n_N\} \le n_N$ gives the bound $1 - \alpha/2$.

The bounded case is attained. Every regime contains the zero function, because it is nonempty and closed under passing to smaller functions. The generator G_0, with budget $g \equiv 0$, never sweeps. At truth K it outputs only points of K, and at truth K' only points of K', so its profile is zero on every legal instance. By Theorem 4.1(b), for every $E \in \mathcal{E}(K)$ and $T \ge 2m_N + 2$ we have $|O_T \cap K'[1 \dots N]| \ge m_N/2 - (\sqrt{2m_N + 2} + 1)/2$, so $\underline{d}(O, K') \ge \alpha/2$.

The unbounded case is attained. Let $f \in \mathcal{R}$ be unbounded. Put $g(0) = 0$ and $g(t) = \min\{f(t), \lfloor\sqrt{t}\rfloor, g(t-1) + 1\}$. Each of the three terms is at least $g(t-1)$, so g is nondecreasing, and it is 1-Lipschitz with $g \le f$ and $g(t) \le \sqrt{t}$. It is also unbounded: if g were eventually constant at some value c, then for large t both $f(t)$ and $\lfloor\sqrt{t}\rfloor$ would exceed c, forcing $g(t) = c + 1$. The profile of G_g is $s \le g \le f$ at truth K and zero at truth K'. Both lie in $\mathcal{R}$, so G_g obeys $\mathcal{R}$. Since g is unbounded, $\min\{g(T), n_N\} = n_N$ for large T, and Theorem 4.1(b) gives $|O \cap K'[1 \dots N]| \ge N - m_N/2 - (\sqrt{2m_N + 2} + 1)/2$ for every $E \in \mathcal{E}(K)$. Dividing by N gives $\underline{d}(O, K') \ge 1 - \alpha/2$. At truth K' this generator produces no trivia.

Round counts. If the regimes constrain B rather than b, the upper bounds still apply because $b \le B$, and the attaining generators never repeat an output, so for them $B = b$. □

The corollary contains the endpoints of the original dichotomy as the regimes of finitely many distinct trivia and of a prescribed unbounded budget of trivial rounds. It needs no inclusion between the classes being compared. The pair $K = \{2, 6, 10, \dots\} \subset K' = \{0, 2, 4, \dots\}$ has $\alpha = 1/2$ and optimal guarantees $1/4$ and $3/4$.

C.4 Proof of the clock

Proof of Corollary 4.3. Let $\alpha \in (0, 1)$, let g be an unbounded budget with $g(T) = o(T)$, let $m_N = \lfloor \alpha N \rfloor$, and note that $(1 - \alpha)N \le n_N \le (1 - \alpha)N + 1$ and that $m_N \ge 2$ for all large N.

Lower bound. Let $G \in \mathcal{G}_g$ and let E_G be its racing record; then $b_K(T) \le g(T)$ there. Suppose $Q_{G,E_G}(T, N) \ge N - m_N/2 - \varepsilon N$. By Theorem 4.1(a), $\min\{g(T), n_N\} \ge n_N - \varepsilon N - \log_2 m_N - 3 =: y$. Since $y < n_N$, this forces $g(T) \ge y \ge (1 - \alpha - \varepsilon)N - \log_2 N - 3$. Hence $T \ge g^{-1}\big((1 - \alpha - \varepsilon)N - \log_2 N - 3\big)$, and this holds simultaneously for every N with $m_N \ge 2$ on the single record E_G.

Upper bound. Let $E \in \mathcal{E}(K)$, put $e_N = (1 + g(\tau_g(m_N)))/2$, which is $o(N)$ by Theorem 4.1(b), and suppose that $\varepsilon N \ge e_N$, which holds for all large N. Let $T^* = \max\{\tau_g(m_N),\ g^{-1}(n_N - \varepsilon N + e_N)\}$. Then $g(T^*) \ge n_N - \varepsilon N + e_N$ and $n_N \ge n_N - \varepsilon N + e_N$, so Theorem 4.1(b) gives $Q_{G_g,E}(T^*, N) \ge m_N/2 + n_N - \varepsilon N = N - m_N/2 - \varepsilon N$. Hence $\tau_N(\varepsilon) \le T^*$ whenever $m_N \ge 1$ and $\varepsilon N \ge e_N$, for every N. Since g is sublinear, $g^{-1}(y)/y \to \infty$, whereas $\tau_g(m_N) \le 2m_N$ for large N. So for all $N \ge N_0$, with N_0 depending on g, α and ε, the maximum is the second term and $\tau_N(\varepsilon) \le g^{-1}\big((1 - \alpha - \varepsilon)N + o(N)\big)$. For small N the first term can dominate; an early burst of sweeps is an example. If $g(T) \le \sqrt{T}$, then $e_N \le (\sqrt{2N + 2} + 1)/2$ and $\tau_g(m_N) \le 2N + 2$. Moreover $n_N - \varepsilon N + e_N \le (1 - \alpha - \varepsilon)N + 1 + (\sqrt{2N + 2} + 1)/2 \le (1 - \alpha - \varepsilon)N + \sqrt{2N + 2}$ for $N \ge 4$, and $\max\{u, v\} \le u + v$, which gives the stated form.

The square-root budget. For $g(T) = \lfloor \sqrt{T} \rfloor$ we have $g^{-1}(y) = \lceil y \rceil^2$ for $y > 0$. For $N \geq 4$ with $m_N \geq 2$ and $\varepsilon N \geq (1 + \sqrt{2N+2})/2$, the minimax recovery time $\inf_{G \in \mathcal{G}_g} \sup_{E \in \mathcal{E}(K)} \tau_N(\varepsilon; G, E)$ therefore lies between $\left[(1-\alpha-\varepsilon)N - \log_2 N - 3\right]_+^2$ and $\left((1-\alpha-\varepsilon)N + \sqrt{2N+2} + 1\right)^2 + 2N + 2$. Both equal $(1-\alpha-\varepsilon)^2 N^2 \big(1 + O(N^{-1/2})\big)$ for fixed α and ε. For small N neither side need hold: the conditions of the upper bound can fail, and the recovery threshold may then never be reached on some records. □

For every budget the converse and the attainment of Theorem 4.1 differ by at most $\log_2 m_N + 3 + (1 + g(\tau_g(m_N)))/2$. This is $o(N)$ for every sublinear budget; the gap is caused by sweeps that interrupt the race while a prefix of K is still being resolved. A linear budget is different in kind: a generator that sweeps a constant fraction of the time also races only a constant fraction of the time, and the recorded term of the frontier is no longer attained.

C.5 Boundary checks and the role of each assumption

The shared record is essential. At $\alpha = 1/2$, let the truth be K' and let the observed core be $D = \{0, 4, 8, \ldots\}$, which also has density $1/2$ in K'. Observing 0 refutes K, the ledger closes, and a racer on K' faces a legal enumeration of D. Against the increasing enumeration of D, racing on K' outputs exactly K and has coverage $1/2$, not $3/4$. The frontier of Theorem 4.1 is a property of an ambiguous record, not of the density of the core.

Safety must be imposed under the smaller reading. If trivia were constrained only at truth K', the sweep would satisfy every regime vacuously and would cover $1 - \alpha/2$ with bounded trivia. The reading at truth K is a legal instance of the family, and it is where the price is paid.

Repetition cannot buy coverage. The caps bound the limit set. A generator that repeats finitely many points of D infinitely often has finitely many distinct trivia, and by Theorem 4.1(a) its coverage is at most $\alpha/2$.

The ambient is fixed. With $F = U$ every output is valid and the verifier is silent. By Theorem 3.1(i), no verifier can change any statement of this appendix inside a fixed formal world. Across formal worlds, Theorem 3.1(ii) applies to the criteria judged in the limit, which include lower-density coverage but not a budget imposed at every round.

Remark C.2 (The half is the price of one observation per output). Suppose each round reveals $r \geq 1$ fresh record points before the generator outputs one point. In (R1), each race round before resolution still claims a new prefix point, while each round resolves at most r prefix points by observation. This gives $|O_{t_M} \cap K[1 \ldots M]| \geq (M - r(1 + w(t_M)))/(1 + r)$. In (R2), let the racing record reveal the r least unused points of K in each round and recycle at rounds 2^j. If the generator outputs a prefix point y for the first time in a round, then y was not among the r least unused points, so all r fresh observations of that round lie below y in the prefix. The pairing is therefore one output to r observations, which gives a cap of $M/(1 + r) + O(\log M)$. The recorded term of the frontier becomes $m_N/(1 + r)$ and the endpoints become $\alpha/(1 + r)$ and $1 - \alpha r/(1 + r)$. The unrecorded term $\min\{g(T), n_N\}$ and the jump $1 - \alpha$ are unchanged, because the ledger does not depend on how fast the record grows. We sketch this remark and use it nowhere else.

C.6 Persistent exploration harvests in every family

The exact frontier needs a shared record. The harvest itself does not. A still-missing point has only finitely many predecessors in the universe, so infinitely many sweeps reach it however far apart they are.

Lemma C.3 (Pointer lemma)**.** Suppose R is an infinite set of sweep rounds and each sweep outputs the least point of $V_t \setminus (S_t \cup O_{t-1})$. Every point x that belongs to every V_t and is never observed belongs to the limit output set.

Proof. If x were never output, it would be available at every sweep, and each sweep would output a new point strictly below x. There are only finitely many such points in the universe, which contradicts the infinitude of R. The argument does not require the outputs to be valuable. □

Proposition C.4 (Universal sweep)**.** Let all candidate targets $H_1, H_2, \ldots$ lie in a common known infinite ambient F. There is a sound generator G_{sw} with $O \supseteq H \setminus C$ on every legal instance, and hence $c \geq 1 - \bar{d}(C, H)$. Its trivial rounds satisfy $B(T) \leq T_0 + \lfloor \sqrt{T} \rfloor$ for an instance-dependent finite T_0.

Proof. For $0 \leq d \leq t$ define $J_{t,d} = \bigcap\{H_i : i \leq d,\ S_t \subseteq H_i\}$ with $J_{t,0} = F$, where every empty intersection is F. Let d_t be the largest $d \leq t$ for which $J_{t,d}$ is infinite, and put $J_t = J_{t,d_t}$; testing infinitude is permitted under our non-effective conventions. Let $U_t = \bigcup\{H_i : S_t \subseteq H_i\}$ and $R = \{k^2 : k \geq 1\}$, and let G_{sw} output the least point of $U_t \setminus (S_t \cup O_{t-1})$ at rounds of R and the least point of $J_t \setminus (S_t \cup O_{t-1})$ at other rounds. Both branches remain inside F and avoid all previous outputs.

Let H_z be the true target and C the presented core. For each $i \leq z$ with $C \not\subseteq H_i$, some point of $C \setminus H_i$ is eventually observed. After all these finitely many witnesses have appeared and $t \geq z$, every surviving candidate with index at most z contains C. Hence $J_{t,z}$ is infinite, $d_t \geq z$ and $J_t \subseteq H_z$; this is the identified-intersection principle of Fact A.3. The true target is always consistent, so $H \subseteq U_t$, and Lemma C.3 gives $H \setminus C \subseteq O$, whence $|O \cap H[1\ldots N]| \geq N - |C \cap H[1\ldots N]|$. After a finite time T_0 every default output lies in H, so only the first T_0 rounds and the at most $\lfloor \sqrt{T} \rfloor$ sweeps can be trivial. □

The transient T_0 is not uniform over instances, which is why a budget prescribed from the first round needs the two-target structure of Theorem 4.1. The Lean development checks the pointer lemma, the hybrid sweep with an eventually valid proposer and its trivia-round bound (Appendix F).

D Compression constructions and counting

D.1 Mathematical background

The case study of Section 5 works inside two monoids and the macro framework of Aksenov et al. (2026) (ABFM). We collect the definitions here for the constructions in this appendix. We use this model as a substrate only: no theorem of ABFM is an ingredient in any proof below.

D.1.1 Monoids and canonical orders

A *monoid* is a set with an associative binary operation and an identity element. Two monoids carry the case study.

The free abelian monoid A_n on generators $a_1, \ldots, a_n$ is the set of formal products $\prod_{i=1}^{n} a_i^{N_i}$ with $N_i \in \mathbb{N}$, multiplied coordinate-wise; it is isomorphic to $(\mathbb{N}^n, +)$, with each element identified with its exponent vector $(N_1, \ldots, N_n)$. The *raw length* (the generator count) is $|w|_G = \sum_{i=1}^{n} N_i$. We order A_n by degree-then-lex: first by $|w|_G$, then lexicographically on the exponent vector.

The free monoid F_n on n letters is the set of finite words over the alphabet, with concatenation as the operation and the empty word as identity; $|w|_G$ is the word length. We order F_n by length-then-lex.

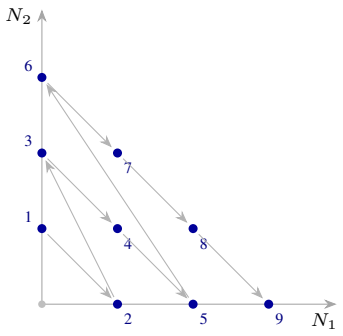
Both orders are order-isomorphic to $\mathbb{N}$ (Figure 5; each has finitely many elements below any given one, and the order is total), which is exactly the structure the density machinery of Appendix A requires: a countable ground set with a canonical order in which $L[1\ldots N]$ is a well-defined initial segment.

D.1.2 The macro framework

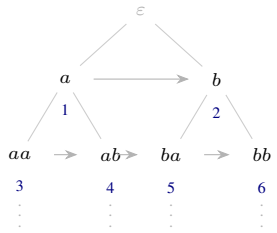
Following ABFM, a *dictionary* D is a finite or infinite set of designated monoid elements, called *macros*. Generation is allowed to use the *augmented generating set* $G \cup D$: a factorization of w is an expression of w as a product of elements of $G \cup D$. The *wrapped length*

$$|w|_{G \cup D} \coloneqq \min\{\, k : w = g_1 \cdots g_k,\ g_j \in G \cup D \,\} \tag{10}$$

is the minimum number of symbols from $G \cup D$ in any factorization of w, and the *raw* (unwrapped) length $|w|_G$ uses G alone. A macro that abbreviates a long element shortens its wrapped length below its raw length: this is the precise sense in which "compressible" means "short under the dictionary."



A_2: elements $a_1^{N_1} a_2^{N_2}$ are lattice points;
degree-then-lex snakes through the diagonals



F_2: words are tree nodes;
length-then-lex reads the tree level by level

Figure 5: The two monoids and their canonical orders (Appendix D.1.1). Left: $A_2 \cong (\mathbb{N}^2, +)$; degree-then-lex enumerates the lattice diagonal by diagonal, so the order is isomorphic to $\mathbb{N}$. Right: F_2; length-then-lex reads the word tree level by level. In both cases $L[1..N]$ is a well-defined initial segment, which is all the density machinery of Appendix A needs.

For A_n the relevant dictionaries are *geometric*: $D_c = \{ a_i^{c^j} : i \le n,\ j \ge 1 \}$ provides single-letter powers at every scale c^j. For this dictionary the wrapped length has a closed form,

$$|w|_{G \cup D_c} = \sum_{i=1}^{n} s_c(N_i), \tag{11}$$

where $s_c(N)$ is the sum of the base-c digits of N (Figure 6) (the digit-sum lemma, Lemma D.2; the minimizing factorization is the base-c representation of each coordinate).

The *compressibility cut* is the threshold

$$\theta_\kappa(L) := \kappa \log_2(1 + L), \tag{12}$$

and the *compressible sublanguage* is

$$H_D := \{ w : |w|_{G \cup D} \le \theta_\kappa(|w|_G) \}, \tag{13}$$

the elements whose wrapped length is logarithmic in their raw length. A *bounded cut* $\theta \equiv 1$ (used in Proposition D.5) collapses $H_{D,1}$ to the nameable corpus $\{\epsilon\} \cup G \cup D$ itself, and the *dictionary-only* length ℓ_s of Appendix D.3 counts dictionary factors alone; Section 5 develops this last variant.

We are explicit about the ledger: the framework supplies definitions (macros, wrapped length, geometric dictionaries) and the motivating asymmetry between A_n and F_n, and nothing more. The proofs in Appendix D are self-contained; no ABFM theorem is invoked.

D.1.3 How the case study uses the framework

The supplementary constructions use these objects in two further ways. In A_n it builds *refinement chains*: nesting the geometric dictionaries by $D_{t+1} \subseteq D_t$ gives a descending chain of compressible languages H_{D_t} separated by digit-sum witnesses (Theorem D.1). In F_n it builds *name-cut pairs*: words $x^i a\, x^{q+1-i} a\, x^{2^m}$ whose unique parsing lets an a-position cut the corpus into a sparse target K and a full world K' at any rational density (Proposition D.5).

D.2 Abelian refinement chains

A descending chain admits singleton tell-tales if each target contains one point excluded by its immediate successor. Dictionary inclusion will provide the nesting. Squaring a geometric base removes its odd scales; the witness below puts weight on two of those scales so that its description cost rises from two factors to twice the base, while the threshold grows only logarithmically.

Theorem D.1 (Abelian refinement chains are identifiable)**.** Fix $n \ge 1$, $\kappa \ge 1$ and an integer $b \ge 16\kappa^2$. Let $b_t = b^{2^t}$ and $D_t = \{a_i^{b_t^j} : 1 \le i \le n,\ j \ge 1\}$. The languages $\{H_{D_t}\}_t$ form a strictly descending

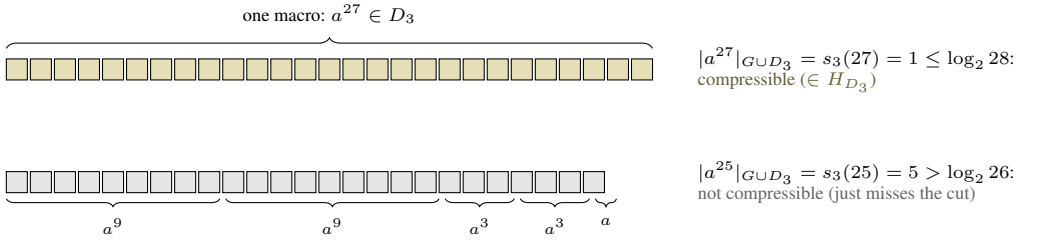


Figure 6: Wrapped length and the compressibility cut, by example (Appendix D.1.2; A_1, geometric dictionary $D_3 = \{a^3, a^9, a^{27}, \dots\}$, $\kappa = 1$). The wrapped length is the base-3 digit sum (11): a^{27} compresses to a single macro and passes the cut, while $a^{25} = a^9 a^9 a^3 a^3 a$ needs five symbols and just misses it. The cut accepts a named power while rejecting a nearby word; Proposition D.6 separately quantifies sparsity in its stated parameter range.

infinite chain with singleton tell-tales $T_t = \{a_1^{b_t + b_t^3}\}$. Consequently they admit exact breadth from full enumeration, with no eventual trivia.

Proof. Write $w = \prod_i a_i^{N_i} \in A_n \cong (\mathbb{N}^n, +)$, identified with its exponent vector $(N_1, \dots, N_n)$. We use three lemmas.

Lemma D.2 (Digit-sum formula)**.** For the geometric dictionary $D_c = \{a_i^{c^j} : i \leq n, j \geq 1\}$ (with $c = b_t$), the wrapped length is $|w|_{G \cup D_c} = \sum_{i=1}^{n} s_c(N_i)$, where $s_c(N)$ is the sum of the base-c digits of N.

Proof. Each coordinate $a_i^{N_i}$ is factored independently (the dictionary words are single-letter powers). For one coordinate: a factoring corresponds to a representation of N_i as a sum of powers of c with nonnegative integer coefficients (multiplicities), and its symbol count is the sum of those multiplicities. Replacing c copies of the macro $a_i^{c^j}$ by one copy of $a_i^{c^{j+1}}$ preserves the represented power and strictly decreases the symbol count, so every minimizing representation has all multiplicities $< c$; the unique representation with all multiplicities $< c$ is the base-c representation of N_i. Hence the minimum symbol count is its digit sum $s_c(N_i)$. Summing over coordinates gives the formula. □

Lemma D.3 (Nesting)**.** $H_{D_{t+1}} \subseteq H_{D_t}$, so $\{H_{D_t}\}_t$ is a descending chain.

Proof. Since $b_{t+1} = b_t^2$, we have $D_{t+1} \subseteq D_t$. A richer dictionary cannot increase the minimal factoring length, so $|w|_{G \cup D_t} \leq |w|_{G \cup D_{t+1}}$ for every w. Therefore every w with $|w|_{G \cup D_{t+1}} \leq \theta_\kappa(|w|_G)$ also satisfies $|w|_{G \cup D_t} \leq \theta_\kappa(|w|_G)$: membership transfers downward and the chain is descending. □

Lemma D.4 (Witness / singleton tell-tale)**.** $w_t := a_1^{b_t + b_t^3} \in H_{D_t} \setminus H_{D_{t+1}}$, and $T_t = \{w_t\}$ is a tell-tale for H_{D_t} in the chain.

Proof. By Lemma D.2, $|w_t|_{G \cup D_t} = s_{b_t}(b_t + b_t^3) = 2$ (the number $b_t + b_t^3$ has base-b_t digits 1 at positions 1 and 3). In base $b_{t+1} = b_t^2$, the same number is $b_t + b_t \cdot b_t^2 = b_t(1 + b_t^2)$, whose base-$b_t^2$ digit sum is $2b_t$ (a digit b_t at positions 0 and 1). The raw length is $|w_t|_G = b_t + b_t^3$, so $\theta_\kappa(|w_t|_G) = \kappa \log_2(1 + b_t + b_t^3) \leq 5\kappa \log_2 b_t$ for large b_t. Then $|w_t|_{G \cup D_t} = 2 \leq \theta_\kappa$, so $w_t \in H_{D_t}$; while $|w_t|_{G \cup D_{t+1}} = 2b_t > 5\kappa \log_2 b_t \geq \theta_\kappa$ by the inequality verified below, so $w_t \notin H_{D_{t+1}}$. As the chain is descending, $\{w_t\}$ separates H_{D_t} from every proper sub-language $H_{D_{t'}}$, $t' > t$: any $H_{D_{t'}} \ni w_t$ would need $t' \leq t$, so no proper sub-language contains the tell-tale. The languages are infinite (e.g. all $a_1^{c^j}$ have wrapped length 1). □

The three lemmas give a strictly descending chain of infinite languages with singleton tell-tales, so $\mathcal{H}_{\mathrm{ab}}$ satisfies Angluin's condition. By Corollary 3.2(a), applied in the single fiber $F = A_n$, exact breadth in the limit is achievable: in the abelian regime valuable generation needs no trivia. (The partial-enumeration behavior of the chain, its relative densities, is open.) □

For completeness, the witness inequality is uniform over the stated parameter range. The function $c/\log_2 c$ is increasing for $c > e$, so for $c \geq 16\kappa^2$ it suffices to check $c = 16\kappa^2$. At that point the inequality is $32\kappa > 20 + 10\log_2 \kappa$. Its difference is positive at $\kappa = 1$ and has derivative $32 - 10/(\kappa \ln 2) > 0$ on $[1, \infty)$. Also $1 + c + c^3 \leq c^5$ and $\kappa \log_2(1 + c + c^3) \geq 2$ for $c \geq 16$, proving both sides of the separation.

D.3 Free-monoid pairs

Proposition D.5 (Free-monoid pairs at every rational density)**.** Work in the free monoid on $\{a, x\}$ and fix integers $1 \leq p < q$. Under the constant name cut one, let $\mu_{m,i} = x^i a x^{q+1-i} a x^{2^m}$ for $m \geq 1$ and $1 \leq i \leq q$, and put $K = \{\epsilon, a, x\} \cup \{\mu_{m,i} : i \leq p\}$ and $K' = \{\epsilon, a, x\} \cup \{\mu_{m,i} : i \leq q\}$. Then $|K \cap K'[1 \ldots N]| = (p/q)N + O(q)$, so Theorem 4.1 and Corollaries 4.2 and 4.3 hold for this pair with $\alpha = p/q$.

The name cut. Under a constant cut of one, the zero-factor representation gives ϵ, and a one-factor representation gives a primitive letter or a dictionary macro. Each word $\mu_{m,i}$ has exactly two occurrences of a, which recover i and m uniquely. Its length is $q + 3 + 2^m$, independent of i at each scale. A prefix of K' contains complete layers and at most one incomplete layer, each with q words and p selected tags. The finite initial set $\{\epsilon, a, x\}$ contributes only a constant, giving $|K \cap K'[1 \ldots N]| = (p/q)N + O(q)$. All languages and differences are infinite; for example a^j, $j \geq 2$, lies outside the nameable corpus. Theorem 4.1 is stated for an arbitrary nested pair through the counts m_N and n_N, so it applies verbatim; $m_N = (p/q)N + O(q)$ gives $\alpha = p/q$ in the corollaries.

The density-zero pair of Proposition 5.1. The one-factor construction gives a fixed tag fraction. For zero density, we want that fraction to compound across many independently tagged blocks. Three features serve this purpose: a marker makes the blocks uniquely recoverable, equal tag lengths preserve the compression test, and exponentially growing scales provide many compressible long descriptions. The proof must also control prefixes that stop within a length layer; a count over complete layers would not establish the stated ordered density.

Fix integers $1 \leq p < q$ and a real $\kappa > 0$. Use the tag words

$$\widehat{\mu}_{m,i} = aa\, x^i a\, x^{q+1-i} a\, x^{2^m}, \qquad m \geq 1, \quad 1 \leq i \leq q,$$

and put $\widehat{D}_s = \{\widehat{\mu}_{m,i} : m \geq 1,\ i \leq s\}$ for $s \in \{p, q\}$. Here the wrapped length is *dictionary-only*:

$$\ell_s(w) = \min\{k : w = d_1 \cdots d_k,\ d_j \in \widehat{D}_s\},$$

with value $+\infty$ when no such factorization exists and $\ell_s(\varepsilon) = 0$. Define

$$\widehat{K}_s = \{w : \ell_s(w) \leq \kappa \log_2(1 + |w|_G)\}, \qquad \widehat{K} = \widehat{K}_p, \quad \widehat{K}' = \widehat{K}_q.$$

In canonical length-then-lex order,

$$\overline{d}(\widehat{K}, \widehat{K}') = 0.$$

More precisely, if the last word of the first N elements of $\widehat{K}'$ has raw length L, their $\widehat{K}$-fraction is $O(1/\log L)$ as $N \to \infty$. The implicit constants may depend on p, q, κ.

Proof of Proposition 5.1. An occurrence of aa marks exactly one block start: every other a is separated from the next a by a nonempty run of x's, including at a concatenation boundary. Thus a word in $\widehat{D}_q^*$ has a unique block decomposition, recovering every tag i and scale m. A block has raw length $h + 2^m$, where $h = q + 5$, independent of its tag. A word with k blocks and total raw length u belongs to $\widehat{K}_q$ exactly when $k \leq \kappa \log_2(1 + u)$; it belongs to $\widehat{K}_p$ exactly when this inequality holds and every tag is at most p.

Write $C(L) = |\{w \in \widehat{K}_q : |w|_G \le L\}|$, and let $C_k(L)$ count those words with exactly k blocks. Let $A_k(L)$ be the corresponding count in $\widehat{K}_p$. For $L \ge 2$, set

$$J(L) = \lfloor \log_2 L \rfloor, \qquad K(L) = \lfloor \kappa \log_2(1+L) \rfloor.$$

Every block scale is at most $J(L)$ and every compressible word has at most $K(L)$ blocks, so

$$C_k(L) \le (qJ(L))^k, \qquad A_k(L) = (p/q)^k C_k(L). \tag{14}$$

The second equality holds because raw length and compressibility depend only on the scales, while all q^k tag assignments produce distinct words.

We first establish the count needed to control small block numbers:

$$C(L) = \exp\big(\Theta(\log L \, \log\log L)\big). \tag{15}$$

The upper bound follows from $C(L) \le (K(L)+1)(qJ(L))^{K(L)}$ for all sufficiently large L. For the lower bound, put

$$M = \lfloor \tfrac{1}{2} \log_2 L \rfloor, \qquad k = \lfloor \kappa M \rfloor, \qquad R = M - \lceil \log_2 k \rceil.$$

For large L, both k and R are positive and of order $\log L$. Fix the first scale to be M and choose each of the other $k-1$ scales independently from $\{1, \dots, R\}$; choose all tags freely. The resulting raw length u satisfies

$$2^M \le u \le hk + 2^M + (k-1)2^R \le hk + 2^{M+1} \le L.$$

Also $k \le \kappa M < \kappa \log_2(1+u)$, so all $q^k R^{k-1}$ words are compressible. Their number proves the lower bound in (15).

Next we compare consecutive length balls. Choose a fixed integer $k_0 \ge 2$ such that

$$\kappa \log_2\left(\frac{k}{k-1}\right) \le 1 \qquad \text{for every } k \ge k_0.$$

For example, $k_0 = \lceil \kappa / \log 2 \rceil + 2$ suffices, where $\log$ is natural. For a compressible word with $k \ge k_0$ blocks and raw length u, delete its first shortest block, whose length d is at most u/k. The remainder is still compressible, because

$$\kappa \log_2(1+u-d) \ge \kappa \log_2(1+u) - \kappa \log_2\left(\frac{k}{k-1}\right) \ge k-1.$$

It has raw length at most $L-1$ when the original word has raw length at most L. Each image has at most $K(L)+1$ possible insertion positions, q deleted tags, and $J(L)$ deleted scales. Hence

$$C(L) \le \sum_{k<k_0} (qJ(L))^k + qJ(L)(K(L)+1)C(L-1).$$

The finite sum is bounded by a fixed power of $\log L$, whereas (15) applied at $L-1$ grows faster than every such power. Consequently,

$$\frac{C(L)}{C(L-1)} = O((\log L)^2). \tag{16}$$

Finally set

$$B = \left\lceil \frac{3 \log\log L}{\log(q/p)} \right\rceil.$$

By (14), the contribution with $k \le B$ is at most

$$(B+1)(qJ(L))^B = \exp\big(O((\log\log L)^2)\big),$$

and the contribution with $k > B$ is at most $(\log L)^{-3} C(L)$. If a shortlex prefix ends at raw length L, it contains every word of raw length less than L, so $N \ge C(L-1)$. Its good-word fraction is therefore at most

$$\frac{\exp(O((\log\log L)^2))}{C(L-1)} + (\log L)^{-3} \frac{C(L)}{C(L-1)} = O(1/\log L),$$

by (15) and (16). This bound holds for every prefix, proving the claimed zero upper density. □

Consequences for coverage. Single blocks of arbitrarily large scale pass the cut, both with retained tags and with a tag greater than p, so $\widehat{K}$, $\widehat{K}'$ and $\widehat{K}' \setminus \widehat{K}$ are infinite. A record enumerating $\widehat{K}$ is legal under either truth. Theorem 4.1 applies verbatim with $m_N = |\widehat{K} \cap \widehat{K}'[1 \dots N]| = o(N)$. By Corollary 4.2, the optimal coverage of $\widehat{K}'$ on such records is 0 in every bounded trivia regime and 1 in every regime that permits unbounded trivia. The race-and-sweep generator needs only the two dictionaries, which the candidate model supplies, and the ambient can be the whole free monoid. Density-one coverage of $\widehat{K}'$ may still omit points of $\widehat{K}$.

D.4 Imprecision of exhaustive generation

Proposition D.6 (Imprecision of exhaustive generation)**.** In A_n let $n \geq 1$, $c \geq 4$ be an integer, $\kappa \geq 0$ and $\delta = n - \kappa - n/\log_2 c > 0$. For the geometric dictionary D_c and the raw-length ball $B_G(r)$ at integer radii $r \geq 1$,

$$\frac{|H_{D_c} \cap B_G(r)|}{|B_G(r)|} \leq 2^{n+\kappa} n^n r^{-\delta}. \tag{17}$$

If additionally $\kappa \log_2 c \geq 1/2$, the ratio is at least $2^{-n-5} r^{1/\log_2 c - n}$ for every $r \geq 1$.

The statement below restates the two bounds with their constants; the counting admits $\kappa = 0$, in which case H contains only the identity, whereas a generation instance uses $\kappa > 0$.

Let $n \geq 1$ and $c \geq 4$ be integers, $\kappa \geq 0$, and $H = H_{D_c} \subseteq A_n$. Set $d = 1/\log_2 c$. For every integer $r \geq 1$,

$$\frac{|H \cap B_G(r)|}{|B_G(r)|} \leq 2^{n+\kappa} n^n r^{\kappa + nd - n}. \tag{18}$$

In particular, when $\kappa < n(1-d)$ the upper exponent $\delta = n - \kappa - nd > 0$ is valid, without a logarithmic loss. If, additionally, $\kappa \log_2 c \geq \frac{1}{2}$, then the lower exponent is also valid:

$$2^{-n-5} r^{d-n} \leq \frac{|H \cap B_G(r)|}{|B_G(r)|}. \tag{19}$$

The extra condition is sufficient for this elementary lower-bound construction; it is not asserted to be necessary.

Proof. Identify $w \in A_n$ with $(N_1, \dots, N_n) \in \mathbb{N}^n$. The digit-sum formula gives

$$|w|_{G \cup D_c} = \sum_i s_c(N_i), \qquad |w|_G = \sum_i N_i.$$

For the upper bound, let

$$\ell = \lfloor \log_c r \rfloor, \quad M = n(\ell + 1), \quad D = \lfloor \kappa \log_2(1 + r) \rfloor.$$

Every coordinate in $B_G(r)$ has at most $\ell + 1$ base-c digits. The padded digit matrix determines the vector uniquely, and membership in $H \cap B_G(r)$ implies that its sum is at most D. Dropping the digit upper bounds and the raw-length constraint only enlarges the count. Stars and bars therefore gives

$$|H \cap B_G(r)| \leq \binom{M + D}{M} \leq 2^{M+D} \leq 2^{n+\kappa} r^{nd+\kappa}.$$

Here $2^M \leq 2^n r^{nd}$ and $2^D \leq (1+r)^\kappa \leq 2^\kappa r^\kappa$. The ball contains the cube $\{0, \dots, \lfloor r/n \rfloor\}^n$, so $|B_G(r)| \geq (\lfloor r/n \rfloor + 1)^n \geq (r/n)^n$. Dividing proves (18).

For the lower bound assume $\kappa \log_2 c \geq \frac{1}{2}$ and first let $m = \lfloor \log_c r \rfloor \geq 4$. Consider vectors with only their first coordinate nonzero, of the form

$$N = c^{m-1} + \sum_{j=0}^{m-4} \varepsilon_j c^j, \qquad \varepsilon_j \in \{0, 1\}, \qquad \sum_{j=0}^{m-4} \varepsilon_j \leq \left\lfloor \frac{m-3}{2} \right\rfloor.$$

The two digit positions immediately below the leading digit are fixed to zero. Thus $c^{m-1} \leq N < c^m \leq r$, and

$$s_c(N) = 1 + \sum_j \varepsilon_j \leq \frac{m-1}{2} \leq \kappa(m-1)\log_2 c \leq \kappa \log_2(1+N).$$

Each point passes the cut at its *own* raw length, so all these points belong to $H \cap B_G(r)$. Complementing all $m-3$ bits shows that at least half of their 2^{m-3} assignments have weight at most $\lfloor (m-3)/2 \rfloor$. Since $r < c^{m+1}$ and $c^d = 2$, this yields

$$|H \cap B_G(r)| \geq 2^{m-4} \geq \frac{r^d}{32}.$$

For $1 \leq r < c^4$, the same weaker bound follows from the identity element: $|H \cap B_G(r)| \geq 1$ and $r^d < 16$. Finally $|B_G(r)| \leq (r+1)^n \leq 2^n r^n$, proving (19) for every $r \geq 1$. □

E Numerical illustration

The computations instantiate the running family (9) at rational $\alpha = p/q$ and run the race-and-sweep generator G_g of Theorem 4.1(b) against its own racing record from Lemma C.1(R2). The same transcript is read under K to count trivia and under K' to count coverage, and no target label is ever given to the generator. The program `code/exchange.py` uses only integer rank arithmetic. At every round and every tracked prefix it asserts four things: freshness; that the trivia under K equal the sweeps and are at most $g(t)$; the converse inequality $Q \leq \Phi_g + \log_2 m_N + 3$ on the racing record; and the attainment inequality of Theorem 4.1(b), in its general form and, for budgets $g(T) \leq \sqrt{T}$, in the form $Q \geq \Phi_g - (\sqrt{2m_N+2}+1)/2$ once $t \geq 2m_N + 2$. A control run checks the attainment inequality on a second, oblivious record that lists K in increasing order. These checks can detect implementation errors. They do not replace the proofs or certify their universal quantifiers, and nothing here is an experiment on a language model or a proof assistant.

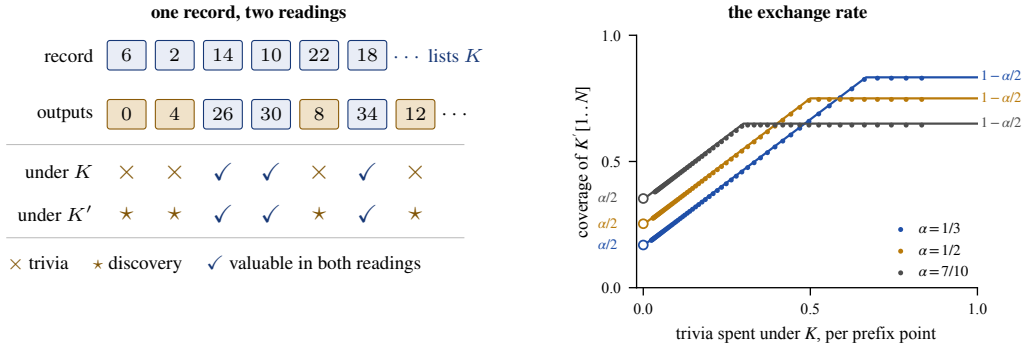


Figure 7: *Left:* with $K' = \{0, 2, 4, \ldots\}$ and $K = \{2, 6, 10, \ldots\}$, the record lists K. The same outputs are read twice, and no observation or verifier answer separates the two rows. *Right:* the frontier of Theorem 4.1 on the prefix $N = 1200$. Lines show $\big(m_N/2 + \min\{b, n_N\}\big)/N$ against the trivia b spent under K. Markers are computed runs of the race-and-sweep generator on its own adversarial record at horizons $T \geq 2m_N + 2$ (Appendix E), and open circles are pure racing. On its own racing record, no generator lies above the line by more than $(\log_2 m_N + 3)/N$.

The frontier. The right panel of Figure 7 shows the fixed prefix $N = 1200$ at $\alpha \in \{1/3, 1/2, 7/10\}$ under $g(T) = \lfloor \sqrt{T} \rfloor$, at logarithmically spaced horizons from $T = 2m_N + 2$ to $T = 10^6$. At every horizon the measured coverage lies within the two error terms of the line $\big(m_N/2 + \min\{b, n_N\}\big)/N$. The largest excess over the converse line and the smallest slack under the attainment line are recorded in `results/exchange.json`. Pure racing ($g = 0$) ends at 0.169, 0.253 and 0.353, against $\alpha/2 = 0.167$, 0.25 and 0.35. The swept runs end at 0.830, 0.746 and 0.644 after 1000 trivia under K, against $1 - \alpha/2 = 0.833$, 0.75 and 0.65. The residual is the $O(\sqrt{N})$ race loss on a finite prefix, not a failure to sweep: every point of $D \cap K'[1 \ldots N]$ has been output.

The clock. Table 2 and Figure 3 record the recovery time $\tau_N(\varepsilon)$ at $\alpha = 1/3$ under $g(T) = \lfloor \sqrt{T} \rfloor$ for N up to 3200, and under $g(T) = \lfloor T^{1/3} \rfloor$ for N up to 200. Every measured value lies between the converse and attainment bounds of Corollary 4.3, in the exact forms $g^{-1}(n_N - \varepsilon N - \log_2 m_N - 3)$ and $\max\{2m_N + 2, g^{-1}(n_N - \varepsilon N + (\sqrt{2m_N+2}+1)/2)\}$; the program asserts both. Under the square-root budget the ratio $\tau_N(\varepsilon)/((1-\alpha-\varepsilon)N)^2$ falls from 1.044 at $N = 100$ to 1.009 at

$N = 3200$ for $\varepsilon = 0.05$. This is the leading constant of Corollary 4.3. Under the cube-root budget the measured times equal $g^{-1}((1-\alpha-\varepsilon)N) = \lceil (1-\alpha-\varepsilon)N \rceil^3$ for $N \geq 100$. This is the same clock read through a different budget: slowing the budget from $T^{1/2}$ to $T^{1/3}$ changes the time to recover $N = 200$ at $\varepsilon = 0.05$ from 15 625 rounds to 1 906 624.

Budgets above the square root. The attainment bound of Theorem 4.1(b) holds for every budget in the form $Q \geq \Phi_g - (1 + g(\tau_g(m_N)))/2$ for $T \geq \tau_g(m_N)$. The program asserts this general form at every round for every budget above, and additionally for $g(T) = \lfloor T^{3/4} \rfloor$, which grows faster than $\sqrt{T}$. It checks this at $\alpha \in \{1/3, 1/2\}$ on both the racing and the increasing record, for prefixes up to $N = 1600$. For example, $\tau_g(m_N) = 663$ at $\alpha = 1/3$ and $N = 1600$, compared with $2m_N + 2 = 1068$. The smallest slack under the bound is positive in every run (`results/exchange.json`, field `general_budget`). The recovery time at $\alpha = 1/3$ and $\varepsilon = 0.05$ is 4107 rounds at $N = 800$ and 10281 at $N = 1600$, which is 1.054 and 1.047 times $((1-\alpha-\varepsilon)N)^{4/3}$. At $\varepsilon = 0.1$ and $N = 1600$ the ratio is 1.051. The slower convergence than under the square-root budget reflects the larger error term $g(\tau_g(m_N))/2$.

Table 2: Recovery times of the race-and-sweep generator on its racing record at $\alpha = 1/3$ under $g(T) = \lfloor \sqrt{T} \rfloor$. The converse bound is $g^{-1}(n_N - \varepsilon N - \log_2 m_N - 3)$, which every generator in $\mathcal{G}_g$ must meet on its own racing record. The attainment bound is $\max\{2m_N + 2, g^{-1}(n_N - \varepsilon N + (\sqrt{2m_N + 2} + 1)/2)\}$. The last column normalizes by the leading term of Corollary 4.3.

N	m_N	ε	converse bound	$\tau_N(\varepsilon)$	attainment bound	$\tau_N(\varepsilon)/((1-\alpha-\varepsilon)N)^2$
100	33	0.05	2 916	3 969	4 489	1.0437
100	33	0.10	2 401	3 364	3 844	1.0476
200	66	0.05	13 225	15 625	17 161	1.0272
200	66	0.10	11 025	13 225	14 641	1.0296
400	133	0.05	56 169	62 001	65 536	1.0190
400	133	0.10	47 089	52 441	55 696	1.0207
800	266	0.05	233 289	247 009	257 049	1.0149
800	266	0.10	196 249	208 849	218 089	1.0162
1600	533	0.05	950 625	984 064	1 008 016	1.0108
1600	533	0.10	801 025	831 744	853 776	1.0118
3200	1066	0.05	3 845 521	3 928 324	3 992 004	1.0088
3200	1066	0.10	3 243 601	3 319 684	3 378 244	1.0096

On the oblivious control record the attainment inequality holds with slack; there racing gains more than half of the core, because the record falls behind the race. The measured curves describe one generator on one adversarial record each. They are not an empirical optimum over generators; the optimum is the content of Theorem 4.1.

F LEAN FORMALIZATION

The accompanying Lean artifact reuses code from `generation-in-the-limit-lib`, developed and maintained by Shuangping Li and Peng Zhang (Li & Zhang, 2026). The supplement includes the relevant shared infrastructure, the existing `FloodHarvest` development, and its bridge to the library's patient-machine construction. These 70 source modules are preserved unchanged, with their Apache-2.0 license and snapshot provenance recorded in `UPSTREAM.md`.

To make the code follow the present paper, we provide topic-based modules under `FloodHarvest/` with descriptive declaration names, such as `finite_trivia_coverage_upper_bound`. Six of them are alias modules: each of their 51 declarations is a direct alias of an existing declaration, with the same assumptions and conclusion, and `PAPER_MAP.md` records all correspondences. The seventh, `FloodHarvest/ExchangeRate.lean`, is new. It contains proofs written for this revision of Theorem 4.1 in full and of the attainment half of Corollary 4.2, built on the unchanged library modules.

Table 3 summarizes the checked components and the steps that remain paper proofs. The new module formalizes Theorem 4.1 for every nested pair given by increasing enumerations of K and D, and

for the running family (9). It covers the converse on the racing record at every round and prefix simultaneously. It defines the race-and-sweep generator G_g as an element generator for every unit-step budget, and checks its freshness, its exactly $g(T)$ trivial rounds at truth K, and the absence of trivia at truth K'. It also checks the sweep invariant, the finite-time race bound, the attainment bound, and the prefix split $K'[1 \dots N] = K[1 \dots m_N] \sqcup D[1 \dots n_N]$. For Corollary 4.2 it checks both attainments: coverage $\alpha/2$ with no trivia, and coverage $1 - \alpha/2$ for every unbounded nondecreasing budget, however slow. The matching caps were already in the library. Lean counts rounds from 0; the statements are otherwise those of the paper. The regime-level assembly of Corollary 4.2, the clock, gluing for arbitrary criteria, and the compression counting remain paper proofs. On the verification side the library checks stripping, ambient stabilization, the exact-breadth fiber characterization and the oracle-free separation.

Table 3: Paper-to-Lean correspondence. Representative declarations are in the namespace `FloodHarvest`; module names refer to files under `FloodHarvest/`. The supplementary `PAPER_MAP.md` gives the full declaration map, including the original library names.

Result	Representative declarations (module)	Directly checked	Remaining paper steps
Lemma A.1	none	—	paper proof (Appendix A)
Theorem 3.1	`oracle_stripping_same_world`, `first_consistent_world_stabilizes`, `approximate_breadth_oracle_stripping` (`Verification`)	stripping on a fiber; stabilization of the ambient guess	gluing for an arbitrary criterion judged in the limit (Appendix B.1)
Corollary 3.2	`verifier_exact_breadth_iff_fiber_tell_tales`, `known_world_exhaustive_sound_coverage`, `no_oracle_sound_coverage_counterexample` (`Verification`)	(a) for exact breadth; (b) with the whole ambient as a parameter; the separation Theorem B.2	approximate breadth through Fact A.2; finite-query eventual sound coverage by gluing; Proposition B.5
Lemma C.1	`race_half_coverage`, `racing_record_is_legal`, `racing_prefix_cap` (`CoverageBounds`); `Exchange.race_greedy_cover_finite` (`ExchangeRate`)	greedy half-coverage; legality of the racing record; the prefix cap; the finite form of (R1)	the code uses zero-based recycling rounds
Theorem 4.1	`Exchange.converse`, `Exchange.converse_budget`, `Exchange.attainment`, `Exchange.tight_converse`, `Exchange.tight_attainment` (`ExchangeRate`)	both parts, for general nested pairs and for the running family: the frontier at every round and prefix, the generator G_g with its trivia accounting at both readings, the sweep invariant and the prefix split	none
Corollary 4.2	`Exchange.tight_bounded_attained`, `Exchange.tight_unbounded_attained` (`ExchangeRate`); `finite_trivia_coverage_upper_bound`, `unrestricted_coverage_upper_bound`, `coverage_gain_requires_infinite_trivia` (`TriviaDichotomy`)	both caps and both attainments: $\alpha/2$ with no trivia, and $1-\alpha/2$ for every unbounded nondecreasing budget; necessity of infinitely many trivia	the assembly over regimes (Appendix C.3)
Corollary 4.3	none	—	paper proof (Appendix C.4); the computations of Appendix E check the inequalities on their runs
Proposition C.4	`pointer_progress`, `patient_hybrid_sweep_coverage_and_trivia`, `hybrid_sweep_trivia_round_bound` (`SparseExploration`)	the pointer lemma; the hybrid sweep with an eventually valid proposer and its trivia-round bound	the identified-intersection default and its validity proof (Appendix C.6)
Proposition 5.1, Proposition D.5	none	—	counting proofs (Appendix D.3)

Result	Representative declarations (module)	Directly checked	Remaining paper steps
Theorem D.1	`abelian_chain_exact_breadth`, `representation_cost_eq_digit_sum`, `abelian_chain_strict_descent` (`AbelianCompression`)	digit sums, nesting, singleton witnesses and the identifiable chain for integer bases in the stated range	none beyond Corollary 3.2(a)
Proposition D.6	`raw_ball_precision_upper_bound`, `compressible_ball_count_bound` (`CompressionPrecision`)	the raw-ball upper bound with explicit constants	the lower bound under $\kappa \log_2 c \geq 1/2$ (Appendix D.4)

The project pins Lean 4.24.0, Mathlib 4.24.0 and all nine dependency revisions. From the supplementary `FloodHarvestLean/` directory, run

```
lake build FloodHarvest
lake env lean PaperAudit.lean
```

The build checks the public entry point and its imported proof modules. The audit checks that the 51 alias declarations have the same types as their source declarations, and collects the transitive axiom dependencies of these and of the 16 main declarations of `ExchangeRate`. They use only subsets of `propext`, `Classical.choice` and `Quot.sound`, with no `sorryAx` or project-specific axioms; classical noncomputable choice is consistent with the non-effective model. The source manifest, dependency lock file, statement map and verification logs are included in the supplement. Results whose declaration entry is "none" remain paper proofs, and the successful checks apply to the encoded statements with the scope specified above.

G Extended context and related work

G.1 Theorem-level comparison

Table 4 records, for each work the results build on, what the paper attributes to it and what the present results add. The attributions are to the cited versions listed in the bibliography; theorem numbers of the cited papers are given only where the text invokes a specific statement.

G.2 Generation in the limit and density

The model of Kleinberg & Mullainathan (2024) separates generation from identification: a generator succeeds if it eventually emits unseen elements of the target, and every countable family is generatable. Breadth asks for more, and Kalavasis et al. (2026) show that exact and approximate breadth revert to the identification conditions of Angluin (1980). Kleinberg & Wei (2025) introduce limit-set density as a quantitative notion of breadth, and Kleinberg & Wei (2026) study partial enumeration, where the stream reveals only a core of the target, and prove that half of the core's relative density is achievable and, against a racing adversary, optimal. The present paper keeps their protocol and adds the validity/value nesting: a verifier for the ambient language, trivia as the elements of the ambient outside the target, and a safety constraint that must hold on every instance of two indistinguishable targets. The recorded term $m_N/2$ of the frontier and the $\alpha/2$ endpoint are the KW race read on the larger target. The unrecorded term, its exchange rate of one with both converse and attainment, and the two corollaries are the new content.

G.3 Errors, hallucination rates and timely generation

Ganju et al. (2026) show that permitting sparse hallucinations overcomes coverage barriers when generation is time-sensitive, and Panigrahi et al. (2026) study hallucination rates, including finite-error and zero-rate regimes, and breadth hierarchies. The error in those works is emission outside the

Table 4: What is attributed and what is added, result by result.

Work	Attributed content	Increment here
Kalavasis et al. (2026)	exact and approximate breadth are equivalent to Angluin's condition and its weak form (Fact A.2)	the localization theorem, which transfers the characterizations fiber by fiber for every criterion judged in the limit; one new query per round; the separation and every-round-safety limits (Theorem 3.1 and Corollary 3.2, Theorem B.2 and Proposition B.5)
Kleinberg & Wei (2025; 2026)	limit-set density; the $\alpha/2$ partial-enumeration guarantee; the racing obstruction; identified intersections (Fact A.3)	the exact finite-time frontier $m_N/2 + \min\{g(T), n_N\}$ on an ambiguous record, with converse and attainment; the two-value law over all trivia regimes; the inverse-budget clock with exact leading constant (Theorem 4.1 and Corollaries 4.2 and 4.3)
Ganju et al. (2026)	sparse errors overcome coverage limits in timely generation, including partial enumeration	an explicit coverage metric, evaluation domain and two-truth safety requirement; the principle that sparse errors can help is not claimed as new
Panigrahi et al. (2026)	finite errors, zero error rates and breadth hierarchies, contemporaneous with this work	a different error notion (trivia inside a verified world), protocol and task; no priority claim on the general idea is made
Aksenov et al. (2026)	macros, wrapped length and the compression framework	self-contained constructions and counting proofs, including a density-zero nested pair; the framework is a carrier, not a theorem dependency or an axiom about value (Proposition 5.1, Theorem D.1 and Propositions D.5 and D.6)
Botta et al. (2025)	query complexity of verifier-assisted decoding with prefix-completability access	the verifier here recognizes the ambient language, and the results are informational, not computational
Colton et al. (2000); Tsoukalas et al. (2025)	interestingness measures in automated discovery	the complementary question of coverage given a perfect validity verifier; no interestingness measure is proposed
Gold (1967); Angluin (1980); Kleinberg & Mullainathan (2024)	identification in the limit; tell-tales; language generation in the limit	the model is an extension of these frameworks to a nested pair; their definitions are used unchanged

target; here the corresponding event is trivia inside a verified world, the ambient verifier removes hallucinations altogether, and the constraint is imposed across two truths. Corollary 4.2 shows that, on an ambiguous record, every notion of few trivia closed under taking fewer has one of only two optimal coverages, and that the finite-count and vanishing-rate regimes fall on different sides.

The quantitative comparison is also important. Ganju et al. (2026, Lemma 1) give an error-rate barrier on the scale of $D^{-1}(t)/t$ for deadline function D, and their Theorem 9 treats timely partial enumeration. Our converse fixes a two-target family and gives one deterministic racing transcript with an explicit prefix count valid simultaneously for every T and N, and our attainment matches it up to $O(\sqrt{N})$. Together they identify the recovery time as the inverse budget at the unrecorded mass. A general error–time tradeoff is not itself a new claim here.

Cai et al. (2026) subsequently simplify deterministic dense generation and establish the optimal randomized density $1 - 1/e$ against oblivious enumeration, including guarantees for multiple orders. Thus randomization in ordinary dense generation is already understood in that setting. Our deterministic shared-core theorem imposes a different, two-truth trivia constraint; extending that constrained objective requires specifying both the probabilistic coverage guarantee and whether its budget holds pathwise or in expectation.

G.4 Verifiers, interestingness and automated discovery

Botta et al. (2025) study how many verifier queries are needed to generate from a language given prefix-completability access; the verifier there supports search, whereas ours recognizes membership in the formal world and is silent about value. Interestingness in automated theory formation has been studied since Colton et al. (2000), and Tsoukalas et al. (2025) learn it from data; those works propose measures, whereas the present results ask what a generator can guarantee before any measure of value is available, with value modeled as membership in a latent target. Proof-search systems (Polu & Sutskever, 2020; Lample et al., 2022; Yang et al., 2023) and program search (Romera-Paredes et al., 2024) are the motivation for the ambient verifier; no claim about their behavior is derived from the theorems.

G.5 Compression as a model of mathematics

Aksenov et al. (2026) propose that mathematical value tracks compressibility under a growing dictionary of macros and contrast abelian and non-abelian monoids. We use their macro framework as a substrate for three explicit description rules (Appendix D.1.2). Section 5 develops the dictionary-only logarithmic construction, which places the entire harvest in the unrecorded region, and Appendix D treats the other rules and the counting proofs. These examples realize both the nested-target obstruction and the tell-tale regime. They neither identify compressibility with mathematical value nor establish a law separating commutative from noncommutative mathematics: the examples also differ in their records and tasks.

H How to read the results

The results rest on a gap between generating value and knowing how much of the valuable target remains unrecorded. A generator can stay within the part of the target that every possible world agrees on and always produce value, but it then omits every point on which the worlds disagree. On an ambiguous record, reaching those points is priced at exactly one trivia per discovery (Theorem 4.1), and the tagged-block construction makes the disputed region carry all of the density (Proposition 5.1). The distinctions below fix what these statements require and what they leave open.

H.1 The words of the title

Table 5 fixes what each word of the title denotes in the model.

Table 5: The title's words and the objects they denote.

Word	Precise reading
valuable mathematics	membership in the latent target H of the pair model, not a universal value function (Appendix A)
trivia	formally valid objects outside the latent target, $F \setminus H$; not invalid proofs or logical errors ((6))
necessity	to guarantee coverage of K' above $\alpha/2$ over all complete records of K, a generator must produce unboundedly many distinct trivia on some indistinguishable truth-K instance; any unbounded amount suffices (Corollary 4.2, (20))
provable	the necessity, its attainment, and the finite-time frontier are theorems in the stated protocol (Theorem 4.1 and Corollaries 4.2 and 4.3)
flood	persistent exploration of the region whose value cannot be told apart, at any prescribed rate; not high-frequency output, and not a recommendation to deposit results in a public repository
harvest	coverage of unrecorded value by the generated set; it does not include a rule for selecting which outputs are valuable (Appendix I)

H.2 The quantifiers

The necessity is a property of a coverage guarantee over indistinguishable instances. For every deterministic generator whose guaranteed coverage exceeds $\alpha/2$, there is a complete record of K on

which it outputs infinitely many points of $K' \setminus K$. Read under K, those points are trivia; read under K', they are discoveries. By Lemma A.1, both readings have the same output set. In the notation of Corollary 4.2,

$$\inf_{E \in \mathcal{E}(K)} \underline{d}(O(G,E), K') > \frac{\alpha}{2} \quad \Longrightarrow \quad \exists E_G \in \mathcal{E}(K) : \left| O(G, E_G) \cap (K' \setminus K) \right| = \infty, \tag{20}$$

which is the proof of Theorem 4.1(a) as $T \to \infty$: there the unrecorded part of each prefix is bounded by $|O \cap (K' \setminus K)|$, so if this set were finite the upper density would be at most $\alpha/2$. The target stays fixed within each instance. The record may depend on the generator, and the statement does not require trivia on every successful run: the attaining generator produces no trivia at all when the larger target is true, and its trivia under the smaller reading occupy a vanishing fraction of rounds.

H.3 Counts, rates and clocks

A trivia regime is any notion of few trivia that allows fewer whenever it allows more. Corollary 4.2 says that only one feature of a regime matters for eventual coverage, namely whether it contains an unbounded function. Finitely many distinct trivia, even without a uniform bound across instances, falls on one side. A prescribed budget of trivial rounds, however slowly it grows, and a vanishing trivia rate fall on the other. The proof explains why rare exploration suffices. Inside the record a statement must be reached before it is observed; outside it, a statement waits however long the generator takes. Infinitely many sweeps therefore recover the whole unrecorded region, while a vanishing fraction of sweep rounds leaves the race for recorded points intact.

Eventual coverage imposes no deadline. The frontier restores time: by round T at most $g(T)$ unrecorded points can have been harvested, and G_g harvests exactly that many. The recovery time is therefore the inverse budget at the unrecorded mass (Corollary 4.3). A slower budget preserves the eventual harvest and postpones it without limit. These temporal counts also differ from spatial precision: the fraction of compressible words in a raw-length ball (Proposition D.6) is a separate counting result, not a bound on the waiting time of the sweep.

H.4 What is fixed by the model

Coverage counts distinct strings under one canonical encoding, so whether two certificates of one statement count once or twice depends on that encoding. Densities depend on the canonical order. Generators are deterministic, computation is unrestricted, and the record may adapt to the outputs. Randomized generators against records fixed independently of their random choices are a different problem (Appendix I). The verifier answers membership in the formal world; it does not answer whether a valid statement is valuable. The observed core is part of the instance and is not, in general, known to the generator. These conventions specify the guarantees being compared.

I Discussion

I.1 Coverage is not selection

The generator of Theorem 4.1 covers the unrecorded region without ever learning where that region is. Its excursions are discoveries under K' and trivia under K, and every verifier query receives the same answer in both worlds. The harvest therefore arrives unlabeled: it can include previously unrecorded value without identifying which outputs deserve attention, and even the density-one coverage of the compression construction (Proposition 5.1) does not resolve this. In the language of Section 2, value is not a function of information. Verification supplies the validity coordinate and the record supplies a core; what remains is the ambiguous region, and any rule that selects outputs there must bring information from outside both. The discernment in Poincaré's epigraph is exactly this missing information.

I.2 Advice within the model

Read as advice to a designer of exploring systems, and only within the model, the theorems say three things. Exploration should be aimed at the region where the plausible accounts of value disagree, because only there does risked trivia buy discovery (Section 4). The rate of exploration sets the

clock but not the eventual harvest, so the budget should be chosen by the time one is willing to wait (Corollary 4.3). And no amount of verification substitutes for a signal about value, because verification resolves only the formal world (Theorem 3.1).

I.3 Open directions

Three directions seem natural. The first is selection with feedback. In the two-reading family a single label on one point of $K' \setminus K$ would settle which reading is true, so feedback matters when the record admits many readings. There, the number of labels needed to select the harvest while keeping its coverage would measure how much taste a generator must borrow. The second is randomization: the ledger (2) holds pathwise, so the gap $1 - \alpha$ should survive randomized generators while the race constant changes; against oblivious records dense generation attains $1 - 1/e$ (Cai et al., 2026), but the two-reading version is open. The third is richer families: the converse of Theorem 4.1 holds whenever the record might be complete, but when the record admits many nested readings, the readings compete for one budget and the exact frontier becomes an open Pareto question.