NEAR-LINEAR ADAPTIVITY GAPS IN PARALLEL MASKED GENERATION

Anonymous authors
Paper under double-blind review

ABSTRACT

Masked generators reduce decoding depth by filling several positions at once, but each round must decide both what values to generate and which positions to fill together. Existing work shows that adapting this choice can help, but leaves open whether the generated values themselves are necessary for fast and accurate generation. We prove that they are, even when the full target distribution is known and every per-position prediction is exact. On a four-symbol target family where every pair of positions is independent, a decoder that uses earlier values samples exactly in three rounds, whereas every randomized schedule that ignores them needs nearly linear depth to attain a fixed small error. We further prove that any short, accurate decoder must make its batch choices reflect the information revealed by those values. In contrast, for uniform sampling without replacement, generated values cannot improve the optimal schedule. These results identify generated values as a necessary scheduling resource, rather than merely a useful signal, and show when this necessity disappears.

1 INTRODUCTION

Masked and discrete-diffusion models can reduce decoding depth by generating several positions in one round (Lou et al., 2024; Sahoo et al., 2024; Nie et al., 2025). This speedup introduces a second decision beyond predicting the value at each position: the decoder must also decide which positions to generate together. Positions selected in the same round see the same incomplete output and cannot use one another's newly generated values. A larger batch is faster, but it may combine values that are individually plausible and jointly inconsistent. Even when the full target distribution is known and every per-position prediction is exact, are already generated values still necessary for choosing a fast and accurate schedule?

Existing work already recognizes that dependence between positions affects parallel generation. Current methods use confidence, attention, estimated dependence, or learned generation orders to select the next batch (Kim et al., 2025; Azangulov et al., 2025; Kim et al., 2026; Ringel et al., 2026; Zoabi et al., 2026; Fox et al., 2026; Wu et al., 2025). WeDLM combines confidence-based streaming decoding with causal attention and prefix caching (Liu et al., 2025). Theory studies the errors caused by parallel updates and the design of efficient generation schedules (Shi et al., 2024; Chen et al., 2026a; Zhang et al., 2026). Other methods allow positions in the same batch to interact through joint generation, verification, repair, or remasking (Li et al., 2026; Chen et al., 2026b; Jin et al., 2026). Together, these works show that dependence matters and that adapting the generation process can help. What remains unresolved is whether using generated values is merely one effective way to choose batches, or whether some targets make it unavoidable for fast and accurate generation.

Our key insight is that knowing every possible dependency pattern is different from knowing which pattern applies to the current sample. The full target distribution provides the former, while generated values can reveal the latter. If the same unfilled positions require different safe batches under different observed values, a schedule that ignores those values must remain cautious across all possibilities. A decoder that uses them only needs to respect the dependencies present in the current sample. Conversely, when the safe choice is the same across all possible histories, observing more values cannot improve the schedule. Figure 1 shows how generated values change which batch is safe.

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.

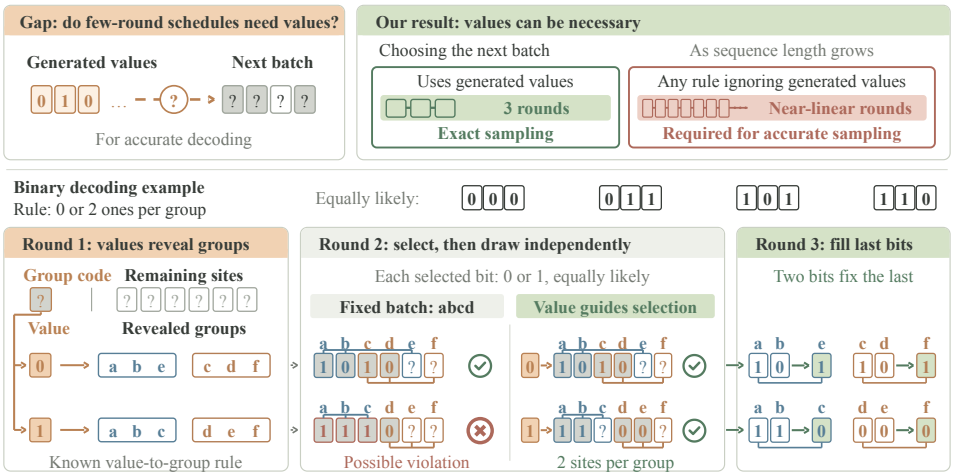

Figure 1: **Generated values determine safe batches.** In the binary example, the first value reveals groups. Sampling two bits per group determines the third; a fixed batch may violate the constraint despite exact predictions. Our separation holds for pairwise-independent four-symbol targets.

Our contributions are:

- **We prove that generated values are necessary for fast and accurate scheduling on some targets.** On a four-symbol target family where every pair of positions is independent, using generated values enables exact sampling in three rounds, while every randomized schedule that ignores them requires nearly linear depth to attain a fixed small error.
- **We quantify what information a fast decoder must use.** We prove that every short, accurate decoder must choose batches that sufficiently distinguish the possible groupings revealed by earlier values.
- **We exactly characterize an important opposite case.** For uniform sampling without replacement, current values provide no scheduling gain; the optimum is fully determined by the remaining count, with an exact round–error tradeoff and optimal batch sizes.

2 Background and Problem

Parallel masked generation. Masked generation begins with a set of unfilled positions. In each round, the model predicts a probability distribution for every remaining position from the values currently visible, and the decoder selects several positions to fill. The selected positions are sampled separately and then fixed. Their new values can affect later rounds, but they cannot affect one another within the current round. Increasing the batch size reduces the number of rounds, but it also commits more values using incomplete information. Work on generation schedules studies how to manage this tradeoff efficiently (Shi et al., 2024; Chen et al., 2026a; Zhang et al., 2026).

What "exact prediction" means. An exact prediction gives the correct current probability for one position; it does not imply that several such predictions can be sampled independently in the same round. Consider three binary positions with four equally likely valid outputs: 000, 011, 101, and 110. Every position is individually uniform, and every pair is independent. Nevertheless, independently generating all three positions produces an invalid output with probability one half. Generating any two positions first is safe because their values determine the third. Every per-position prediction is correct in this example. The error comes only from placing too many positions in the same batch.

Two ways to choose the next batch. A content-adaptive schedule may use values generated in earlier rounds when selecting the next batch. Selection based on predictions or internal states that depend on these values is also content-adaptive. A preplanned schedule may use the target distribution, position identities, and its own randomness, but its decisions cannot depend on generated values. Both know the same target and use the same exact per-position predictions. We therefore

vary only the information used to choose a batch. Algorithms that count valid completions or query richer conditional information have access to stronger sampling operations (Anari et al., 2024; 2025). Joint generation, verification, repair, and remasking also change how values in a batch are produced (Li et al., 2026; Chen et al., 2026b; Jin et al., 2026). WeDLM instead combines confidence-based streaming selection with causal attention and prefix caching (Liu et al., 2025). The comparison here isolates batch selection under a fixed product transition.

The remaining gap. Existing results show that adaptive schedules can be useful, but usefulness does not establish necessity. The unresolved question is whether a target can force every fast and accurate decoder to use already generated values when choosing later batches, even when the target is fully known and all per-position predictions are exact. We measure speed by the number of decoding rounds and accuracy by the difference between the generated and target distributions. Section 3 first identifies when a batch of separately sampled positions is safe. Sections 4 and 5 then prove that generated values are necessary on one target family and provide no scheduling advantage on another.

3 When Is a Parallel Commit Safe?

Let P have full support on a nonempty valid set $\mathcal{V} \subseteq \prod_i \mathcal{X}_i$. At an extendible history h, let $U(h)$ be the hidden sites. A batch $B \subseteq U(h)$ independently draws the exact current singleton at each selected site and commits the draws irreversibly.

A safe batch has rectangular support. Let $S_i(h)$ be the conditional support of coordinate i, and let

$$E_{h,B} := \{y_B : h \text{ followed by } y_B \text{ has a completion in } \mathcal{V}\}.$$

The product transition can draw every point in $\prod_{i\in B} S_i(h)$, giving the exact criterion below.

Lemma 3.1 (Support rectangularity)**.** The batch B stays within the extendible target support with probability one if and only if

$$\boxed{E_{h,B} = \prod_{i\in B} S_i(h).} \tag{1}$$

Proof. The batch-law support is $\prod_{i\in B} S_i(h)$, while $E_{h,B}$ contains exactly its extendible assignments; safety is equality of these sets. □

Rectangularity concerns support, not independence. Minimal unsafe batches form a history-dependent conflict hypergraph, and pairwise checks may miss the first obstruction. The three-bit example in Section 2 is exactly such a size-three conflict: every pair is possible, but the product of all three singletons leaves the valid set with probability $1/2$.

Support safety is not full-law correctness. Fix any completion kernel after a history becomes nonextendible. This makes the output law Q well defined; irreversibility makes $Q(\mathcal{V})$ independent of the chosen completion. Staying extendible is still weaker than sampling the right law along every reached history. Our upper bounds therefore prove $Q = P$ by the chain law, while lower bounds use $\mathrm{TV}(P, Q) \geq 1 - Q(\mathcal{V})$.

How observations change concrete actions. Observed values can guide scheduling by changing which concrete batch is best. Consider two histories that reveal the same sites. Suppose relabeling their values maps one residual law to the other while keeping every site identity fixed. Each batch then has corresponding outcomes with equal probabilities and the same continuation value at both histories. We call these histories site-labeled equivalent.

Proposition 3.2 (When realized values cannot help)**.** Suppose that, for every set of revealed sites, all extendible value histories are site-labeled equivalent. Then, for every remaining round budget, all those histories assign the same continuation value to every concrete batch. Consequently, a deterministic policy that uses the revealed sites but not their values attains the optimal history-aware target-support probability. If site symmetries also make all same-size revealed sets and same-size batches equivalent, an optimal policy can be chosen to depend only on the remaining count and round budget.

Proof. Induct on the remaining round budget. The claim is immediate with no rounds left. For the inductive step, site-labeled equivalence pairs the outcomes of any concrete batch, with the same probabilities, and pairs their child histories. By induction, paired children have the same optimal continuation value. Thus every concrete batch has the same value at all equivalent histories, so these histories share a maximizer that does not use revealed values. Under the additional site symmetry, batches of the same size also have the same value, leaving only the batch size to choose. □

This proposition isolates a label-homogeneity condition that equalizes concrete action values. Uniform injections realize it exactly and collapse to count schedules. The next section constructs histories whose values locate conflicts hidden from every unconditional coordinate pair.

4 Pairwise-Identical Fixed-Alphabet Targets

We now construct histories that change the value of concrete batch actions while leaving every unconditional coordinate pair uninformative. We first state a compact realization that gives the clean linear rate, then lift the same mechanism to actual coordinates with alphabet at most four. Throughout, depth counts at most T nonempty rounds and permits early stopping. Both policy classes use the same irreversible product of exact current singleton conditionals.

4.1 One principle, two realizations

The result has two realizations of the same phenomenon. A compact target gives the clean linear depth gap. A second target spreads the same hidden structure over ordinary binary and four-valued coordinates, showing that the gap is not created by one large symbol. In both cases, every unconditional coordinate pair agrees with an independent baseline; only access to realized values changes the attainable depth.

Theorem 4.1 (Pairwise-identical adaptivity separation)**.** For every sufficiently large N divisible by three, there are a compact target P_N^{C} and a fixed-alphabet target P_N^{FA}, of actual length $L = \Theta(N \log N)$, that can be chosen simultaneously so that:

1. every two distinct compact coordinates of P_N^{C} are independent, and every two distinct actual coordinates of P_N^{FA} are independent; hence each target has the same unconditional one- and two-coordinate tables as its coordinatewise-independent baseline;
2. both constrained targets have exact history-aware depth three, whereas their independent baselines have exact depth one; and
3. for every $0 \le \varepsilon \le 1/20$, every randomized value-oblivious policy attaining TV error at most ε needs more than $N/64$ rounds for either target.

Consequently, the compact separation is linear in its coordinate count, while the fixed-alphabet separation is $\Omega(N) = \Omega(L/\log L)$ in its actual length.

These axes isolate access to realized values as the controlled difference. The two policy classes face the same fixed, known target and use the same product transition and full-law TV metric. The alphabet-four construction places the separation on the actual scheduled coordinates, and the lower bound covers arbitrary fixed interleavings. Realized values alone can therefore create a near-linear depth gap despite identical unconditional pair tables.

4.2 Compact route intuition

Let $q = 2^k = \Theta(N \log N)$. Fix a family $(H_z)_{z \in \mathbb{F}_q}$ of perfect three-uniform hypermatchings on the N labeled payload groups. This family is part of the target and is not resampled by the decoder. The hidden value $Z = z$ selects one matching H_z, called the route, which specifies the triples of groups that share a parity constraint.

Choose distinct field points $\xi_1, \xi_2, \xi_3, \lambda_1, \ldots, \lambda_N$ and draw

$$f(t) = U + Vt + Zt^2, \qquad (U, V, Z) \sim \mathrm{Unif}(\mathbb{F}_q^3).$$

The compact target has anchors $Y_a = f(\xi_a)$. Payload coordinate i contains the tag $T_i = f(\lambda_i)$ and a bit C_i. Conditional on $Z = z$, every triple of H_z carries an independent uniform even-parity bit vector. The route is invisible to any coordinate pair: two evaluations have rank two, two parity shares are fair and independent, and two tags leave Z uniform. Thus all compact pairs have the same law as in the independent baseline.

Three jointly revealed anchors interpolate f, expose Z, and locate the matching. Round two reveals two vertices per matching triple; round three reveals the parity-determined last vertices. Each round is the product of exact current singletons, giving full-law equality.

4.3 A GUARDED FOUR-SYMBOL LIFTING

Each compact payload coordinate becomes a group of smaller coordinates, each holding a tag bit and a parity share. The extra tag bit preserves the guard under arbitrary interleavings. Fix a binary basis of $\mathbb{F}_q$, with coordinate functionals $\rho_1, \dots, \rho_k$. The anchors become $3k$ binary coordinates. With the redundant functional $\rho_{k+1} = \rho_1 + \cdots + \rho_k$, payload group i contains coordinates

$$W_{i,b} = \big(\rho_b(T_i), C_{i,b}\big) \in \{0,1\}^2, \qquad 1 \le b \le k+1. \tag{2}$$

This gives

$$L = 3k + N(k+1) = \Theta(N \log N) \tag{3}$$

actual coordinates. For each route triple $e = \{i, j, \ell\} \in H_Z$, its $3(k+1)$ shares are uniform subject to one total even-parity constraint. The observed alphabet is therefore two on anchors and four on payload coordinates.

Any k of the $k+1$ tag functionals recover the field evaluation. A block with only one hidden coordinate therefore determines the quadratic, giving the guard below.

Lemma 4.2 (Pre-crossing guard)**.** At a target-extendible history h before the observed tags determine the quadratic, every candidate route block has at least two hidden actual coordinates. If x is an unrevealed payload coordinate with tag D_x and share C_x, then

$$\boxed{P(D_x, C_x \mid h) = P(D_x \mid h) \otimes \mathrm{Ber}(1/2).}$$

In plain terms, even if tags drawn in the current batch reveal the route, they do not coordinate the shares drawn beside them. Each share remains fair, and another hidden share remains to carry the parity rule. Appendix B proves the lemma and shows why the naive k-coordinate split lacks this protection.

Lemma 4.3 (Actual-coordinate pairwise identity)**.** Every pair of distinct actual coordinates under P_N^{FA} is independent. Thus P_N^{FA} and $\overline{P}_N^{\mathrm{FA}} := \bigotimes_R P_N^{\mathrm{FA}}(R)$ have identical complete unconditional pair tables.

Proof. The anchor bits are jointly uniform. An anchor–payload pair factors because one anchor evaluation and one distinct payload evaluation leave the route uniform. Within a payload group, distinct tag functionals are independent and the two shares form a proper parity marginal. Across groups, conditioning on the route again leaves both tags and both shares pairwise uniform. These cases cover all pairs of actual coordinates; Appendix B gives the full conditioning argument. □

The exact sampler has three rounds. Round one reveals the anchors, recovering f, the route, and every payload tag. Round two reveals all but one actual coordinate per route block with the correct proper parity marginal; round three fills the now-determined coordinate. These product kernels multiply to the full target law. Proposition B.1 gives randomized polynomial preprocessing and polynomial-time sampling and singleton queries.

4.4 WHY ARBITRARY SHORT VALUE-OBLIVIOUS SCHEDULES FAIL

Each anchor or tag bit is a linear form in the bits of (U, V, Z); the scalar rank is the rank of the revealed forms. Conditioning on the policy tape fixes all batches. If any batch overcommits fresh tag information before or upon reaching full scalar rank, at least half the support mass is lost. Otherwise the successive tag draws jointly reveal a uniform quadratic. The guard then reduces the schedule to group-completion rounds: two groups in one true block finishing last together satisfy the required parity only half the time.

Overcommitting fresh tag information. At the start of a batch, suppose m tag bits are not already determined but together add only $\Delta < m$ ranks. The product transition draws 2^m equiprobable strings, of which 2^Δ extend the affine fiber, so scalar consistency has probability $2^{-(m-\Delta)} \le 1/2$.

Reaching full scalar rank. Otherwise every batch has zero scalar nullity until the scalar rank becomes $3k$. Across all consistent tag outcomes, the successive fresh bits reveal a uniform quadratic and hence a uniform route. By the guard, shares observed before this crossing do not reweight the affine scalar posterior, and shares drawn in the crossing batch remain fair after conditioning on its tags.

From actual coordinates to hidden triples. Project the coordinate schedule to $g(i) = \max_b \mathrm{round}(W_{i,b})$, the last reveal round of payload group i. A repeated maximum on a true route triple (a terminal tie) places at least two of that block's last actual coordinates in one product batch. For example, suppose the rank-completing batch also closes such a block with two coordinates remaining. Before the batch, the guard leaves both coordinates hidden; conditioning on all same-batch tags still leaves their shares as independent fair singleton draws. The target permits only one of the two parity classes, so the block survives with probability $1/2$. The same argument applies after the route is known. Disjoint route blocks use disjoint shares, so these factors multiply. If Y_z counts the repeated maxima for route z, then

$$Q(\mathrm{supp}\, P_N^{\mathrm{FA}}) \le \frac{1}{q} \sum_{z \in \mathbb{F}_q} 2^{-Y_z}. \tag{4}$$

The preceding analysis concerns one route and one projected schedule. We now fix a single route family that makes the loss unavoidable for every short schedule at once. For fixed $c : [N] \to [T]$, $T \le N/64$, a uniform perfect 3-hypermatching has at least five terminal ties with probability greater than $1/9$. With $q > 144N \log(N+1)$ independent matchings, Chernoff makes the probability of fewer than $q/18$ such routes at most $e^{-q/72}$. There are at most N^N short colorings and $N^N e^{-q/72} < 1$; hence one fixed family has $|\{z : Y_{H_z}(c) \ge 5\}| \ge q/18$ simultaneously for every $T \le N/64$ and $c : [N] \to [T]$. Appendix A gives the moment calculation and constants. Substitution into (4) yields

$$Q(\mathrm{supp}\, P_N^{\mathrm{FA}}) \le 1 - \frac{1}{18}(1 - 2^{-5}) = \frac{545}{576}.$$

The target puts mass one on its support, so $\mathrm{TV}(P_N^{\mathrm{FA}}, Q) \ge 31/576 > 1/20$. Averaging over the random tape covers randomized policies. The compact proof has the identical terminal-tie branch and a field-valued scalar-overcommit branch; full details appear in Appendices A and B.

Finally, no history-aware two-round policy can be exact. Its first batch is chosen at the empty history and hence, after conditioning on its random tape, is value-oblivious; complete output forces the second batch to be its complement. The preceding bound applies for sufficiently large N. The three-round upper bound is therefore sharp.

4.5 What accurate short actions must reveal

A short decoder need not reconstruct the entire hidden route. It must, however, make different batch choices for enough different groupings: if the same completion pattern is reused for too many routes, some route triple closes unsafely. To quantify this requirement, record the round in which each payload group finishes. Safety on the true route requires one uniquely last-finishing group in every true route triple. We call this completion-round vector the *action cover*.

Different routes can share a cover. For $x \in [T]^N$, let $\mathcal{L}(x)$ contain the routes whose triples all have a unique last group. The fixed family satisfies $|\mathcal{L}(x)| \le \Lambda_T$, where $\Lambda_T := \lceil \max\{168T \log T, e\, q\, e^{-N/(28T)}\} \rceil$. Set $\bar{\Lambda}_T := \min\{q, \Lambda_T\}$ and $G_T := \log(q/\bar{\Lambda}_T)$, the route information eliminated by this list.

Let $\mathsf{C} \in [T]^N$ be the compact cover. In the lifting, let $\mathsf{A} \in [T]^{N \times (k+1)}$ record actual-coordinate rounds and $\mathsf{G}_i = \max_b \mathsf{A}_{i,b}$. Thus $\mathsf{X} = \mathsf{C}$ in the compact target and $\mathsf{X} = \mathsf{G}$ in the lifting.

Theorem 4.4 (Action-cover information)**.** The route family can be chosen so that, for every $2 \le T \le N/64$, every randomized exact history-aware policy using at most T rounds satisfies

$$I(Z;\mathsf{X}) \ge \min\left\{\log \frac{q}{168T\log T}, \frac{N}{28T} - 1\right\} - \log 2,$$

where $\mathsf{X} = \mathsf{C}$ for the compact target and $\mathsf{X} = \mathsf{G}$ for the lifting; moreover $I(Z;\mathsf{A}) \ge I(Z;\mathsf{G})$. Conversely, for the compact target and $T \ge 4$, an exact policy attains

$$I(Z;\mathsf{C}) \le \log\left\lceil \frac{q}{\lfloor (T-2)/2 \rfloor} \right\rceil.$$

Exactness places the true route in a short surviving list, giving the entropy lower bound. Grouping routes with a common safe coloring gives the compact upper construction. The information measured is in the batch choices; several routes can share one cover. Full proofs are in Appendices A.6 and B.5. The next theorem allows output error; its corollary applies the bound to limited observation channels.

Theorem 4.5 (Action information–error frontier)**.** For either realization, let P be its target law, Q the output law of any $T \le N/64$ policy, and X its action cover. If $d = \mathrm{TV}(P,Q) \le 1/4$, there is a uniform continuation route $Z^\star$, equal to the realized route on valid outputs, such that

$$I(Z^\star;\mathsf{X}) \ge (1-2d)G_T - h_2(2d), \qquad d \ge \frac{1}{2}\left[1 - \frac{B + \log 2}{G_T}\right]_+,$$

where the second bound applies when $I(Z^\star;\mathsf{X}) \le B$ and h_2 is binary entropy in nats.

Proof. The first batch with scalar overcommitment or an actual-coordinate terminal tie fails with probability at least $1/2$, so it is reached with probability at most $2d$. A coupling with the target law places uniform $Z^\star$ outside the surviving list with probability at most $2d$, and list-Fano gives the theorem. Appendix C gives the full proof. □

Corollary 4.6 (Bounded observation channels)**.** If the complete observation transcript O, conditioned on the fixed route table $\mathcal{H} = (H_z)_z$ and independent policy tape R, satisfies $I(Z^\star;\mathsf{O} \mid \mathcal{H}, R) \le B$, it obeys the same error bound. Fixed B and error below $1/4$ require $\Omega(N) = \Omega(L/\log L)$ rounds; a constant new-information budget per round requires $\Omega(\log N)$.

Appendix C.1 gives the causal decision interface and proof.

4.6 Robustness under support-preserving errors

The exact construction also tolerates probability error inside the correct local support. At every target-extendible history, replace each exact singleton by a kernel within TV δ; for the positive result require the approximate and exact singleton supports to agree.

Corollary 4.7 (Support-preserving robustness)**.** Fix $0 < c < 1/20$ and set $\delta_L = c/\sqrt{L}$. The approximate three-round history-aware policy has $\mathrm{TV}(P_N^{\mathrm{FA}}, \widehat{Q}) \le c + o(1) < 1/20$, whereas every randomized value-oblivious policy using at most $N/64$ rounds has $\mathrm{TV}(P_N^{\mathrm{FA}}, \widehat{Q}) \ge 31/576 - o(1) > 1/20$. Thus the $\Omega(L/\log L)$ depth gap survives an inverse-square-root singleton-error window.

The support condition separates calibration error inside the valid local set from leakage into impossible symbols. If each of the $N/3$ deterministic final shares leaks probability δ to an impossible value, output TV is exactly $1 - (1-\delta)^{N/3}$. Appendix D proves the corollary using pathwise KL for the three-round chain and a shadow-extendible continuation for arbitrary value-oblivious schedules.

Thus unconditional pair tables can hide nearly a linear factor in depth even with imperfect singleton probabilities. We next give a symmetric class where values cannot change any concrete action value.

5 The Sharp Opposite: Uniform Permutations

Uniform all-different targets give the complementary mechanism: despite strong dependence, label homogeneity equalizes concrete action values. Here T counts exactly T nonempty rounds, rather than at most T as in Section 4.

5.1 AN EXACT FULL-POLICY FRONTIER

Let

$$\mathcal{I}_{n,q} := \{x \in [q]^n : x_i \neq x_j \text{ for } i \neq j\}, \qquad q \geq n,$$

and let $P_{n,q}$ be uniform on $\mathcal{I}_{n,q}$. At a valid history with r hidden coordinates, exactly $q - n + r$ symbols remain unused. Every hidden coordinate has the same uniform singleton conditional on this set. Consequently, a product batch of size b remains all-different with probability

$$p(r, b) = \frac{(q - n + r)_b}{(q - n + r)^b}. \tag{5}$$

Proposition 5.1 (Fixed-schedule identities)**.** Consider a count schedule with remaining counts $n = r_1 > r_2 > \cdots > r_{T+1} = 0$ and batch sizes $b_s = r_s - r_{s+1}$. Define

$$S(\mathbf{r}) := \frac{(q)_n}{\prod_{s=1}^{T}(q - n + r_s)^{b_s}}.$$

For any randomized site-selection rule that follows these counts,

$$Q_{\mathbf{r}}(\mathcal{I}_{n,q}) = S(\mathbf{r}), \qquad \mathrm{TV}(P_{n,q}, Q_{\mathbf{r}}) = 1 - S(\mathbf{r}), \qquad \mathrm{KL}(P_{n,q} \| Q_{\mathbf{r}}) = -\log S(\mathbf{r}). \tag{6}$$

Conditional on validity, the output law is exactly $P_{n,q}$.

The falling factorials telescope to $(q)_n$; every valid output has the same product probability even along a site-adaptive path, so the error is pure support leakage.

Let $V_t(r)$ be the maximum valid completion probability from r hidden coordinates in exactly t nonempty rounds. Then

$$V_t(r) = \max_{1 \leq b \leq r - t + 1} p(r, b) V_{t-1}(r - b), \qquad V_0(0) = 1, \quad V_0(r > 0) = 0. \tag{7}$$

Theorem 5.2 (Uniform full-policy collapse)**.** Among all policies using sites, values, and randomness, the largest valid mass is $V_T(n)$, attained by a deterministic count schedule from (7). Every such policy satisfies

$$\mathrm{TV}(P_{n,q}, Q) = 1 - Q(\mathcal{I}_{n,q}).$$

Valid histories with the same remaining count share their factors and continuation values. Bellman induction therefore collapses every adaptive action to (7), realizing Proposition 3.2 exactly.

5.2 THE PERMUTATION PHASE DIAGRAM

Set $q = n$ and $M_{n,T} := V_T(n)$. Fixed rounds give exponentially small valid mass, while sublinear T gives $-\log M_{n,T} = \Theta(n/T)$; exact forms are in Appendix H. On the linear scale, let $\Phi(\tau)$ be the explicit function in (19).

Theorem 5.3 (Linear-round limit and optimal profile)**.** If $T_n/n \to \tau \in (0, 1]$, then

$$M_{n,T_n} \longrightarrow \Phi(\tau).$$

Every sequence of exact optimal schedules converges to the same threshold profile over batch sizes $1, \ldots, K$, characterized explicitly in Appendix I; mass above K vanishes.

Optimal schedules use larger batches first; in the linear-round limit, they end with singletons. Appendix I specifies the cutoff $K = K(\tau)$, with $K = 1$ at $\tau = 1$. Uniform values only relabel a common geometry; the anchors in Theorem 4.1 locate the grouping.

6 EXPERIMENTS

Setup. We compare exact sampling laws, the effect of sequence length on scheduling, and sensitivity to errors in singleton probabilities. For $N \in \{192, 384, 768, 1536\}$, we evaluate six fixed schedules (round-robin, contiguous blocks, and four balanced random permutations) at ten round budgets, using 8,192 independently sampled routes per comparison. We report the most favorable schedule. Appendix E gives the full experimental protocol and statistical bounds.

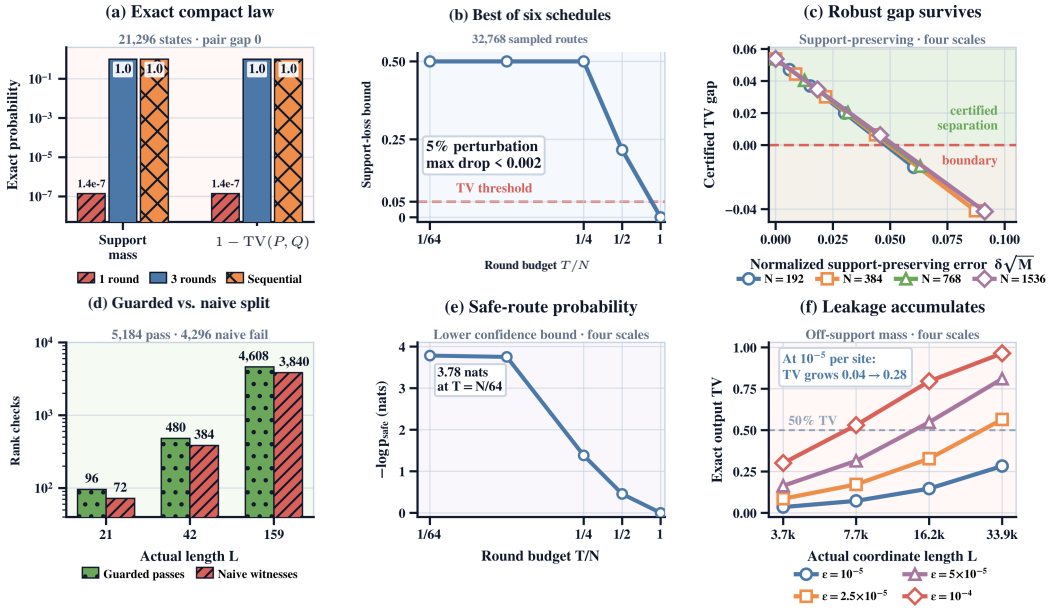

Figure 2: **Parallel sampling with hidden dependencies. (a,d)** Exact enumeration compares sampling laws and tests the role of redundant tags. **(b,e)** Safe routes remain rare at short budgets across four scales, even for the best of six fixed schedules. **(c,f)** Support-preserving error bounds and the accumulation of off-support leakage.

Exact sampling and the role of the guard. Exact enumeration finds identical pair tables for the target and independent baseline, yet one-round sampling fails to recover the target joint law. Three rounds are exact (panel a). The guarded tags still determine the quadratic when any one tag bit in a block is withheld; the naive split loses this property (panel d). The extra tag bit therefore preserves the guard needed by the construction.

The scheduling obstruction persists across scales. At $T = N/64$, none of the six schedules yields a sampled safe route at any tested size, and the smallest sampled support-loss bound across schedules is 0.5. It stays at 0.5 up to $T = N/4$ and exceeds 0.2 at $T = N/2$ at every size. The frontiers in panels (b,e) align under T/N, so the obstruction persists across the tested sizes at comparable relative budgets.

How the support of the error affects robustness. Support-preserving errors leave the sampled support-loss frontier nearly unchanged at $\delta = 0.05$ (panel b). The theoretical bound in panel (c) gives a positive gap when the error times the square root of the fair-draw count is small. Corollary 4.7 requires approximate singletons to keep the exact support because leakage outside it accumulates with length: a mass of 10^{-5} on an impossible tag at each round-two payload draw raises TV from 0.035 to 0.283 across the four lengths (panel f).

7 Conclusion

We prove that generated values can be necessary for fast, accurate parallel masked generation even with exact per-position predictions. Our pairwise-independent four-symbol construction admits exact sampling in three rounds, while every randomized schedule that ignores generated values requires nearly linear depth at a fixed small error. For uniform sampling without replacement, batch sizes chosen in advance are optimal. These results show how generated values can reduce decoding depth by revealing which positions can safely be sampled together.

AI USE STATEMENT

In this work, we used generative AI tools for language polishing and improving readability. The authors reviewed the AI-assisted revisions and take responsibility for the final content of this work.

REFERENCES

Giries Abu Ayoub, Mario Barbara, Lluís Pastor-Pérez, Tanja Bien, Aneesh Barthakur, Alaa Maalouf, and Loay Mualem. Supportive token revealing for fast diffusion language model decoding. *arXiv preprint arXiv:2606.04236*, 2026. URL https://arxiv.org/abs/2606.04236.

Nima Anari, Ruiquan Gao, and Aviad Rubinstein. Parallel sampling via counting. *arXiv preprint arXiv:2408.09442*, 2024. URL https://arxiv.org/abs/2408.09442.

Nima Anari, Carlo Baronio, CJ Chen, Alireza Haqi, Frederic Koehler, Anqi Li, and Thuy-Duong Vuong. Parallel sampling via autospeculation. *arXiv preprint arXiv:2511.07869*, 2025. URL https://arxiv.org/abs/2511.07869.

Iskander Azangulov, Teodora Pandeva, Niranjani Prasad, Javier Zazo, and Sushrut Karmalkar. Parallel sampling from masked diffusion models via conditional independence testing. *arXiv preprint arXiv:2510.21961*, 2025. URL https://arxiv.org/abs/2510.21961.

K. A. Bush. Orthogonal arrays of index unity. *The Annals of Mathematical Statistics*, 23(3):426–434, 1952. URL https://www.jstor.org/stable/2236685.

Panagiotis Cheilaris, Balázs Keszegh, and Dömötör Pálvölgyi. Unique-maximum and conflict-free coloring for hypergraphs and tree graphs. *SIAM Journal on Discrete Mathematics*, 27(4):1775–1787, 2013. doi: 10.1137/120880471. URL https://doi.org/10.1137/120880471.

Sitan Chen, Kevin Cong, and Jerry Li. Optimal inference schedules for masked diffusion models. In *Proceedings of the Thirty-Ninth Conference on Learning Theory*, volume 336 of *Proceedings of Machine Learning Research*, pp. 1279–1311. PMLR, 2026a. URL https://proceedings.mlr.press/v336/chen26e.html.

Zigeng Chen, Gongfan Fang, Xinyin Ma, Ruonan Yu, and Xinchao Wang. DMax: Aggressive parallel decoding for dLLMs. *arXiv preprint arXiv:2604.08302*, 2026b. URL https://arxiv.org/abs/2604.08302.

David Fox, Sam Bowyer, Song Liu, Laurence Aitchison, Raul Santos-Rodriguez, and Mengyue Yang. Learning generation orders for masked discrete diffusion models via variational inference. *arXiv preprint arXiv:2602.23968*, 2026. URL https://arxiv.org/abs/2602.23968.

Yijie Jin, Jiajun Xu, Yuxuan Liu, Chenkai Xu, Yi Tu, Jiajun Li, Dandan Tu, Xiaohui Yan, Kai Yu, Pengfei Liu, and Zhijie Deng. Multi-block diffusion language models. *arXiv preprint arXiv:2606.29215*, 2026. URL https://arxiv.org/abs/2606.29215.

Bumjun Kim, Dongjae Jeon, Moongyu Jeon, and Albert No. DAPD: Dependency-aware parallel decoding via attention for diffusion LLMs. *arXiv preprint arXiv:2603.12996*, 2026. URL https://arxiv.org/abs/2603.12996.

Jaeyeon Kim, Kulin Shah, Vasilis Kontonis, Sham M. Kakade, and Sitan Chen. Train for the worst, plan for the best: Understanding token ordering in masked diffusions. In *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pp. 30749–30768. PMLR, 2025. URL https://proceedings.mlr.press/v267/kim25ah.html.

Ian Li, Zilei Shao, Benjie Wang, Rose Yu, Guy Van den Broeck, and Anji Liu. Breaking the factorization barrier in diffusion language models. *arXiv preprint arXiv:2603.00045*, 2026. URL https://arxiv.org/abs/2603.00045.

Aiwei Liu, Minghua He, Shaoxun Zeng, Sijun Zhang, Linhao Zhang, Chuhan Wu, Wei Jia, Yuan Liu, Xiao Zhou, and Jie Zhou. WeDLM: Reconciling diffusion language models with standard causal attention for fast inference. *arXiv preprint arXiv:2512.22737*, 2025. URL https://arxiv.org/abs/2512.22737.

Aaron Lou, Chenlin Meng, and Stefano Ermon. Discrete diffusion modeling by estimating the ratios of the data distribution. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 32819–32848. PMLR, 2024. URL https://proceedings.mlr.press/v235/lou24a.html.

Omer Luxembourg, Haim Permuter, and Eliya Nachmani. Plan for speed: Dilated scheduling for masked diffusion language models. *arXiv preprint arXiv:2506.19037*, 2025. URL https://arxiv.org/abs/2506.19037.

Shen Nie, Fengqi Zhu, Zebin You, Xiaolu Zhang, Jingyang Ou, Jun Hu, Jun Zhou, Yankai Lin, Ji-Rong Wen, and Chongxuan Li. Large language diffusion models. *arXiv preprint arXiv:2502.09992*, 2025. URL https://arxiv.org/abs/2502.09992.

Heqiang Qi, Wei Huang, Mingyuan Bai, and Xiangming Meng. Cluster-level attention-guided parallel decoding for masked diffusion language models. *arXiv preprint arXiv:2605.29607*, 2026. URL https://arxiv.org/abs/2605.29607.

Liran Ringel, Ameen Ali, and Yaniv Romano. Dependency-guided parallel decoding in discrete diffusion language models. *arXiv preprint arXiv:2604.02560*, 2026. URL https://arxiv.org/abs/2604.02560.

Subham Sekhar Sahoo, Marianne Arriola, Yair Schiff, Aaron Gokaslan, Edgar Marroquin, Justin T. Chiu, Alexander Rush, and Volodymyr Kuleshov. Simple and effective masked diffusion language models. In *Advances in Neural Information Processing Systems*, volume 37, 2024. doi: 10.52202/079017-4135. URL https://proceedings.neurips.cc/paper_files/paper/2024/hash/eb0b13cc515724ab8015bc978fdde0ad-Abstract-Conference.html.

Adi Shamir. How to share a secret. *Communications of the ACM*, 22(11):612–613, 1979. doi: 10.1145/359168.359176. URL https://doi.org/10.1145/359168.359176.

Jiaxin Shi, Kehang Han, Zhe Wang, Arnaud Doucet, and Michalis K. Titsias. Simplified and generalized masked diffusion for discrete data. In *Advances in Neural Information Processing Systems*, volume 37, 2024. URL https://arxiv.org/abs/2406.04329.

Elad Verbin and Wei Yu. The streaming complexity of cycle counting, sorting by reversals, and other problems. In *Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 11–20, 2011. doi: 10.1137/1.9781611973082.2. URL https://doi.org/10.1137/1.9781611973082.2.

Chengyue Wu, Hao Zhang, Shuchen Xue, Zhijian Liu, Shizhe Diao, Ligeng Zhu, Ping Luo, Song Han, and Enze Xie. Fast-dLLM: Training-free acceleration of diffusion LLM by enabling KV cache and parallel decoding. *arXiv preprint arXiv:2505.22618*, 2025. URL https://arxiv.org/abs/2505.22618.

Shaorong Zhang, Longxuan Yu, Rob Brekelmans, Luhan Tang, Salman Asif, and Greg Ver Steeg. Generation order and parallel decoding in masked diffusion models: An information-theoretic perspective. *arXiv preprint arXiv:2602.00286*, 2026. URL https://arxiv.org/abs/2602.00286.

Yangyang Zhong, Yanmei Gu, Zhengqing Zang, Xiaomeng Li, Yuqi Ding, Xibei Jia, Yuting Shen, Zhenzhong Lan, Liwang Zhu, Weiping Liu, Junlin Zhou, Haisheng Liu, Zhong Xin Yu, Pengxin Luo, Donglian Qi, Yunfeng Yan, and Junbo Zhao. Parallelism and generation order in masked diffusion language models: Limits today, potential tomorrow. *arXiv preprint arXiv:2601.15593*, 2026. URL https://arxiv.org/abs/2601.15593.

Tamim Zoabi, Ameen Ali, Liran Ringel, and Lior Wolf. Mean-field parallel decoding for discrete diffusion language models. *arXiv preprint arXiv:2606.15805*, 2026. URL https://arxiv.org/abs/2606.15805.

A Proof of the Pairwise-Identical Separation

This appendix proves the compact part of Theorem 4.1, including the single fixed route family that works against all short value-oblivious schedules. We keep the construction and constants explicit because the lower bound must cover arbitrary anchor–payload interleavings. Appendix B uses the same family for the alphabet-four lifting. Throughout this appendix, write $P_N = P_N^{\mathrm{C}}$.

A.1 A universal family of hidden hypermatchings

Let N be a sufficiently large multiple of three. A perfect 3-hypermatching is a partition of $[N]$ into $N/3$ triples. Given a coloring $c : [N] \to [T]$, call a triple a terminal tie when its largest color occurs at least twice. Let $Y_H(c)$ count terminal ties in H.

Lemma A.1 (Terminal-tie density)**.** If $T \le N/64$ and H is a uniformly random perfect 3-hypermatching, then

$$\Pr\{Y_H(c) \ge 5\} > 1/9$$

for every coloring $c : [N] \to [T]$.

Proof. Write $x_t = |c^{-1}(t)|/N$ and $F_t = \sum_{s \le t} x_s$. For three i.i.d. vertices, the probability that the maximum color is repeated is

$$p_{\mathrm{iid}} = \sum_t \big(x_t^3 + 3x_t^2 F_{t-1}\big) \ge \sum_t x_t^2 F_t.$$

Cauchy–Schwarz and the right Riemann sum for $\sqrt{u}$ give

$$\sum_t x_t^2 F_t \ge \frac{1}{T}\left(\sum_t x_t \sqrt{F_t}\right)^2 \ge \frac{1}{T}\left(\int_0^1 \sqrt{u}\, du\right)^2 = \frac{4}{9T}.$$

Conditioning the draws to be distinct changes any event probability by at most $3/N$. Thus a uniformly random triple is a terminal tie with probability

$$p_{\mathrm{wor}} \ge \frac{4}{9T} - \frac{3}{N} \ge \frac{229}{9N}.$$

Generate H by a uniform permutation followed by consecutive blocks of three. A fixed triple is a block with probability $p_1 = 2/((N-1)(N-2))$; two fixed disjoint triples are both blocks with probability

$$p_2 = \frac{2}{(N-1)(N-2)} \frac{2}{(N-4)(N-5)}.$$

Writing $Y = Y_H(c)$, we have

$$\mu := \mathbb{E}Y = \frac{N}{3} p_{\mathrm{wor}} \ge \frac{229}{27}, \qquad \mathbb{E}Y^2 \le \mu + 2\mu^2,$$

where the second inequality overcounts all pairs as disjoint and uses $p_2/p_1^2 \le 2$ for $N \ge 12$. Paley–Zygmund at threshold $\mu/2$ yields

$$\Pr\{Y \ge \mu/2\} \ge \frac{\mu}{4(1+2\mu)} \ge \frac{229}{1940} > \frac{1}{9}.$$

Since $\mu/2 > 4$, this event implies $Y \ge 5$. □

Lemma A.2 (One family works for all short schedules)**.** There are a power of two $q = \Theta(N \log N)$ and a fixed multiset $(H_z)_{z \in \mathbb{F}_q}$ of perfect 3-hypermatchings such that, for every coloring $c : [N] \to [T]$ with $T \le N/64$,

$$\big|\{z : Y_{H_z}(c) \ge 5\}\big| \ge q/18.$$

Proof. Let $m = \lceil 144N \log(N+1) \rceil$ and choose $q = 2^{\lceil \log_2(m+1) \rceil}$, so $m < q \le 2(m+1)$. Draw q hypermatchings independently. For a fixed coloring, Lemma A.1 and a multiplicative Chernoff bound give

$$\Pr\left\{ \left| \{z : Y_{H_z}(c) \ge 5\} \right| < q/18 \right\} \le e^{-q/72}.$$

All colorings with at most $N/64$ labels are overcounted by the N^N maps $[N] \to [N]$. Since $q > 144N \log(N+1)$, $N^N e^{-q/72} < 1$. A union bound therefore leaves a positive-probability choice satisfying the claim simultaneously for all such colorings. Repeated hypermatchings cause no problem, so the resulting object may be a multiset. □

Fix this family for the rest of the proof. Because $q > N + 3$, the distinct field points used in Section 4.2 exist.

A.2 Pairwise identity and the exact three-round policy

For distinct $r, s \in \mathbb{F}_q$, the two rows $(1, r, r^2)$ and $(1, s, s^2)$ are linearly independent. Hence $(f(r), f(s))$ is uniform on $\mathbb{F}_q^2$. Given any route, one or two bits of an even-parity triple are independent fair bits; bits in different triples are independent, and these one- and two-bit laws do not depend on the route. The bits are also independent of U, V. These facts separately cover anchor–anchor, anchor–payload, and payload–payload pairs and prove that every two complete coordinates are independent. The coordinatewise product baseline therefore has exactly the same unconditional one- and two-coordinate tables.

The three anchor evaluations form an invertible Vandermonde transform of (U, V, Z), so they are mutually independent and can be sampled together in one exact product round. Once observed, they reveal Z, every scalar tag, and the concrete hypermatching H_Z. In a second round, reveal two fixed vertices of every triple. These bits are mutually independent across triples and pairwise independent within each triple. In the third round, each final bit is a point mass determined by parity. The chain law is exactly P_N. Conversely, a one-round decoder outputs the coordinatewise product baseline, which is not P_N: four scalar evaluations lie on one quadratic under P_N, but only with probability $1/q$ under the baseline.

A.3 The sharp two-round barrier

We first prove the compact two-round barrier. Condition on the policy's independent random tape and write B for its first batch, $b = |B|$. On every extendible incomplete branch, the second and final batch must contain the entire complement of B.

If $b \ge 4$, the product transition draws b independent uniform scalar evaluations. Exactly q^3 of their q^b tuples lie on one quadratic, so target-support survival is at most $q^{3-b} \le 1/q$.

Suppose $b \le 2$ and condition on a first history h. The observed scalar evaluations leave a uniform affine fiber of dimension $3 - b$ for (U, V, Z). Any one or two observed payload bits are jointly fair with a law independent of the route, so they neither shrink nor reweight this fiber. Any unrevealed evaluation is a nonconstant linear functional on it, because any set of at most three distinct Vandermonde rows is linearly independent. The final product round therefore draws the $N + 3 - b$ remaining scalar tags independently and uniformly. Exactly q^{3-b} such vectors extend the fiber, giving scalar-consistency probability $q^{3-b}/q^{N+3-b} = q^{-N}$.

Finally let $b = 3$ and consider an extendible first history. Its three scalar evaluations determine (U, V, Z) and hence H_Z. The first batch contains at most three payload bits. When $N \ge 6$, the route has $N/3 \ge 2$ triples; leaving at most one hidden bit in every triple would require at least $2N/3 \ge 4$ revealed payload bits. Thus some triple retains two or three hidden bits. Their singleton conditionals are fair, but their joint support obeys one parity equation, so the final independent product draw satisfies it with probability $1/2$. Nonextendible first histories contribute zero support mass.

All three cases are bounded by $1/2$. Averaging over the random tape gives $Q(\mathcal{V}_N) \le 1/2$, and the event characterization of TV gives $\mathrm{TV}(P_N, Q) \ge 1/2$. Together with the exact three-round policy, this proves that the exact history-aware depth of P_N is three for $N \ge 6$.

The threshold is genuine. Write $X_i = (T_i, C_i)$ for a compact payload coordinate. For $N = 3$, first reveal $\{Y_1, X_1, X_2\}$: their three scalar evaluations are jointly uniform and the two bits are

independent fair bits. They determine (U, V, Z) and the unique route triple, so the complementary batch $\{Y_2, Y_3, X_3\}$ is point-mass conditional and completes exactly. Hence the exact history-aware depth of P_3 is two.

A.4 The scalar crossing round

Fix a deterministic value-oblivious policy. Since reveals are irreversible, it induces a fixed ordered partition $\mathcal{B} = (B_1, \dots, B_T)$ of all coordinates. Every complete coordinate contains a scalar evaluation at a distinct field point. Let s_* be the first round whose cumulative number of revealed coordinates is at least three, let M_{s_*} be that number, and put $\kappa = M_{s_*} - 3$.

Before s_*, at most two scalar evaluations and at most two payload bits have been observed. The scalar observations restrict (U, V, Z) to a uniform affine fiber of dimension at least one. The observed bits do not shrink this fiber because every one- or two-bit parity marginal is uniform and route-independent. Any new evaluation functional is nonconstant on the fiber: otherwise at most three distinct Vandermonde rows would be linearly dependent. Every new scalar singleton is consequently uniform.

Suppose $d < 3$ coordinates precede the crossing batch and the batch contains m coordinates. The product transition draws m independent uniform scalars. Exactly q^{3-d} of their q^m tuples extend the current fiber to a common quadratic. The scalar-consistency probability is therefore $q^{3-d-m} = q^{-\kappa}$. Full target support is a subset of this event, so its probability is at most $q^{-\kappa}$. Notice that this is a crossing-round support bound; no assumption about how an invalid prefix is later completed is needed.

A.5 Parity loss after rank three

Now assume $\kappa = 0$. The first three scalar evaluations are jointly uniform and identify a uniformly random $\Theta = (U, V, Z)$. Color each payload site by its reveal round. For a triple whose latest color is unique, the last bit is sampled after the other two and is deterministic, so the triple survives with probability one. If the latest color is repeated, the last batch contains two or three fair singleton draws constrained by one unseen parity equation, and survival is exactly $1/2$. These factors are independent across route triples. Conditional on $Z = z$, support survival is thus $2^{-Y_{H_z}(c)}$, including when the scalar crossing batch itself contains payload coordinates. Averaging the uniform route before parity acceptance gives

$$Q_{\mathcal{B}}(\operatorname{supp} P_N) = q^{-1} \sum_{z \in \mathbb{F}_q} 2^{-Y_{H_z}(c)}.$$

If $T \le N/64$, Lemma A.2 makes at least $q/18$ routes satisfy $Y_{H_z}(c) \ge 5$. Hence

$$Q_{\mathcal{B}}(\operatorname{supp} P_N) \le 1 - \frac{1}{18}(1 - 2^{-5}) = \frac{545}{576}.$$

When $\kappa \ge 1$, the bound $1/q$ is smaller for all sufficiently large N. A randomized value-oblivious policy becomes a deterministic ordered partition after conditioning on its independent random tape, so the same uniform bound survives averaging.

Finally, P_N assigns mass one to its valid support. The event characterization of total variation gives, for every such short policy law Q,

$$\mathrm{TV}(P_N, Q) \ge 1 - Q(\operatorname{supp} P_N) \ge \frac{31}{576} > \frac{1}{20}.$$

This proves the randomized value-oblivious lower bound and completes Theorem 4.1.

A.6 Action-cover information

We prove Theorem 4.4. For a coloring $c : [N] \to [T]$, write

$$\mathcal{L}(c) := \{z \in \mathbb{F}_q : Y_{H_z}(c) = 0\}, \qquad L_T := \max_c |\mathcal{L}(c)|.$$

Lemma A.3 (Exact transcripts have no terminal tie)**.** If a possibly randomized full-history policy samples P_N exactly and C is its actual payload reveal-round transcript, then $Z \in \mathcal{L}(\mathsf{C})$ almost surely. Consequently

$$I(Z;\mathsf{C}) \geq \log q - \log L_T.$$

Proof. At every positive-probability history, exactness requires the selected product batch to remain in the extendible support almost surely. Suppose scalar rank is $d < 3$ and the next batch contains m complete coordinates. Its scalar projections are independent uniform evaluations. If $d + m > 3$, only a $q^{3-d-m} < 1$ fraction lie on one quadratic, contradicting support safety. Thus every positive path crosses scalar rank three with exactly three cumulative coordinates. Before that crossing there are at most two payload bits, whose law is route-independent; the crossing evaluations therefore determine the target's uniform route without prior reweighting.

A terminal tie cannot finish before the crossing because that would require all three bits of a route triple among at most two coordinates. At the crossing it can occur only when all three triple bits are drawn together, or when one was observed and the other two are drawn together. In either case the product bits satisfy the newly identified parity equation with probability $1/2$. After the crossing the route is already known, and at the maximum round of any terminal-tie triple at most one bit was previously observed while at least two fair singleton bits are drawn independently; again parity succeeds with probability $1/2$. Hence an exact policy has $Y_{H_Z}(\mathsf{C}) = 0$ on every positive-probability path.

Conditioned on $\mathsf{C} = c$, the support of Z is contained in $\mathcal{L}(c)$, so $H(Z \mid \mathsf{C} = c) \leq \log L_T$. Averaging and using $H(Z) = \log q$ proves the information bound. This positive-history argument already includes external policy randomness and does not condition on individual zero-probability random tapes. □

Lemma A.4 (A random route is rarely safe for a fixed coloring)**.** If H is a uniform perfect 3-hypermatching and $2 \leq T \leq N/64$, then every $c : [N] \to [T]$ satisfies

$$\Pr\{Y_H(c) = 0\} \leq \exp\left(-\frac{N}{21T}\right).$$

Proof. Generate H from consecutive triples of a uniform random permutation and expose the first $K = \lfloor N/6 \rfloor$ triples. At every such step at least $R \geq N/2$ vertices remain. For their current color proportions x_t and $F_t = \sum_{s \leq t} x_s$, three i.i.d. vertices have a repeated maximum with probability

$$\sum_t (x_t^3 + 3x_t^2 F_{t-1}) \geq \sum_t x_t^2 F_t \geq \frac{1}{T}\left(\sum_t x_t \sqrt{F_t}\right)^2 \geq \frac{4}{9T}.$$

Conditioning the draws to be distinct loses at most their collision probability $3/R \leq 6/N$. Since $T \leq N/64$, the next exposed triple is therefore a terminal tie with conditional probability at least $1/(3T)$, even after all earlier triples were safe. For sufficiently large N, $K \geq N/7$, and

$$\Pr\{Y_H(c) = 0\} \leq (1 - 1/(3T))^K \leq e^{-N/(21T)}.$$

□

Lemma A.5 (One family has uniformly short surviving lists)**.** The fixed family in Lemma A.2 can be chosen so that, for every $2 \leq T \leq N/64$ and every $c : [N] \to [T]$,

$$|\mathcal{L}(c)| \leq \Lambda_T := \left\lceil \max\left\{168T \log T, eq \exp\left(-\frac{N}{28T}\right)\right\}\right\rceil.$$

Proof. Sample the q hypermatchings independently. For fixed T, c, Lemma A.4 makes $|\mathcal{L}(c)|$ stochastically dominated by $\mathrm{Binomial}(q, p_T)$ with $p_T = e^{-N/(21T)}$. Put $a = N/(21T)$ and $\ell = \Lambda_T$. The two terms in Λ_T ensure

$$\log(eq/\ell) \leq 3a/4, \qquad a\ell/4 \geq 2N \log T.$$

The union-over-subsets binomial tail therefore gives

$$\Pr\{|\mathcal{L}(c)| \geq \ell\} \leq \binom{q}{\ell} p_T^{\ell} \leq e^{-2N \log T}.$$

There are at most T^N colorings, so the failure probability for this T is at most T^{-N}. Summing over T gives at most

$$\sum_{T=2}^{N/64} T^{-N} < 2^{-N} + \frac{2^{1-N}}{N-1}.$$

The failure probability of the simultaneous $Y_z \geq 5$ certificate in Lemma A.2 is at most $N^N e^{-q/72} < N^{-N}$ under $q > 144N \log(N+1)$. Their sum is below one for sufficiently large N, so one sampled family has both properties. □

Lemmas A.3 and A.5 give $I(Z; \mathsf{C}) \geq \log(q/\Lambda_T)$; the displayed simplified bound in Theorem 4.4 follows from $\lceil x \rceil \leq 2x$ here.

For the converse direction, take any k routes and form the graph joining two payload sites whenever they share a triple in one of those routes. Its maximum degree is at most $2k$, so a greedy proper coloring uses at most $2k+1$ colors. Every triple of every selected route then receives three distinct colors and has a unique latest vertex. Partition the q routes into groups of size at most $k = \lfloor (T-2)/2 \rfloor$. Reveal the anchors first, use the proper coloring assigned to the recovered route's group, and reveal one color class per round. Scalar tags are then point masses; in the true route, every color class contains at most one bit from each triple, and each latest bit is fixed by parity. The resulting policy is full-law exact in at most $1 + (2k+1) \leq T$ rounds, while its payload action transcript depends only on the route group. Therefore

$$I(Z; \mathsf{C}) \leq \log \left\lceil \frac{q}{k} \right\rceil,$$

which proves the upper bound in Theorem 4.4 and completes the proof.

B Proof of the Fixed-Alphabet Lifting

This appendix proves the alphabet-four part of Theorem 4.1. Put $q = 2^k$, $K = k+1$, and view $\mathbb{F}_q$ as a k-dimensional vector space over $\mathbb{F}_2$. Let $\rho_1, \ldots, \rho_k$ be a dual basis and $\rho_K = \rho_1 + \cdots + \rho_k$. Any $K - 1 = k$ members of this frame span the full dual space; for $k \geq 2$, any two are distinct nonzero forms and hence linearly independent.

B.1 Coordinate law and pairwise factorization

Sample $f(t) = U + Vt + Zt^2$ uniformly through $(U, V, Z) \in \mathbb{F}_q^3$. Anchor coordinate $A_{a,b} = \rho_b(f(\xi_a))$, for $a \in [3]$ and $b \in [k]$, is binary. Payload coordinate

$$W_{i,b} = (D_{i,b}, C_{i,b}), \qquad D_{i,b} = \rho_b(f(\lambda_i)), \qquad b \in [K],$$

is four-valued. Conditional on $Z = z$, every edge $e \in H_z$ carries an independent share vector, uniform subject only to

$$\bigoplus_{i \in e} \bigoplus_{b \in [K]} C_{i,b} = 0.$$

The three field evaluations at the anchor points are jointly uniform, so their $3k$ basis bits are mutually independent. Conditional on $Z = z$, two evaluations at distinct field points are jointly uniform: after fixing Z, the two equations have full rank in (U, V). Two frame forms of one uniform field element are independent because they are linearly independent. Any one or two shares have the uniform law under a parity constraint on $3K \geq 3$ bits, and this law is route-independent. Conditioning on Z therefore shows that two payload coordinates are independent uniform elements of $\mathbb{F}_2^2$, including within one group, and that every anchor–payload pair factors. This proves Lemma 4.3.

Round one of the upper policy reveals all anchors. They recover the three field evaluations and hence (U, V, Z). In round two, for each edge of the known H_Z, reveal any $3K - 1$ of its coordinates. Their tags are point masses and their shares are jointly independent fair bits. Round three reveals the remaining coordinate, whose tag and share are point masses. The three kernels multiply to the target law.

B.2 COMPUTATIONAL REALIZATION

Proposition B.1 (Polynomial realizability)**.** The fixed-alphabet target has an $O(L^2)$-bit description. Randomized polynomial preprocessing produces a table with all required universal properties with probability $1 - e^{-\Omega(N)}$. Given a good fixed table, the exact three-round sampler uses $O(L \operatorname{polylog} L)$ online bit time and polynomial space, and arbitrary-history exact singletons are polynomial-time computable.

Proof. Represent each H_z by a permutation of the N payload labels, grouped into consecutive triples. Since $q = \Theta(N \log N)$, the indexed table uses

$$O(qN \log N) = O(N^2 \log^2 N) = O(L^2)$$

bits. Sampling all q independent matchings takes $O(qN)$ word operations. The union bounds in Lemmas A.2 and A.5 make the bad-table probability at most

$$N^N e^{-q/72} + 2^{-N} + \frac{2^{1-N}}{N-1} = e^{-\Omega(N)}.$$

Thus randomized polynomial preprocessing produces a table with the required all-colorings properties with probability $1 - e^{-\Omega(N)}$.

After round one, a fixed 3×3 inverse Vandermonde solve recovers (U, V, Z) in $\operatorname{poly}(k)$ bit operations. The resulting Z indexes one table row. Scanning its triples, evaluating all payload tags with schoolbook field arithmetic, sampling the independent shares, and maintaining one parity accumulator per block costs $O(Nk^2 + L) = O(L \log L)$ bit operations, apart from ordinary batch-list bookkeeping. The output and online working storage are polynomial, with the $O(L^2)$-bit table read-only.

Arbitrary-history exact singletons are also polynomial-time computable. Enumerate the q^3 candidates (U, V, Z), reject anchor/tag inconsistencies, and evaluate the observed-share likelihood block by block under H_Z. Proper block marginals are uniform and completed blocks require one parity check, so exact weights have $O(L)$-bit dyadic representations. A conservative implementation uses $\widetilde{O}(q^3 L^2) = L^{O(1)}$ bit operations per query, establishing the claimed polynomial realizability. □

B.3 THE TWO-UNSEEN GUARD

Associate every revealed anchor or payload coordinate with its binary scalar tag form on the $3k$-bit representation of (U, V, Z). For a set S of revealed coordinates, let $\mathcal{S}(S)$ be the span of these forms.

Lemma B.2 (Two-unseen guard)**.** If $\dim \mathcal{S}(S) < 3k$, then every expanded block $e \times [K]$, for any triple e of payload groups, contains at least two coordinates outside S.

Proof. If at most one coordinate were missing, each of the three groups would expose at least $K - 1 = k$ frame forms. These recover all three field evaluations. The three-point Vandermonde map then recovers (U, V, Z), forcing $\dim \mathcal{S}(S) = 3k$, a contradiction. □

At a target-extendible history below full scalar rank, the observed shares in every candidate parity block form a proper marginal. Each observed share contributes one likelihood factor $1/2$, independently of the candidate route and scalar parameter. The scalar posterior remains uniform on its affine fiber. After adding any one hidden coordinate, Lemma B.2 leaves another share hidden in every candidate block; consequently

$$P(D_x, C_x \mid h) = P(D_x \mid h) \otimes \operatorname{Ber}(1/2) \tag{8}$$

for every unrevealed payload coordinate x. This is the crossing property that fails for the unguarded $K = k$ split.

B.4 SCALAR NULLITY AND ROUTE CROSSING

Fix a deterministic value-oblivious coordinate schedule and a positive, extendible pre-crossing history. Let $\mathcal{S}$ be the current scalar span. In the next batch, let m count tag forms not individually determined

by $\mathcal{S}$, and let Δ be their collective rank increment. By (8), the product transition draws their m scalar bits independently and uniformly. Exactly 2^{Δ} of the 2^m strings extend the current affine fiber. Hence

$$\Pr\{\text{scalar prefix remains extendible} \mid h\} = 2^{-(m-\Delta)}.$$

If any batch has $m - \Delta \geq 1$, irreversibility gives target-support mass at most $1/2$.

It remains to consider schedules with zero nullity until the first round $s_\star$ at which scalar rank reaches $3k$. Successive novel forms then expose one uniform (U, V, Z), while all pre-crossing shares remain independent fair bits and do not reweight its affine fiber. Equation (8) also makes the crossing shares independent fair bits after conditioning on their consistent tag outcomes.

For the arbitrary coordinate-level schedule, define the group color

$$g(i) = \max_{b \in [K]} \{\text{round in which } W_{i,b} \text{ is revealed}\}.$$

No group-atomic condition is imposed. If a route edge has a repeated maximum under g, at least two actual coordinates occur in that edge's latest round. A block cannot finish before $s_\star$ by the guard. If it finishes at the crossing, it starts that round with at least two hidden coordinates, whose product shares satisfy the parity equation with probability $1/2$. If it finishes later, the known route leaves at least two fair final singletons under one unseen parity equation, again giving $1/2$. Different route edges use disjoint shares, so these factors multiply.

Let $Y_z(g)$ count group-level terminal ties of H_z. Additional coordinate-level failures can only decrease support mass, and therefore

$$Q(\operatorname{supp} P_N^{\mathrm{FA}}) \leq q^{-1} \sum_{z \in \mathbb{F}_q} 2^{-Y_z(g)}. \tag{9}$$

Lemma A.2 applies to this arbitrary group-max coloring. For $T \leq N/64$, at least $q/18$ routes have $Y_z(g) \geq 5$, so (9) is at most $545/576$. Together with the nullity branch, this proves the deterministic schedule bound. Conditioning on an independent random tape and averaging proves it for randomized value-oblivious policies. Since the target support has probability one under the target, the event characterization of TV yields

$$\mathrm{TV}(P_N^{\mathrm{FA}}, Q) \geq 31/576 > 1/20.$$

For an at-most-two-round full-history policy, condition on its tape. The first batch is fixed at the empty history, and every terminating second batch must be its complement. It is therefore a two-color value-oblivious schedule and cannot be exact for sufficiently large N. The three-round chain law proves exact full-history depth three. Finally, $k = \Theta(\log N)$ and (3) give $N = \Theta(L/\log L)$, completing the fixed-alphabet part of Theorem 4.1.

B.5 Action-cover information for actual coordinates

We prove the fixed-alphabet lower statement in Theorem 4.4. Consider any possibly randomized exact full-history policy using at most T nonempty rounds. Let $\mathsf{A}_{i,b}$ be the reveal round of actual payload coordinate $W_{i,b}$, and define

$$\mathsf{G}_i := \max_{b \in [K]} \mathsf{A}_{i,b}.$$

As in Appendix A.6, put $\mathcal{L}(g) = \{z : Y_{H_z}(g) = 0\}$.

Lemma B.3 (Exact actual transcripts are route-safe)**.** Every such policy satisfies

$$Z \in \mathcal{L}(\mathsf{G}) \quad \text{almost surely.}$$

Proof. At every positive-probability history, each batch selected with positive conditional probability must be target-support safe. Otherwise the product of its singleton supports assigns positive mass to a nonextendible outcome, which irreversibility cannot repair.

Below full scalar rank, Lemma B.2 leaves at least two actual coordinates hidden in every candidate route block, and (8) makes each newly drawn share fair independently of its atomic tag. A batch with positive scalar nullity therefore has a nonextendible tag outcome of positive probability and

cannot occur on an exact positive branch. Thus every exact path reaches full scalar rank without overcommitment, and no route block can finish strictly before that crossing.

Suppose a true-route edge has a repeated maximum under G. Its block-completion batch contains at least two of that block's last actual coordinates. If completion occurs at scalar crossing, their shares remain independent fair bits after conditioning on all same-batch tags by (8); after crossing, the route is already known and the same fair-singleton statement follows directly from the unfinished parity block. In either case, flip one last fair share while keeping all tags and other batch values fixed. This outcome has positive product-transition probability but violates the completed true-route parity constraint, so it is nonextendible. Support safety is contradicted. Hence every true-route edge has a unique group maximum on every positive exact path.

The argument conditions on positive history–action events rather than on individual realizations of a possibly nonatomic random tape. It therefore covers randomized, value-dependent actions and history-dependent early stopping. □

The exact output marginal makes Z uniform, so $H(Z) = \log q$. Since G is a deterministic function of A, data processing and Lemma B.3 give

$$I(Z;\mathsf{A}) \geq I(Z;\mathsf{G}),$$
$$H(Z \mid \mathsf{G}) \leq \mathbb{E}\log|\mathcal{L}(\mathsf{G})| \leq \log \Lambda_T,$$

where Λ_T is the uniform list bound in Lemma A.5. Hence

$$I(Z;\mathsf{A}) \geq I(Z;\mathsf{G}) \geq \log(q/\Lambda_T),$$

and $\lceil x \rceil \leq 2x$ yields the displayed fixed-alphabet bound in Theorem 4.4. This is a statement about actual coordinate actions; G is their deterministic proof projection.

C Approximate Action-Cover Information

We prove Theorem 4.5. Fix the route table and write $\mathcal{V} = \operatorname{supp} P$. At an extendible history h, let $J(y \mid h, B)$ be the target joint conditional on batch B, and $K(y \mid h, B)$ the product of its exact singleton conditionals. A batch has an *actual terminal tie* on a route block if it reveals all remaining shares of that block and at least two were hidden when the batch began. A repeated maximum in the group cover implies such a tie; a tie can also occur within one last-finishing group.

Domination on extendible outcomes. For every y extending h to a target-consistent history,

$$K(y \mid h, B) \leq J(y \mid h, B). \tag{10}$$

To see this for the lifting, let m count individually undetermined tag bits, Δ be their rank increment, and $\nu = m - \Delta$. If scalar rank remains below full rank after the batch, the posterior is uniform on an affine fiber and the guard leaves every candidate block incomplete. Thus consistent tags have product-to-joint probability ratio $2^{-\nu}$, and proper share marginals contribute ratio one. If the batch reaches full rank, its tags determine the route. By Lemma B.2, each block completed in that batch begins with at least two hidden shares. For each of these t blocks, the target parity constraint contributes an additional factor $1/2$ to the ratio, giving $K/J = 2^{-\nu-t}$. After full rank, tags are fixed; a sole remaining share is deterministic, while each actual terminal tie contributes $1/2$. The ratio is then 2^{-t}.

For the compact target the corresponding ratio is $q^{-\nu}2^{-t}$, where ν counts new undetermined evaluations minus their field-rank increment, and is zero after full rank. Before full rank there are at most two observed coordinates, so the tag posterior is again uniform on its affine fiber. The atomic singleton (T_i, C_i) retains its conditional share law: if two shares of a block are already observed, its final tag can determine the route and its final share is then deterministic. If at least two shares remain, each is fair conditional on its own tag, and completing the block in one batch contributes $1/2$. Uncompleted blocks have proper parity marginals. These factors establish (10) for both realizations.

A coupling with an unchanged target marginal. Draw $X^\star \sim P$ independently of the policy tape R. While the actual history remains extendible, keep its revealed values equal to $X^\star$. Given

this history and the selected batch B, accept $Y^\star = X^\star_B$ as the actual batch with probability $K(Y^\star \mid h, B)/J(Y^\star \mid h, B)$, using auxiliary randomness unavailable to the scheduler. On rejection, draw the actual batch from K restricted and normalized to nonextendible outcomes, and subsequently use the policy's prescribed completion rule. Keep $X^\star$ unchanged.

Conditional on an extendible actual history, the unobserved part of $X^\star$ retains the target posterior: all earlier acceptance factors depend only on the revealed values and selected batches. Hence each legal outcome has actual mass $J(y)K(y)/J(y) = K(y)$; the remaining mass is exactly the mass that K assigns to nonextendible outcomes. Induction gives actual output law Q. The presampled $X^\star$ still has law P, so its route $Z^\star$ is uniform and independent of R. On every valid output no rejection occurred, and the realized route equals $Z^\star$.

First threats and list error. A threat is a batch with positive scalar nullity or an actual terminal tie. Positive nullity has scalar consistency probability $q^{-\nu} \le 1/2$ in the compact target or $2^{-\nu} \le 1/2$ in the lifting. Otherwise, factor the batch kernel by exposing tags first for the analysis. If these tags first determine the route, test for an actual tie before exposing shares; if the route is already known, test at the batch's start. Conditional on a tie, the unseen parity survives with probability at most $1/2$. This factorization preserves the atomic singleton kernel and the selected batch. The first-threat events are disjoint, and irreversibility makes every failed extension final. Thus, for the event $\mathcal{U}$ that a threat is reached,

$$\Pr(\mathcal{U}) \le 2Q(\mathcal{V}^c) \le 2\,\mathrm{TV}(P, Q) = 2d. \tag{11}$$

Without a threat, all outcomes are extendible and agree with $X^\star$, and no true route block has a repeated group maximum. Consequently,

$$\Pr\{Z^\star \notin \mathcal{L}(\mathsf{C})\} \le 2d, \qquad \Pr\{Z^\star \notin \mathcal{L}(\mathsf{G})\} \le 2d. \tag{12}$$

Here C is the compact cover and G the lifting's group cover; their list-error events are both contained in $\mathcal{U}$.

Write p for the relevant list-error probability in (12) and recall $\bar{\Lambda}_T$ from Section 4.5. The list form of Fano's inequality and Lemma A.5 give

$$H(Z^\star \mid \mathrm{cover}) \le h_2(p) + (1-p)\log \bar{\Lambda}_T + p \log q.$$

Since $H(Z^\star) = \log q$,

$$I(Z^\star; \mathrm{cover}) \ge (1-p)G_T - h_2(p).$$

For $d \le 1/4$, the right side is decreasing on $p \le 2d \le 1/2$, yielding the two bounds in Theorem 4.5. Because G is a deterministic function of A, data processing gives $I(Z^\star; \mathsf{A}) \ge I(Z^\star; \mathsf{G})$. Finally $h_2(2d) \le \log 2$; rearranging $B \ge (1-2d)G_T - \log 2$ proves the explicit information–error frontier.

C.1 Bounded observation interfaces

Condition on the fixed route table $\mathcal{H} = (H_z)_z$ and use the coupling above, in which $Z^\star$ is independent of the scheduler tape R. Let O be the complete transcript of values, scores, oracle responses, and other route-correlated side information used for decisions. Causality makes the action transcript, and hence its cover, a deterministic function of $(\mathcal{H}, R, \mathsf{O})$. Data processing gives

$$I(Z^\star; \mathrm{cover}) \le I(Z^\star; \mathsf{O}, R \mid \mathcal{H}) = I(Z^\star; \mathsf{O} \mid \mathcal{H}, R) \le B.$$

Combining this with Theorem 4.5 proves Corollary 4.6.

For an explicit depth consequence, suppose $d \le \varepsilon < 1/4$ and put

$$D = \frac{B + h_2(2\varepsilon)}{1 - 2\varepsilon} + \log 2, \qquad x = \frac{qe^{-D}}{168}.$$

The list bound forces, up to integer rounding,

$$T \ge \min\left\{ \frac{N}{64}, \frac{x}{W_0(x)}, \frac{N}{28(D+1)} \right\},$$

where W_0 is the principal Lambert function; the first term covers $T > N/64$. Fixed B and $\varepsilon < 1/4$ make all three terms $\Theta(N)$. If instead $I(Z^\star; \mathsf{O}_t \mid \mathcal{H}, R, \mathsf{O}_{<t}) \le b_t$, the chain rule gives $B \le \sum_t b_t$. For $b_t \le b = O(1)$, the action-information bound forces $T = \Omega(\log N)$.

Finite or quantized transcripts yield such a budget directly.

D SUPPORT-PRESERVING ROBUSTNESS OF THE GUARDED LIFTING

Write $P = P_N^{\mathrm{FA}}$, $K = k + 1$, and

$$L = 3k + NK, \qquad M_{\mathrm{FA}} = 3k + \frac{N}{3}(3K - 1) = L - \frac{N}{3}.$$

At every target-extendible history h, let $p_x(\cdot \mid h)$ be the exact singleton and let $\widehat{p}_x(\cdot \mid h)$ satisfy

$$\sup_{h,x} \mathrm{TV}(p_x(\cdot \mid h), \widehat{p}_x(\cdot \mid h)) \le \delta < \frac{1}{6}. \tag{13}$$

The adaptive upper bound additionally assumes equal local supports. Define

$$\psi(\delta) = -\frac{1}{2}\log(1 - 4\delta^2), \qquad \theta(\delta) = \frac{4\delta}{1 - 2\delta}, \qquad s(\delta) = \frac{1 + \theta(\delta)^2}{2},$$

and

$$\beta_k(\delta) = \min\left\{1, \sqrt{\frac{3k\psi(\delta)}{2}}\right\}.$$

D.1 THE THREE-ROUND UPPER BOUND

Proposition D.1 (Approximate three-round chain)**.** Under equal local supports, the three-round policy $A^\star$ satisfies

$$\mathrm{KL}(P \| \widehat{Q}_{A^\star}) \le M_{\mathrm{FA}}\psi(\delta), \qquad \mathrm{TV}(P, \widehat{Q}_{A^\star}) \le \sqrt{\frac{M_{\mathrm{FA}}\psi(\delta)}{2}}.$$

Proof. Along every P-positive path, $A^\star$'s batches are products of exact current singletons. The chain rule therefore expresses the global forward KL as the target expectation of the sum of local singleton KL terms. Exactly M_{FA} kernels are nondegenerate and uniform on two atoms; equal support and (13) bound each forward KL by $-\frac{1}{2}\log(1 - 4\delta^2) = \psi(\delta)$. The deterministic final shares contribute zero. Pinsker gives the TV statement. □

D.2 A SHADOW CONTINUATION FOR VALUE-OBLIVIOUS SCHEDULES

The lower bound must remain meaningful after an approximate rollout first leaves target support. The following device prevents invalid prefixes from being treated as if they still had exact target conditionals.

Lemma D.2 (Shadow-extendible continuation)**.** Fix a deterministic value-oblivious schedule with zero scalar nullity through full rank. Couple its approximate rollout to a killed process that agrees with it until the first nonextendible prefix and, after killing, completes all undrawn novel scalar-tag bits as independent fair bits. If μ_Z is the route induced by the resulting $3k$ novel bits, then

$$\mathrm{TV}(\mu_Z, \mathrm{Unif}(\mathbb{F}_q)) \le \beta_k(\delta).$$

The killed continuation leaves the actual valid-support mass unchanged.

Proof. At a live zero-nullity history, every novel exact tag bit is fair, so (13) places its approximate success probability in $[1/2 - \delta, 1/2 + \delta]$. At a dead history the shadow bit is exactly fair; marginalizing nuisance shares and the live/dead flag preserves the same interval. Hence the forward KL from $3k$ independent fair bits to the shadow-bit law is at most $3k\psi(\delta)$. The route is a function of those bits, so data processing and Pinsker prove the claim. Killing occurs only after the actual prefix is already invalid, and subsequent values cannot restore target support. □

D.3 THE ROBUST VALUE-OBLIVIOUS LOWER BOUND

Proposition D.3 (Approximate value-oblivious obstruction)**.** Every randomized value-oblivious policy using $T \le N/64$ rounds obeys

$$\mathrm{TV}(P, \widehat{Q}) \ge \left[\min\left\{\frac{1}{2} - 2\delta^2, \frac{1 - s(\delta)^5}{18} - \beta_k(\delta)\right\}\right]_+ .$$

Proof. Condition first on the policy tape. If a batch has positive scalar nullity, choose a nontrivial relation among $r \ge 2$ new tag bits. Each bit has Fourier bias at most 2δ, and product sampling makes the relation survive with probability at most $\frac{1}{2}\{1 + (2\delta)^r\} \le \frac{1}{2} + 2\delta^2$.

Otherwise apply Lemma D.2. At a terminal tie, at least two final coordinates of a true parity block are drawn independently. Conditional on their tags, their share biases have magnitude at most $\theta(\delta)$, so a tied block survives with conditional probability at most $s(\delta)$. Ordering block-completion events and using the tower property gives route-z survival at most $s(\delta)^{Y_z}$. The fixed family has at least $q/18$ routes with $Y_z \ge 5$, while Lemma D.2 changes route expectations by at most $\beta_k(\delta)$. Thus valid-support mass is at most

$$1 - \frac{1 - s(\delta)^5}{18} + \beta_k(\delta).$$

The two branches give the displayed minimum. For randomized policies, average invalid-support mass over the tape; no averaging of per-tape TV is needed. Finally, P has unit mass on its support. □

For $\delta_L = c/\sqrt{L}$, Proposition D.1 gives $\mathrm{TV}(P, \widehat{Q}_{A^\star}) \le c + o(1)$. Since $k = \Theta(\log L)$, $\beta_k(\delta_L) = o(1)$ and $s(\delta_L) = 1/2 + o(1)$; Proposition D.3 gives $31/576 - o(1)$. This proves Corollary 4.7.

Equal support separates two error regimes. If each of the $N/3$ deterministic last shares places probability δ on one impossible value, every local singleton is within TV δ, but the exact output TV is

$$1 - (1 - \delta)^{N/3}.$$

This example shows that local errors outside the support can accumulate with dimension.

E EXPERIMENTAL DETAILS

Figure 2 combines exact enumeration of sampling laws and tag constraints with route sampling at larger sizes. We describe the comparisons, statistical bounds, and perturbations below.

Exact enumeration. Exact arithmetic enumerates all 4,683 ordered interleavings at $(N, q) = (3, 7)$ and, for panel (a), 21,296 target states at $(N, q) = (6, 11)$. Panel (d) compares guarded and naive tag splits at $(k, N) = (3, 3), (4, 6), (5, 24)$. For each route block, we omit each tag bit in turn and test whether the remaining tag forms retain rank $3k$, which is necessary and sufficient to recover (U, V, Z). All 5,184 guarded cases retain full rank; all 4,296 matched cases of the naive split lose rank.

Schedules and route sampling. For each $N \in \{192, 384, 768, 1536\}$, we evaluate ten round budgets: the distinct values in $\{2, 4, 8\} \cup \{N/64, N/32, N/16, N/8, N/4, N/2, N\}$. At every budget, the six value-oblivious schedules are round-robin, contiguous equal-sized blocks, and four balanced random permutations fixed before route sampling. Each comparison uses 8,192 independently sampled uniform perfect 3-hypermatchings. A route is safe when every route triple has a unique last-finishing group.

Safe-route probability bound. Let $\widehat{p}$ be the observed safe-route fraction for one schedule, $R = 8192$ the number of sampled routes, $M = 240$ the total number of comparisons across sizes, budgets, and schedules, and $\alpha = 0.05$. We use the familywise Hoeffding upper bound, valid simultaneously with probability at least $1 - \alpha$,

$$p_{\mathrm{safe}} \le u := \min\left\{1, \widehat{p} + \sqrt{\frac{\log(M/\alpha)}{2R}}\right\}.$$

Thus $-\log u$ is a lower confidence bound on $-\log p_{\mathrm{safe}}$ for that fixed cover. Panel (e) uses the largest u among the six schedules at each size and budget, then shows the median and full range across sizes. This measures how much the cover restricts the surviving routes. At $T = N/64$, all six schedules have $\widehat{p} = 0$ at every tested size. The simultaneous upper bound is therefore 0.022745 for each comparison under the uniform route distribution, corresponding to $-\log(0.022745) = 3.783$ nats.

Table 1: Safe routes at $T = N/64$. "Safe" is the largest count among the six schedules.

N	k	actual L	T	routes	safe
192	18	3,702	3	8,192	0
384	19	7,737	6	8,192	0
768	20	16,188	12	8,192	0
1536	21	33,855	24	8,192	0

Sampled support loss. For a fixed cover c, let $Y_r(c)$ count terminal ties on sampled route r. Panel (b) evaluates the following support-loss expression, using the same sampled routes for both perturbation levels.

$$\widehat{b}_\delta(c) = \min\left\{\frac{1}{2}, 1 - \frac{1}{R}\sum_{r=1}^{R}\left(\frac{1+4\delta^2}{2}\right)^{Y_r(c)}\right\}.$$

Here δ bounds each independent share's probability bias from $1/2$. A terminal tie leaves at least two such shares, whose probability of satisfying parity is at most $(1+4\delta^2)/2$. The plot takes the smallest value over the six schedules at each size and budget, then plots the median and full range across sizes. Here $\delta = 0$ uses exact singleton probabilities and $\delta = 0.05$ perturbs probabilities within the support; the displayed curve uses the latter. Setting $\delta = 0.05$ changes the sampled frontier by less than 0.002; at $T = N/2$, the bound remains between 0.215 and 0.219 across sizes.

Perturbation models. Panel (c) uses the support-preserving bounds in Appendix D to plot the certified gap against the normalized error $\delta\sqrt{M}$, where M is the number of fair draws. Panel (f) moves probability ϵ to an impossible tag at each of the M_2 round-two payload draws, giving exact TV $1-(1-\epsilon)^{M_2}$. The resulting error grows with the number of draws, as in the leakage example following Corollary 4.7.

F Decoder-Interface Map

The main comparison holds the singleton-product transition fixed and changes whether batch choices may use generated values. Confidence, learned order, attention, dependence, support, and mean-field schedulers all use current-state signals to choose sites (Fox et al., 2026; Luxembourg et al., 2025; Zhong et al., 2026; Abu Ayoub et al., 2026; Qi et al., 2026; Ringel et al., 2026; Zoabi et al., 2026). They therefore belong on the history-aware side whenever those signals change the selected batch. Theorem 4.4 concerns the resulting batch choices and applies regardless of their internal representation.

WeDLM combines confidence-based streaming selection with topological reordering, causal attention, and prefix caching (Liu et al., 2025). Its selection rule uses current predictions; its attention and caching design addresses inference cost. Joint generation, remasking, and concurrent blocks alter the commit interface (Li et al., 2026; Chen et al., 2026b; Jin et al., 2026). This distinction keeps history information, transition power, and implementation cost separate.

The construction combines finite-field sharing, orthogonal arrays, hidden hypermatchings, and unique-maximum coloring (Bush, 1952; Shamir, 1979; Verbin & Yu, 2011; Cheilaris et al., 2013) into a fixed four-symbol target that supports the same-target policy separation under arbitrary interleavings.

G Finite Uniform-Injection Frontier

We prove Proposition 5.1 and Theorem 5.2. Consider a valid partial injection with r hidden sites. Exactly $m = q - n + r$ symbols remain unused. Fixing one hidden site to a particular unused symbol

leaves $(m-1)_{r-1}$ completions out of $(m)_r$, so every singleton conditional is uniform on the same m-symbol set. An independent batch of size b is therefore uniform on m^b ordered tuples, of which $(m)_b$ are collision-free. This proves (5).

For a count schedule $r_1 > \cdots > r_{T+1} = 0$, set $b_s = r_s - r_{s+1}$ and $m_s = q - n + r_s$. The chain rule gives

$$Q_{\mathbf{r}}(\mathcal{I}_{n,q}) = \prod_{s=1}^{T} \frac{(m_s)_{b_s}}{m_s^{b_s}}.$$

Since $m_{s+1} = m_s - b_s$, the numerators telescope: $\prod_s (m_s)_{b_s} = (q)_n$. This is the stated $S(\mathbf{r})$.

For the stronger full-law identity, fix a valid completion $x \in \mathcal{I}_{n,q}$ and condition on the policy's independent random tape. The complete values determine a unique site-selection path. Along it, the batch values prescribed by x have probability $m_s^{-b_s}$, independent of the selected identities and the random tape. Thus

$$Q_{\mathbf{r}}(x) = \prod_s m_s^{-b_s} = \frac{S(\mathbf{r})}{(q)_n} = S(\mathbf{r}) P_{n,q}(x).$$

The total deficit on the valid support and the total mass outside it both equal $1 - S(\mathbf{r})$, which gives the TV identity in (6); the likelihood ratio is constant on the target support, giving the KL identity.

For full-history policies, induct on the number t of rounds remaining. All valid histories with r hidden sites expose the same number $m = q - n + r$ of unused values. If the first batch has size $1 \le b \le r - t + 1$, it survives with probability $p(r, b)$; conditional on survival, every child has $r - b$ hidden sites and the same continuation value $V_{t-1}(r - b)$. Randomizing the action or conditioning it on realized values only forms a convex combination of these action values. Taking the largest one yields exactly (7), and choosing a fixed maximizing batch size at each state attains it.

It remains to identify TV for an arbitrary adaptive policy. For any valid x, its realized count path satisfies

$$Q(x) = \prod_s m_s^{-b_s} \le \frac{1}{\prod_s (m_s)_{b_s}} = \frac{1}{(q)_n} = P_{n,q}(x).$$

The inequality also holds after averaging over external randomness. Hence the absolute value can be removed on the valid support, and

$$\mathrm{TV}(P_{n,q}, Q) = \frac{1}{2}\big(1 - Q(\mathcal{I}_{n,q}) + Q(\mathcal{I}_{n,q}^{\mathsf{c}})\big) = 1 - Q(\mathcal{I}_{n,q}).$$

Combining this identity with the Bellman optimum proves Theorem 5.2.

H Fixed and Growing Permutation Rounds

Set $q = n$. It is convenient to minimize the negative log-validity. Define

$$C_1(r) = r \log r, \qquad C_t(r) = \min_{t-1 \le k \le r-1} \{(r-k)\log r + C_{t-1}(k)\}.$$

Then $M_{n,T} = n! e^{-C_T(n)}$.

H.1 Fixed-round regime

Let $u_0 = 0$ and $u_{j+1} = e^{u_j - 1}$. We claim, for every fixed t,

$$C_t(r) = r \log r - r u_{t-1} + O_t(r^{-1}). \tag{14}$$

The claim is exact for $t = 1$. Suppose it holds for $t - 1$, put $a = u_{t-2}$ and $v = e^{a-1} = u_{t-1}$, and write $k = rx$. Up to an $O_t(k^{-1})$ term, the part of the objective beyond $r \log r$ is

$$r\phi_a(x), \qquad \phi_a(x) = x \log x - ax.$$

The strictly convex function ϕ_a has unique minimizer v and value $-v$. Rounding rv supplies the upper bound in (14). Any sequence of exact integer minimizers must have $k/r \to v$; otherwise a convergent subsequence would attain a strictly larger limit of ϕ_a. The minimizers are then bounded

below by a positive multiple of r, making the induction error $O_t(r^{-1})$ and giving the matching lower bound. Stirling's formula now yields, for every fixed T,

$$M_{n,T} = \sqrt{2\pi n}\, e^{-n(1-u_{T-1})}\big(1 + O_T(n^{-1})\big). \tag{15}$$

Equivalently,

$$\log M_{n,T} = -n(1-u_{T-1}) + \tfrac{1}{2}\log(2\pi n) + O_T(n^{-1}).$$

H.2 Growing-round regime

For a batch of size b at remaining count r, define its collision hazard

$$h(r,b) := -\log\frac{(r)_b}{r^b} = \sum_{j=0}^{b-1} -\log(1-j/r).$$

The exact frontier says that $-\log M_{n,T}$ is the minimum of $\sum_s h(r_s,b_s)$ over T-batch schedules. Since $-\log(1-x) \ge x$,

$$h(r,b) \ge \frac{b(b-1)}{2r}.$$

Writing $c_s = b_s - 1$ and $E = n - T = \sum_s c_s$, weighted Cauchy–Schwarz gives

$$-\log M_{n,T} \ge \frac{E^2}{2nT} + \frac{E}{2n}. \tag{16}$$

For an upper bound, reverse the differences of $R_k = \max\{k, \lceil nk^2/T^2\rceil\}$, $k = 0,\dots,T$, into a legal schedule. The complementary inequality

$$h(r,b) \le \frac{b(b-1)}{2(r-b+1)}$$

shows that every nonsingleton stage costs at most $18n/T^2$, apart from a terminal stage bounded by $2n/T^2$. Hence

$$-\log M_{n,T} \le 18n/T. \tag{17}$$

If $T \le n/2$, then $E \ge n/2$, and (16)–(17) give $-\log M_{n,T} = \Theta(n/T)$.

Finally, when $E \le n/2$, use E initial size-two batches followed by $n-2E$ singletons. Its cost is at most

$$\sum_{j=0}^{E-1} \frac{1}{n-2j-1} \le \frac{E}{n-2E+1}.$$

Together with the elementary lower bound $-\log M_{n,T} \ge E/n$, this proves $M_{n,T} \to 1$ exactly when $E/n = (n-T)/n \to 0$, completing the growing-round statements in Section 5.2.

I Sharp Linear-Round Limit and Profile

We prove Theorem 5.3. Write $H_{n,T} = -\log M_{n,T}$. By the finite frontier,

$$H_{n,T} = \min \sum_s h(r_s,b_s), \qquad h(r,b) = \sum_{j=0}^{b-1} -\log(1-j/r), \tag{18}$$

where $r_{s+1} = r_s - b_s$, $r_1 = n$, $r_{T+1} = 0$, and all $b_s \ge 1$.

For $0 < \tau < 1$, let $K \ge 2$ be the unique integer satisfying

$$\frac{2}{K+1} < \tau \le \frac{2}{K},$$

and define

$$I(\tau) := \sum_{b=2}^{K-1} \frac{b-1}{2}\log\frac{b+1}{b-1} - \frac{K-1}{2}\log(K\tau - 1), \qquad \Phi(\tau) := e^{-I(\tau)}, \tag{19}$$

with $\Phi(1) = 1$.

I.1 Ordering and the continuum program

An exact optimum is front-loaded. Indeed, adjacent batches a, b beginning at r contribute $a \log r + b \log(r-a)$ to the equivalent denominator cost. Reversing them changes the contribution to $b \log r + a \log(r-b)$. The function $x \mapsto \log(1-x/r)/x$ is strictly decreasing, so $a > b$ is strictly better before b. Repeated swaps put every optimum in nonincreasing order.

Represent a schedule by its token control

$$\beta_n(u) = b_s \quad \text{for} \quad (r_s - b_s)/n < u \le r_s/n.$$

The intervals partition $(0, 1]$ and satisfy the exact round identity

$$\int_0^1 \frac{du}{\beta_n(u)} = T/n. \tag{20}$$

The linearized collision cost on such an interval is

$$\int_{(r-b)/n}^{r/n} \frac{b-1}{2u}\, du,$$

suggesting the continuum program

$$\inf_{\beta:\,(0,1]\to\mathbb{N}} \left\{ \int_0^1 \frac{\beta(u)-1}{2u}\, du : \int_0^1 \frac{du}{\beta(u)} = \tau \right\}. \tag{21}$$

Lemma I.1 (Discrete-to-continuum value limit)**.** If $T_n/n \to \tau \in (0,1)$, then H_{n,T_n} converges to the value of (21).

Proof. The elementary bound

$$h(r,b) \ge L(r,b) := \frac{b(b-1)}{2r}$$

and any bounded-cost comparison schedule imply, along a minimizing sequence, $b(b-1) \le O(n)$ for every batch. Hence the largest token interval has length $o(1)$. On every truncated domain $u \ge \delta > 0$, $b/r = o(1)$ uniformly, and

$$0 \le -\log(1-x) - x \le \frac{x^2}{2(1-x)}$$

makes the sum of $L(r,b)$ a Riemann lower approximation to the objective in (21). Introducing a multiplier $\mu > 0$ and using (20) yields the liminf bound

$$\liminf H_{n,T_n} \ge \sup_{\mu>0} \left\{ \int_0^1 a_\mu(u)\, du - \mu\tau \right\}, \quad a_\mu(u) := \min_{b\ge 1} \left\{ \frac{b-1}{2u} + \frac{\mu}{b} \right\}. \tag{22}$$

Near zero, $a_\mu(u) = \mu$, so removing the truncation is harmless.

For the converse, use the bounded piecewise-constant minimizer identified below. Its finitely many interval lengths specify real counts of each batch size. Round the counts for sizes $b \ge 3$, then use sizes two and one to enforce exactly $\sum b_s = n$ and $\sum 1 = T_n$. Their limiting intervals have positive length, so these corrected counts remain nonnegative. Arrange the batches in decreasing order. Every nonsingleton batch then occurs at $r = \Omega(n)$, and uniformly for bounded b,

$$h(r,b) = \frac{b(b-1)}{2r} + O(r^{-2}).$$

The $O(n)$ remainders sum to $o(1)$, while the leading terms converge to the continuum integral. This supplies the matching limsup and strong duality in (22). □

I.2 SOLVING THE SCALAR DUAL

For fixed μ, u, let $G_b(u) = (b-1)/(2u) + \mu/b$. Its consecutive difference is

$$G_{b+1}(u) - G_b(u) = \frac{1}{2u} - \frac{\mu}{b(b+1)}.$$

Thus a minimizing action b is characterized, up to endpoint ties, by

$$\frac{b(b-1)}{2\mu} \le u \le \frac{b(b+1)}{2\mu}. \tag{23}$$

Let K be the largest active batch size. The $b = 1$ interval has length $1/\mu$, each complete intermediate b-interval has length b/μ, and the last interval ends at one. Consequently,

$$\tau = \int_0^1 \frac{du}{\beta_\mu(u)} = \frac{1}{K} + \frac{K-1}{2\mu}.$$

This equation has the unique solution

$$\frac{2}{K+1} < \tau \le \frac{2}{K}, \qquad \mu = \frac{K(K-1)}{2(K\tau - 1)}.$$

Integrating $(b-1)/(2u)$ over the intervals in (23) gives

$$\sum_{b=2}^{K-1} \frac{b-1}{2} \log \frac{b+1}{b-1} - \frac{K-1}{2} \log(K\tau - 1) = I(\tau).$$

Lemma I.1 therefore implies $H_{n,T_n} \to I(\tau)$, and exponentiation proves $M_{n,T_n} \to \Phi(\tau)$. The endpoint $\tau = 1$ follows from the bounds in Appendix H because $n - T_n = o(n)$.

The lengths of the minimizing action intervals are exactly

$$\theta_1 = 1/\mu, \qquad \theta_b = b/\mu \ (2 \le b < K), \qquad \theta_K = 1 - \frac{K(K-1)}{2\mu} = 2 - K\tau.$$

It remains to show that every sequence of exact finite optima, including arbitrary tie choices, converges to these lengths.

I.3 WHY EVERY EXACT OPTIMIZER HAS THE SAME PROFILE

Fix any exact optimal schedules with $T_n/n \to \tau \in (0,1)$, and retain the notation β_n. Their costs are bounded by the value limit, so the largest batch is $O(\sqrt{n})$ and the largest token interval is $o(1)$. Define the nonnegative dual gap

$$\Delta_\mu(u, b) := \frac{b-1}{2u} + \frac{\mu}{b} - a_\mu(u).$$

The truncated Riemann argument from Lemma I.1, now using optimal-value convergence, gives for every $\eta > 0$,

$$\int_\eta^1 \Delta_\mu(u, \beta_n(u))\, du \longrightarrow 0. \tag{24}$$

On a compact subset of the singleton region $(0, 1/\mu)$, action one is the unique minimizer and every other integer action has a uniform positive gap. Equation (24) therefore forces a positive linear number of singleton tokens. Front-loading puts these singletons at the end of every exact schedule. Hence every nonsingleton batch begins with $r \ge c_\mu n$, while its size is $O(\sqrt{n})$; uniformly, $b/r = o(1)$.

This bootstrap makes the Riemann comparison global. Singleton intervals have zero collision functional, and for each nonsingleton batch,

$$\left| h(r,b) - \int_{(r-b)/n}^{r/n} \frac{b-1}{2u}\, du \right| \le \frac{5(b/r)}{6(1-b/r)} L(r,b).$$

Summing and using $\sum L(r,b) \le H_{n,T_n} = O(1)$ shows that the exact hazard and continuum functional differ by $o(1)$. Strong duality and (20) now yield the global certificate

$$\int_0^1 \Delta_\mu(u, \beta_n(u))\, du \longrightarrow 0. \tag{25}$$

There are only finitely many action-tie thresholds in $(0, 1]$. Away from arbitrarily small neighborhoods of those points and of zero, the incorrect actions have a uniform positive dual gap; large actions are included because $\Delta_\mu(u, b) \ge (b-1)/2 - \mu \to \infty$. Equation (25) therefore implies $\beta_n \to \beta_\mu$ in Lebesgue measure. Since a size-b batch occupies b/n token measure, let $N_{b,n}$ be the number of size-b batches in an exact optimal schedule. Then

$$\frac{bN_{b,n}}{n} = |\{u : \beta_n(u) = b\}| \longrightarrow \theta_b$$

for every $1 \le b \le K$, and the token measure assigned to $b > K$ tends to zero.

For $\tau = 1$, the identities $\sum_b bN_{b,n} = n$ and $\sum_b N_{b,n} = T_n$ give

$$\frac{1}{n}\sum_{b\ge 2} bN_{b,n} \le \frac{2}{n}\sum_{b\ge 2}(b-1)N_{b,n} = \frac{2(n-T_n)}{n} \longrightarrow 0.$$

Thus the singleton token fraction tends to one, giving the profile with $K = 1$. The argument used only exact optimality, not a rule for breaking finite ties, proving the full quantifier in Theorem 5.3.