CRITICAL CIRCULATIONS GOVERN SECOND-ORDER COMPLEXITY UNDER DYNAMIC ACCESS

Anonymous authors
Paper under double-blind review

ABSTRACT

First-order sequential-design theory replaces an adaptive policy by its long-run occupation measure. This forgets order. Under dynamic access constraints, where choosing an experiment changes what can be chosen next, order can determine the first unbounded correction to sample complexity. We characterize that correction. The zero-dual-slack edges form a critical graph, whose signed circulations generate a linear space of relative likelihood drifts. With a full-support dual optimum and a strongly connected critical graph, its rank yields a dichotomy: full relative rank permits restoring control and $O(\log b)$ excess, whereas a missing direction with nondegenerate noise forces $\Theta(\sqrt{b})$ excess. The criterion is computable by a nullspace and matrix-rank calculation on the first-order critical graph. An exact identity charges every excess sample either to terminal likelihood imbalance or to leaving the critical face. For a first-order-equivalent pair, we sharpen the two regimes to $V^{\mathrm{cert}}_{\mathrm{Flex}}(b) = rb/d + O(1)$, $V^{\mathrm{cert}}_{\mathrm{Lock}}(b) = rb/d + \kappa_r\sqrt{b} + o(\sqrt{b})$. The gap persists in fixed confidence. A second pair keeps the statewise action menus and critical cycle rank fixed but changes the routing; its second-order cost changes again. Optimal frequencies and cycle count are therefore insufficient. What matters is where critical circulations point in relative likelihood space.

1 INTRODUCTION

First-order sequential-design theory sees an adaptive policy through its long-run experiment frequencies (Chernoff, 1959; Nitinawarat et al., 2013; Garivier & Kaufmann, 2016). That summary determines the leading information rate but discards execution order. The loss is harmless when every experiment can be repeated at will. It is consequential when a sensing action moves a robot, reconfigures an instrument, or otherwise changes which experiment is available next. Two access systems can then share the same observation laws, optimal frequencies, and first-order rate while requiring different numbers of samples beyond the leading term.

This difference is not a negligible constant. At confidence level δ, the evidence threshold is of order $\log(1/\delta)$ (Dragalin et al., 1999), so a $\sqrt{\log(1/\delta)}$ correction is the first growing term that separates architectures declared equivalent by first-order theory. We show when dynamic access creates this diffusion-scale cost and when feedback removes it. The resulting question is architectural, not merely algorithmic: can the access mechanism reorder first-order-optimal experiments so as to repair a random evidence deficit, or must any repair leave the optimal face?

Model access by a directed multigraph: an edge both executes an experiment and sets the next access state. For a fixed true hypothesis i, the first-order benchmark is a stationary-flow max–min problem with value $C_i^\star$, in the spirit of occupation-measure formulations for controlled Markov systems (Graves & Lai, 1997; Karthik et al., 2024). Its dual assigns each edge a nonnegative slack. Edges that are tight for every optimal dual solution form the *critical graph*, which supports every first-order optimal occupation.

The edge set alone is not enough. Executable perturbations must preserve flow conservation. If B_i° is the critical incidence matrix and $\mathbf{D}_i^\circ$ maps edge usage to mean likelihood drift, then the zero-cost relative drift directions form $\bar{\mathcal{L}}_i^\circ = \bar{\mathbf{D}}_i(\ker B_i^\circ)$. Topology supplies the signed circulations; the observation model decides where they point. The rank of this *critical circulation image*—not the

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.

number of cycles—measures how many likelihood imbalances can be corrected without sacrificing first-order rate.

This exposes what occupation measures omit. A first-order optimum identifies a point, or a face, of stationary flows. A causal policy must also move locally along that face in response to fluctuations. Those zero-slack motions are precisely signed critical circulations, and their projected drifts form the tangent directions that matter statistically. Two systems can therefore have the same optimal point while offering different local corrective motion. The second-order problem is controlled by this tangent geometry.

Contributions. First, a relative-interior optimal dual solution gives, for every causal policy with finite expected certification time,

$$C_i^\star \left(\mathbb{E}_i^\pi T_b - \frac{b}{C_i^\star} \right) = \mathbb{E}_i^\pi I_{T_b}^\lambda + \mathbb{E}_i^\pi D_{T_b} + O(1). \tag{1}$$

Here $I_{T_b}^\lambda$ is the terminal relative-likelihood imbalance and D_{T_b} is the accumulated cost of leaving the critical face. These are the only two sources of growing excess.

Second, a relative-interior theorem turns the geometry into a rank test. Full relative rank yields $V_i^{\mathrm{cert}}(b) = b/C_i^\star + O(\log b)$. If the image misses a direction with nondegenerate critical noise, every policy must absorb its $\sqrt{b}$-scale fluctuation as terminal imbalance or positive slack, giving $V_i^{\mathrm{cert}}(b) = b/C_i^\star + \Theta(\sqrt{b})$. Operationally, one extracts the universally tight edges from a relative-interior dual optimum, forms a basis of their signed circulations, and computes the rank of the projected drift matrix.

Third, the distinction is sharp. Flex and Lock have the same hypotheses, observation laws, state dynamics, optimal aggregate frequencies, and first-order rate d/r, but circulation-image ranks $r-1$ and zero:

$$V_{0,\mathrm{Flex}}^{\mathrm{cert}}(b) = \frac{rb}{d} + O(1), \qquad V_{0,\mathrm{Lock}}^{\mathrm{cert}}(b) = \frac{rb}{d} + \kappa_r \sqrt{b} + o(\sqrt{b}).$$

The constant $\kappa_r > 0$ is explicit, and the gap transfers to fixed confidence. A routing pair then keeps the statewise menus and critical cycle rank fixed, ruling out both aggregate frequencies and cycle count as substitutes for the circulation image.

Ex ante structural test. The classification is available before policy design or simulation. The first-order program and projected circulation basis either yield all restoring directions or a missing contrast. Full span makes the remaining problem algorithmic; a noisy missing contrast is a policy-independent obstruction. Thus the rank test separates a failure of control design from a limitation of the access architecture.

Relation to prior work. Controlled sensing with Markov observations, action costs, usage constraints, or time-varying sources changes the first-order design problem (Nitinawarat & Veeravalli, 2015; Vaidhiyan & Sundaresan, 2015; Li et al., 2019; Hsu & Wang, 2025); classical second-order theory refines boundary crossing for repeatable experiments or controlled random walks (Keener, 1984; Lalley & Lorden, 1986; Dragalin et al., 2000). Here access is endogenous and order determines the correction scale itself. The governing object is the drift image of zero-slack circulations, not another optimal frequency or cycle count. Appendix F gives the detailed comparison.

2 Problem Setup and First-Order Benchmark

2.1 Access-constrained sequential identification

Let Θ be a finite hypothesis set. Fix a true hypothesis $i \in \Theta$, let $\mathcal{J}_i := \Theta \setminus \{i\}$, and write $r := |\mathcal{J}_i|$. Experiment availability is specified by a finite directed multigraph $G = (\mathcal{Z}, E)$. An edge $e = (z, a, z')$ represents an experiment a that is admissible in access state z and moves the access state to z'. The initial state $Z_0 = z_0$ is fixed.

Let $(\mathcal{F}_t)_{t\geq 0}$ denote the history filtration, including any policy randomization. At time t, a causal policy chooses $E_t = (Z_{t-1}, A_t, Z_t)$ before observing Y_t; thus E_t is $\mathcal{F}_{t-1}$-measurable. Under

hypothesis $\theta \in \Theta$,

$$\mathbb{P}_\theta(Y_t \in A \mid \mathcal{F}_{t-1}) = P_\theta^{E_t}(A)$$

for every measurable A. For $j \in \mathcal{J}_i$, define

$$X_e^j(y) := \log \frac{dP_i^e}{dP_j^e}(y), \qquad d_{i,j}(e) := \mathbb{E}_i[X_e^j(Y)] = D_{\mathrm{KL}}(P_i^e \| P_j^e),$$

and collect these coordinates as $X_e \in \mathbb{R}^r$ and $d_i(e) \in \mathbb{R}^r$. The cumulative log-likelihood vector is

$$S_t := \sum_{s=1}^{t} X_{E_s}(Y_s). \tag{2}$$

For $b > 0$, let

$$T_b := \inf\left\{ t \geq 0 : \min_{j \in \mathcal{J}_i} S_t^j \geq b \right\}, \qquad V_i^{\mathrm{cert}}(b) := \inf_{\pi} \mathbb{E}_i^\pi[T_b]. \tag{3}$$

The fixed-i certification problem isolates the structural cost of accumulating evidence against all alternatives. For the Flex/Lock example, Section 6 later translates the resulting separation to the usual unknown-hypothesis fixed-confidence formulation.

Standing assumptions. Unless stated otherwise, $\Theta, \mathcal{Z}, E$ are finite, the full access graph G is strongly connected, $C_i^\star > 0$, and $|X_e^j| \leq L$ almost surely for every e, j. Section 5 adds assumptions on the optimal dual set and the critical graph.

2.2 First-order occupation problem

Let $B \in \mathbb{R}^{|\mathcal{Z}| \times |E|}$ be the incidence matrix, with column $-\mathbf{e}_z + \mathbf{e}_{z'}$ for $e = (z, a, z')$. The normalized stationary flow polytope is

$$\mathcal{W}(G) := \{ w \in \mathbb{R}_+^{|E|} : Bw = 0,\ \mathbf{1}^\top w = 1 \}. \tag{4}$$

Let $\mathbf{D}_i \in \mathbb{R}^{r \times |E|}$ have columns $d_i(e)$. The first-order max–min evidence rate is

$$C_i^\star := \max_{w \in \mathcal{W}(G)} \min_{j \in \mathcal{J}_i} (\mathbf{D}_i w)_j. \tag{5}$$

Its dual representation is

$$C_i^\star = \min_{\lambda \in \Delta_r,\ h \in \mathbb{R}^{\mathcal{Z}}} \max_{e=(z,a,z') \in E} \left\{ \lambda^\top d_i(e) + h(z') - h(z) \right\}, \tag{6}$$

where $\Delta_r := \{\lambda \geq 0 : \mathbf{1}^\top \lambda = 1\}$. We normalize $h(z_{\mathrm{ref}}) = 0$ and denote the normalized optimal dual set by $\mathcal{D}_i^\star$. Without access constraints, analogous max–min information rates underlie classical sequential design and controlled sensing (Chernoff, 1959; Nitinawarat et al., 2013; Nitinawarat & Veeravalli, 2015).

3 The First-Order Critical Access Structure

3.1 The zero-slack access graph

For $q = (\lambda, h) \in \mathcal{D}_i^\star$ and $e = (z, a, z')$, define

$$\Delta_q(e) := C_i^\star - \left[\lambda^\top d_i(e) + h(z') - h(z) \right] \geq 0. \tag{7}$$

The critical edge set consists of the edges that are tight for every optimal dual solution:

$$E_i^\circ := \{ e \in E : \Delta_q(e) = 0 \text{ for every } q \in \mathcal{D}_i^\star \}, \qquad G_i^\circ := (\mathcal{Z}, E_i^\circ). \tag{8}$$

Proposition 1. *Let $q^\circ = (\lambda^\circ, h^\circ) \in \mathrm{i}(\mathcal{D}_i^\star)$. Then:*

1. *$\Delta_{q^\circ}(e) = 0$ if and only if $e \in E_i^\circ$;*
2. *every optimizer $w^\star$ of equation 5 satisfies $\mathrm{supp}(w^\star) \subseteq E_i^\circ$;*
3. *if $\mathcal{D}_i^\star$ contains an optimal multiplier with full support, then every relative-interior optimal multiplier has full support and every primal optimizer satisfies*

$$\mathbf{D}_i w^\star = C_i^\star \mathbf{1}. \tag{9}$$

The proof is in Appendix A.1. Fix $q^\circ = (\lambda^\circ, h^\circ) \in \mathrm{i}(\mathcal{D}_i^\star)$ from now on and write $\Delta(e) := \Delta_{q^\circ}(e)$. By Proposition 1, $\Delta(e) > 0$ exactly on the noncritical edges.

3.2 CRITICAL CIRCULATIONS AND RELATIVE LIKELIHOOD DRIFT

Let B_i° and $\mathbf{D}_i^\circ$ be the restrictions of B and $\mathbf{D}_i$ to E_i°. Since a common shift of all likelihood coordinates does not affect their relative values, work in the quotient

$$\mathsf{Q}_i := \mathbb{R}^r / \operatorname{span}\{\mathbf{1}\}, \tag{10}$$

with dual $\mathsf{Q}_i^* := \{v \in \mathbb{R}^r : v^\top \mathbf{1} = 0\}$. For $x \in \mathbb{R}^{E_i^\circ}$, write $\bar{\mathbf{D}}_i x = [\mathbf{D}_i^\circ x]$. The image of the signed circulation space is

$$\bar{\mathcal{L}}_i^\circ := \bar{\mathbf{D}}_i(\ker B_i^\circ) \subseteq \mathsf{Q}_i. \tag{11}$$

Thus $\bar{\mathcal{L}}_i^\circ$ records the relative mean-likelihood directions generated by signed circulations on the critical access graph.

Normalized critical operation is described by

$$\mathcal{W}_i^\circ := \{w \geq 0 : B_i^\circ w = 0,\ \mathbf{1}^\top w = 1\}, \qquad \bar{\mathcal{K}}_i^\circ := \bar{\mathbf{D}}_i(\mathcal{W}_i^\circ). \tag{12}$$

The set $\bar{\mathcal{K}}_i^\circ$ is compact and convex, whereas $\bar{\mathcal{L}}_i^\circ$ is linear. For every $w \in \mathcal{W}_i^\circ$, zero dual slack and flow conservation imply

$$(\lambda^\circ)^\top \mathbf{D}_i^\circ w = C_i^\star. \tag{13}$$

Critical circulations can therefore change relative evidence while preserving the optimal λ°-weighted first-order rate. If the full-support condition in Proposition 1 holds, any primal optimizer belongs to $\mathcal{W}_i^\circ$ and has zero relative drift; hence $0 \in \bar{\mathcal{K}}_i^\circ$.

Critical cycle-rank obstruction. The circulation space gives a purely topological upper bound on the number of independent relative-likelihood directions available at zero first-order cost.

Corollary 2. *Assume that $\mathcal{D}_i^\star$ contains an optimal multiplier with full support and that G_i° is strongly connected. Define the critical cycle rank*

$$c_i := \dim \ker B_i^\circ = |E_i^\circ| - |\mathcal{Z}| + 1. \tag{14}$$

Then

$$\bar{\mathcal{L}}_i^\circ = \bar{\mathbf{D}}_i \left(\ker B_i^\circ \cap \mathbf{1}^\perp\right), \tag{15}$$

and consequently

$$\dim \bar{\mathcal{L}}_i^\circ \leq c_i - 1. \tag{16}$$

In particular,

$$0 \in \operatorname{int}_{\mathsf{Q}_i} \bar{\mathcal{K}}_i^\circ \quad \Longrightarrow \quad c_i \geq r. \tag{17}$$

The canonical definition of the critical graph yields more than this dimension bound. It also rules out a one-sided boundary inside the circulation image.

Proposition 3 (Relative interior and rank)**.** *Assume that $\mathcal{D}_i^\star$ contains a full-support multiplier and that G_i° is strongly connected. Then*

$$0 \in \operatorname{ri}_{\bar{\mathcal{L}}_i^\circ} \bar{\mathcal{K}}_i^\circ. \tag{18}$$

Equivalently, the cone of supporting contrasts at the origin is the annihilator of the critical circulation image:

$$\left\{v \in \mathsf{Q}_i^* : v^\top \mathbf{D}_i^\circ w \geq 0\ \forall w \in \mathcal{W}_i^\circ\right\} = \left(\bar{\mathcal{L}}_i^\circ\right)^\perp. \tag{19}$$

Consequently,

$$0 \in \operatorname{int}_{\mathsf{Q}_i} \bar{\mathcal{K}}_i^\circ \quad \Longleftrightarrow \quad \dim \bar{\mathcal{L}}_i^\circ = r - 1. \tag{20}$$

Why a rank test suffices. Suppose that v supports $\bar{\mathcal{K}}_i^\circ$ at the origin. Min-cost circulation duality produces a potential ψ such that

$$a_v(e) := v^\top d_i(e) - [\psi(z') - \psi(z)] \geq 0, \qquad e = (z, a, z') \in E_i^\circ.$$

Perturbing the relative-interior dual solution to $(\lambda - \varepsilon v, h + \varepsilon\psi)$ preserves feasibility and optimality for small $\varepsilon > 0$. Since a critical edge is tight for every dual optimum, $a_v(e) = 0$ throughout the critical graph. Summing around signed circulations cancels ψ, so v annihilates the circulation image. Thus the origin has no one-sided boundary inside that image, and full relative rank is exactly the restoring condition.

Proofs of Corollary 2 and Proposition 3 are in Appendix A.2. The conclusion is sharper than a cycle count: distinct critical cycles can induce the same relative drift, so the effective dimension is the rank after projection through $\bar{\mathbf{D}}_i$.

Free-control specialization. With one access state, $B_i^\circ = 0$. If $\mathcal{A}_i^\circ$ is the set of critical experiments, then

$$\begin{aligned}\bar{\mathcal{L}}_i^\circ &= \operatorname{span}\{[d_i(a)] : a \in \mathcal{A}_i^\circ\},\\ \bar{\mathcal{K}}_i^\circ &= \operatorname{conv}\{[d_i(a)] : a \in \mathcal{A}_i^\circ\}.\end{aligned} \tag{21}$$

Thus the classical span/convex-hull criterion for repeatable experiments is the one-state case of the circulation-image criterion. In particular, at least r critical experiments are needed to restore all $r-1$ relative directions.

4 An Exact Decomposition of the Excess Stopping Time

Fix $q^\circ = (\lambda, h) \in \mathrm{i}(\mathcal{D}_i^\star)$ and write $\Delta(e) := \Delta_{q^\circ}(e)$. Consider the two nonnegative processes $I_t^\lambda := \lambda^\top S_t - \min_j S_t^j, \qquad D_t := \sum_{s=1}^t \Delta(E_s)$. The first records relative-evidence imbalance, while the second records the cumulative dual slack incurred by noncritical edges.

4.1 Exact accounting

Let $O_b := \min_j S_{T_b}^j - b, \qquad 0 \le O_b \le L$. By the definition of the dual slack,

$$\mathcal{M}_t := \lambda^\top S_t + h(Z_t) - h(Z_0) - C_i^\star t + D_t \tag{22}$$

is a bounded-increment martingale.

Theorem 4. *For every causal policy π with $\mathbb{E}_i^\pi[T_b] < \infty$,*

$$C_i^\star \mathbb{E}_i^\pi[T_b] = b + \mathbb{E}_i^\pi[O_b] + \mathbb{E}_i^\pi[I_{T_b}^\lambda] + \mathbb{E}_i^\pi[D_{T_b}] + \mathbb{E}_i^\pi[h(Z_{T_b}) - h(Z_0)]. \tag{23}$$

Consequently,

$$C_i^\star \left(\mathbb{E}_i^\pi[T_b] - \frac{b}{C_i^\star} \right) = \mathbb{E}_i^\pi[I_{T_b}^\lambda] + \mathbb{E}_i^\pi[D_{T_b}] + O(1), \tag{24}$$

uniformly over b and causal policies with finite expected stopping time.

Appendix A.3 proves the theorem by optional stopping applied to equation 22. Since the overshoot and potential terms are bounded, equation 24 leaves only two possible sources of growing excess: terminal relative-evidence imbalance and cumulative dual slack.

4.2 A universal finite-threshold bound

Corollary 5. *Under the standing assumptions,*

$$\frac{b}{C_i^\star} - O(1) \le V_i^{\mathrm{cert}}(b) \le \frac{b}{C_i^\star} + O(\sqrt{b}). \tag{25}$$

The lower bound follows from exact accounting. A length-$\Theta(\sqrt{b})$ macro-tour approximating an optimal flow gives the upper bound; the construction and tail estimate are in Appendix A.4.

5 Second-Order Consequences of the Critical Drift Geometry

In this section, assume in addition that $\mathcal{D}_i^\star$ contains a full-support optimal multiplier and that G_i° is strongly connected. For the fixed relative-interior optimizer $q^\circ = (\lambda, h)$, this gives $\lambda_j > 0$ for every j and $0 \in \bar{\mathcal{K}}_i^\circ$. For $f : \mathcal{Z} \to \mathbb{R}$, write $\operatorname{span}(f) := \max_{z\in\mathcal{Z}} f(z) - \min_{z\in\mathcal{Z}} f(z)$.

For $v \in \mathsf{Q}_i^*$, let

$$\nu_{i,v}(e) := \operatorname{Var}_i\big(v^\top X_e(Y)\big), \qquad e \in E.$$

The contrasts missed by every signed critical circulation form

$$\mathcal{U}_i^\circ := \left\{ v \in \mathsf{Q}_i^* : v^\top \mathbf{D}_i^\circ x = 0 \ \forall x \in \ker B_i^\circ \right\}. \tag{26}$$

This is the annihilator of $\bar{\mathcal{L}}_i^\circ$, and therefore

$$\dim \mathcal{U}_i^\circ = (r-1) - \dim \bar{\mathcal{L}}_i^\circ \ge r - c_i. \tag{27}$$

For $v \in \mathcal{U}_i^\circ$, define its smallest stationary variance rate on the critical graph by

$$\underline{\nu}_i(v) := \min_{w \in \mathcal{W}_i^\circ} \sum_{e \in E_i^\circ} w_e \nu_{i,v}(e). \tag{28}$$

Theorem 6 (Critical-rank dichotomy)**.** *Under the assumptions above:*

(a) *If*

$$\dim \bar{\mathcal{L}}_i^\circ = r - 1, \tag{29}$$

then there exists a causal critical-edge policy satisfying

$$\mathbb{E}_i[T_b] \le \frac{b}{C_i^\star} + O(\log b), \tag{30}$$

and hence $V_i^{\mathrm{cert}}(b) = b/C_i^\star + O(\log b)$.

(b) *If there exists* $v \in \mathcal{U}_i^\circ \setminus \{0\}$ *with*

$$\underline{\nu}_i(v) > 0, \tag{31}$$

then for some $c > 0$*,*

$$\liminf_{b \to \infty} \inf_{\pi:\, \mathbb{E}_i^\pi T_b < \infty} \frac{\mathbb{E}_i^\pi I_{T_b}^\lambda + \mathbb{E}_i^\pi D_{T_b}}{\sqrt{b}} \ge c, \tag{32}$$

and consequently

$$V_i^{\mathrm{cert}}(b) = \frac{b}{C_i^\star} + \Theta(\sqrt{b}). \tag{33}$$

Operational test. Let the columns of N_i° form a basis of $\ker B_i^\circ \cap \mathbf{1}^\perp$, and set

$$R_i^\circ := \bar{\mathbf{D}}_i N_i^\circ. \tag{34}$$

Then $\mathrm{rank}(R_i^\circ) = r - 1$ is exactly the restoring criterion. If the rank is smaller, any nonzero $v \in \ker[(R_i^\circ)^\top] \cap \mathsf{Q}_i^*$ is a missing contrast; checking $\underline{\nu}_i(v) > 0$ certifies the diffusion regime. The classification therefore requires only a first-order primal–dual solve, a circulation-space basis, and a rank and variance check.

Concretely, N_i° is obtained as a nullspace basis of the stacked constraints $B_i^\circ x = 0$ and $\mathbf{1}^\top x = 0$. A singular-value decomposition of R_i° then either certifies full relative rank or returns a left-null vector v. In the latter case, the corresponding critical-edge potential ψ is recovered from the coboundary equations, and $\underline{\nu}_i(v)$ is a linear program over $\mathcal{W}_i^\circ$. Thus the same finite-dimensional calculation produces either the directions used by a restoring controller or a witness for the diffusion lower bound.

Corollary 7 (Monotonicity under critical augmentation)**.** *Consider two systems with the same access states and alternative coordinates. Suppose that their critical graphs satisfy* $E_i^\circ \subseteq \widetilde{E}_i^\circ$*, and that common edges have the same drift vectors. Then*

$$\bar{\mathcal{L}}_i^\circ \subseteq \widetilde{\bar{\mathcal{L}}}_i^\circ, \qquad \dim \bar{\mathcal{L}}_i^\circ \le \dim \widetilde{\bar{\mathcal{L}}}_i^\circ. \tag{35}$$

If a newly available zero-mass circulation has projected drift outside $\bar{\mathcal{L}}_i^\circ$*, the rank increases by at least one.*

Indeed, an old zero-mass circulation extends to the augmented edge set by assigning zero flow to every new edge, so its relative drift is unchanged. Thus eliminating a rank deficit of dimension k requires at least k new independent projected cycle drifts. Merely increasing the number of critical cycles can leave the second-order regime unchanged if their drifts are redundant. The proviso that the added edges remain critical is essential: an edge with positive dual slack may enable repair, but every use is then charged directly by the accounting identity.

The dichotomy is a property of one linear map. Full rank means that critical circulations can oppose every local evidence imbalance. A missing direction has zero mean around every critical cycle; if it remains noisy, every policy must pay order $\sqrt{b}$ in terminal imbalance or positive slack.

If

$$\underline{\nu}_i(v) > 0 \qquad \forall v \in \mathcal{U}_i^\circ \setminus \{0\}, \tag{36}$$

then Theorem 6 gives the diffusion-scale classification

$$V_i^{\mathrm{cert}}(b) - \frac{b}{C_i^\star} = \begin{cases} o(\sqrt{b}), & \dim \bar{\mathcal{L}}_i^\circ = r-1, \\ \Theta(\sqrt{b}), & \dim \bar{\mathcal{L}}_i^\circ < r-1. \end{cases}$$

The first line uses the sharper $O(\log b)$ bound in part (a).

Proof mechanism. Under full rank, Proposition 3 supplies a uniform inward drift in every relative-evidence direction. A feedback choice among critical cycles stabilizes the imbalance while paying no dual slack, giving $\mathbb{E}_i I_{T_b}^\lambda = O(\log b)$. If rank is deficient, a missing contrast is a graph coboundary on the critical edges. For a potential ψ and the corresponding martingale M^v,

$$v^\top d_i(e) = \psi(z') - \psi(z), \qquad \langle M^v \rangle_t \geq \underline{\nu}_i(v)t - K_v\{N_{\mathrm{nc}}(t) + 1\}.$$

Its mean motion is bounded but its quadratic variation grows linearly. A $\sqrt{b}$-scale fluctuation therefore survives to stopping unless the policy pays comparable noncritical slack. Exact accounting converts these two mechanisms into the bounds above. Appendix A.5 and Appendix A.6 give the proofs.

Since $\dim \mathcal{U}_i^\circ \geq r - c_i$, the noisy-deficiency conclusion applies in particular when $c_i < r$ and every nonzero missing contrast has positive minimal critical variance.

6 Same First-Order Benchmark, Different Second-Order Complexity

We now hold the observation model, state dynamics, first-order rate, and optimal aggregate frequencies fixed, and change only which experiments can be executed along the state trajectory.

6.1 A first-order-equivalent Flex/Lock pair

Fix $r \geq 2$ and let $\Theta = \{0, 1, \ldots, r\}$, with experiments $\mathcal{A} = \{a_1, \ldots, a_r\}$. Choose $p, q \in (0,1)$, $p \neq q$. Under H_0, every experiment has distribution $\mathrm{Bern}(p)$. Under H_j, only a_j changes distribution:

$$P_0^{a_k} = \mathrm{Bern}(p), \qquad P_j^{a_k} = \begin{cases} \mathrm{Bern}(q), & k = j, \\ \mathrm{Bern}(p), & k \neq j. \end{cases}$$

Let $d = D_{\mathrm{KL}}(\mathrm{Bern}(p) \| \mathrm{Bern}(q))$, and let σ^2 be the variance under $\mathrm{Bern}(p)$ of the corresponding log-likelihood ratio. Then $d_0(a_j) = d\mathbf{e}_j$. Both systems follow the deterministic rotor $z_1 \to \cdots \to z_r \to z_1$. Flex admits every a_j at every state; Lock admits only a_k at z_k. Thus Lock forces uniform sampling, while Flex can realize any aggregate marginal $u \in \Delta_r$. In both systems,

$$C_{0,\mathrm{Flex}}^\star = C_{0,\mathrm{Lock}}^\star = \frac{d}{r}, \qquad u_j^\star = \frac{1}{r}. \tag{37}$$

Every edge is critical in both systems. Flex circulations realize every $u \in \Delta_r$, whereas Lock has a unique stationary circulation with drift parallel to $\mathbf{1}$. Consequently,

$$\begin{aligned} &\bar{\mathcal{L}}_{0,\mathrm{Flex}}^\circ = \mathsf{Q}_0, \qquad \dim \bar{\mathcal{L}}_{0,\mathrm{Flex}}^\circ = r-1, \\ &\bar{\mathcal{L}}_{0,\mathrm{Lock}}^\circ = \{0\}, \qquad \underline{\nu}_0(v) = \frac{\sigma^2}{r} \sum_{j=1}^r v_j^2 > 0 \quad (v \in \mathsf{Q}_0^* \setminus \{0\}). \end{aligned} \tag{38}$$

Theorem 6 therefore yields

$$V_{0,\mathrm{Flex}}^{\mathrm{cert}}(b) = \frac{rb}{d} + O(\log b), \qquad V_{0,\mathrm{Lock}}^{\mathrm{cert}}(b) = \frac{rb}{d} + \Theta(\sqrt{b}).$$

The next result sharpens the Flex remainder to $O(1)$ and identifies the Lock constant.

6.2 SHARP CERTIFICATION SEPARATION

Let $Z_1, \dots, Z_r$ be i.i.d. standard Gaussians, let $\xi_1, \xi_2, \dots$ be i.i.d. informative log-likelihood increments under H_0, and define

$$\kappa_r := \frac{r\sigma}{d^{3/2}} \mathbb{E}\left[\max_{1\le j\le r} Z_j\right], \quad \tau_b := \inf\left\{n \ge 0 : \sum_{m=1}^{n} \xi_m \ge b\right\}. \tag{39}$$

Theorem 8. *For the Flex/Lock pair of Section 6.1,*

$$\frac{T_{b,\mathrm{Lock}} - \frac{rb}{d}}{\sqrt{b}} \Longrightarrow -\frac{r\sigma}{d^{3/2}} \min_{1\le j\le r} Z_j. \tag{40}$$

Moreover,

$$V^{\mathrm{cert}}_{0,\mathrm{Flex}}(b) = r\mathbb{E}_0[\tau_b] = \frac{rb}{d} + O(1), \tag{41}$$

while

$$V^{\mathrm{cert}}_{0,\mathrm{Lock}}(b) = \frac{rb}{d} + \kappa_r\sqrt{b} + o(\sqrt{b}), \tag{42}$$

Consequently,

$$V^{\mathrm{cert}}_{0,\mathrm{Lock}}(b) - V^{\mathrm{cert}}_{0,\mathrm{Flex}}(b) = \kappa_r\sqrt{b} + o(\sqrt{b}). \tag{43}$$

Flex samples coordinates in turn and freezes each at its first crossing, reducing the problem to r renewal times. Lock cannot react to the weakest coordinate; its minimum Gaussian fluctuation at the first-order crossing produces $\kappa_r\sqrt{b}$. Appendix B makes this first-crossing argument uniform and proves the expectation expansion.

6.3 SAME MENUS AND CYCLE RANK, DIFFERENT ROUTING

To isolate routing, set $r = 2$ and use states z_0, z_1, z_2. The action menus are identical: both actions are available at z_0, only a_1 at z_1, and only a_2 at z_2. Only the destinations differ:

state	action	Split route	Coupled route
z_0	a_1	z_1	z_2
z_0	a_2	z_2	z_1
z_1	a_1	z_0	z_0
z_2	a_2	z_0	z_0

The Split cycles execute (a_1, a_1) and (a_2, a_2); the Coupled cycles execute (a_1, a_2) and (a_2, a_1).

Proposition 9 (Routing separation)**.** *The Split and Coupled systems have the same first-order rate $d/2$, the same unique optimal aggregate action frequencies $(1/2, 1/2)$, and the same critical cycle rank 2. Nevertheless,*

$$\dim \bar{\mathcal{L}}^{\circ}_{0,\mathrm{Split}} = 1, \qquad \bar{\mathcal{L}}^{\circ}_{0,\mathrm{Coupled}} = \{0\}, \tag{44}$$

and

$$V^{\mathrm{cert}}_{0,\mathrm{Split}}(b) = \frac{2b}{d} + O(1), \tag{45}$$

$$V^{\mathrm{cert}}_{0,\mathrm{Coupled}}(b) = \frac{2b}{d} + \kappa_2\sqrt{b} + o(\sqrt{b}). \tag{46}$$

Both graphs have cycle rank two, but the Split cycle drifts span the relative direction while every Coupled cycle has zero relative drift. Thus an adaptive choice at the root is useful only if it routes the process onto a cycle pointing in the needed direction. Appendix C proves the sharp expansions.

What the comparisons identify. The two pairs and the policy ablation remove three distinct confounds. Flex/Lock holds fixed the observation laws, optimal aggregate frequencies, and first-order rate. Split/Coupled additionally holds fixed the statewise menus and critical cycle rank, isolating where those cycles point after projection. Finally, the Flex ablation holds the access graph itself fixed: restoring policies remove the diffusion term, while round robin does not. Together these comparisons separate availability, statistical direction, and actual feedback use—all three are needed for bounded second-order excess.

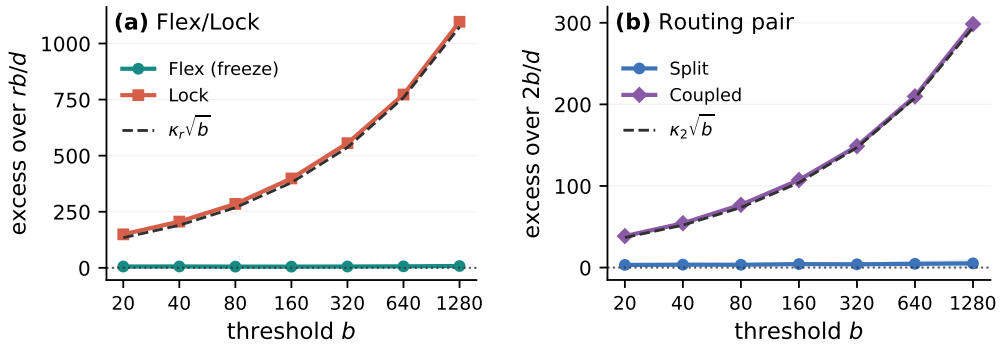

Figure 1: Finite-threshold excess for $p = 0.3$, $q = 0.6$, based on 20,000 trajectories; shaded bands are pointwise 95% Monte Carlo intervals. **(a)** For $r = 4$, optimal Flex has bounded excess, whereas Lock follows the predicted $\kappa_4\sqrt{b}$ correction. **(b)** Split and Coupled have the same statewise menus and critical cycle rank; changing only the routing again separates bounded and $\sqrt{b}$-scale excess. Dashed curves are the theoretical diffusion terms.

Numerical illustration. We simulate both separation pairs under the same Bernoulli model for $b \in \{20, \dots, 1280\}$, using 20,000 trajectories per policy. Flex uses the coordinatewise-freeze policy that attains equation 41; Split uses a weakest-coordinate restoring cycle, while Lock and Coupled cannot alter their round composition. A weighted fit of excess to $\alpha\sqrt{b} + \beta$ over the four largest thresholds gives $\widehat{\alpha} = 0.11, 30.01, 0.03, 8.21$ for Flex, Lock, Split, and Coupled, against theoretical values $0, 30.00, 0, 8.22$. Moreover, round robin on the Flex graph recovers the Lock coefficient 30.01: access to restoring cycles matters only when the policy uses them. Figure 1 shows the two pairwise comparisons; Appendix E gives Monte Carlo uncertainty, policy ablations, and distributional diagnostics.

6.4 Fixed-confidence consequence

Let $V^{\mathrm{fc}}_{0,\mathsf{A}}(\delta)$ be the minimum $\mathbb{E}_0\tau$ over procedures that terminate under every hypothesis and have worst-case error at most δ in system $\mathsf{A} \in \{\mathrm{Flex}, \mathrm{Lock}\}$. Write $B_\delta := \log(1/\delta)$.

Corollary 10. *As $\delta \downarrow 0$,*

$$V^{\mathrm{fc}}_{0,\mathrm{Flex}}(\delta) = \frac{r}{d}B_\delta + O(1), \tag{47}$$

$$V^{\mathrm{fc}}_{0,\mathrm{Lock}}(\delta) = \frac{r}{d}B_\delta + \kappa_r\sqrt{B_\delta} + o(\sqrt{B_\delta}), \tag{48}$$

$$V^{\mathrm{fc}}_{0,\mathrm{Lock}}(\delta) - V^{\mathrm{fc}}_{0,\mathrm{Flex}}(\delta) = \kappa_r\sqrt{\log(1/\delta)} + o\Big(\sqrt{\log(1/\delta)}\Big). \tag{49}$$

Binary data processing gives the common first-order lower bound. A stopped change-of-measure argument transfers the sharp Lock certification correction to arbitrary δ-correct procedures; capped coordinatewise certification gives the Flex upper bound. Appendix D contains both arguments.

Discussion.

The design implication is not simply to add cycles, but to add critical cycles whose projected drifts enlarge the circulation image. The rank test separates topological freedom from statistically useful direction and can therefore diagnose an access architecture before optimizing its policy. Our dichotomy uses a full-support dual optimum, a strongly connected critical graph, and nondegenerate noise in every missing direction. Rank-deficient systems with degenerate noise, or critical graphs with several communicating components, may exhibit different intermediate behavior; resolving those cases requires combining circulation geometry with the active alternative set and component transitions.

AI USE STATEMENT

The authors used generative AI tools for language editing, manuscript organization, drafting and revising expository passages, and improving the clarity and consistency of the mathematical presentation. The tools were also used as critical readers to stress-test the exposition, identify unclear assumptions or possible internal inconsistencies, and anticipate alternative interpretations or potential reviewer misunderstandings. In addition, they assisted with literature discovery and bibliographic formatting. Candidate references and bibliographic information were subsequently checked against the corresponding papers, archival records, or publisher pages.

Generative AI tools provided limited assistance with revising and executing the authors' Monte Carlo simulation and plotting code, as well as organizing and summarizing the resulting numerical diagnostics. All code changes were inspected by the authors, and all reported numerical values and figures were generated by executing the reviewed source code and checked for consistency with the stated experimental setup and theoretical predictions. The authors reviewed all AI-assisted material and take full responsibility for the final text, claims, code, results, and citations in this paper.

REFERENCES

Herman Chernoff. Sequential design of experiments. *The Annals of Mathematical Statistics*, 30 (3):755–770, 1959. doi: 10.1214/aoms/1177706205. URL https://doi.org/10.1214/aoms/1177706205.

Kobi Cohen and Qing Zhao. Active hypothesis testing for anomaly detection. *IEEE Transactions on Information Theory*, 61(3):1432–1450, 2015. doi: 10.1109/TIT.2014.2387857. URL https://doi.org/10.1109/TIT.2014.2387857.

Aditya Deshmukh, Venugopal V. Veeravalli, and Srikrishna Bhashyam. Sequential controlled sensing for composite multihypothesis testing. *Sequential Analysis*, 40(2):259–289, 2021. doi: 10.1080/07474946.2021.1912525. URL https://doi.org/10.1080/07474946.2021.1912525.

Vladimir P. Dragalin, Alexander G. Tartakovsky, and Venugopal V. Veeravalli. Multihypothesis sequential probability ratio tests—part I: Asymptotic optimality. *IEEE Transactions on Information Theory*, 45(7):2448–2461, 1999. doi: 10.1109/18.796383. URL https://doi.org/10.1109/18.796383.

Vladimir P. Dragalin, Alexander G. Tartakovsky, and Venugopal V. Veeravalli. Multihypothesis sequential probability ratio tests—part II: Accurate asymptotic expansions for the expected sample size. *IEEE Transactions on Information Theory*, 46(4):1366–1383, 2000. doi: 10.1109/18.850677. URL https://doi.org/10.1109/18.850677.

Aurélien Garivier and Emilie Kaufmann. Optimal best arm identification with fixed confidence. In *Proceedings of the 29th Annual Conference on Learning Theory*, volume 49 of *Proceedings of Machine Learning Research*, pp. 998–1027. PMLR, 2016. URL https://proceedings.mlr.press/v49/garivier16a.html.

Todd L. Graves and Tze Leung Lai. Asymptotically efficient adaptive choice of control laws in controlled markov chains. *SIAM Journal on Control and Optimization*, 35(3):715–743, 1997. doi: 10.1137/S0363012994275440. URL https://doi.org/10.1137/S0363012994275440.

Chia-Yu Hsu and I-Hsiang Wang. Tradeoffs among action taking policies matter in active sequential multi-hypothesis testing: The optimal error exponent region. *IEEE Transactions on Information Theory*, 71(9):6546–6565, 2025. doi: 10.1109/TIT.2025.3588527. URL https://doi.org/10.1109/TIT.2025.3588527.

P. N. Karthik, Vincent Y. F. Tan, Arpan Mukherjee, and Ali Tajer. Optimal best arm identification with fixed confidence in restless bandits. *IEEE Transactions on Information Theory*, 70(10):7349–7384, 2024. doi: 10.1109/TIT.2024.3419924. URL https://doi.org/10.1109/TIT.2024.3419924.

Emilie Kaufmann, Olivier Cappé, and Aurélien Garivier. On the complexity of best-arm identification in multi-armed bandit models. *Journal of Machine Learning Research*, 17(1):1–42, 2016. URL https://www.jmlr.org/papers/v17/kaufman16a.html.

Robert W. Keener. Second order efficiency in the sequential design of experiments. *The Annals of Statistics*, 12(2):510–532, 1984. doi: 10.1214/aos/1176346503. URL https://doi.org/10.1214/aos/1176346503.

Steven P. Lalley and Gary Lorden. A control problem arising in the sequential design of experiments. *The Annals of Probability*, 14(1):136–172, 1986. doi: 10.1214/aop/1176992620. URL https://doi.org/10.1214/aop/1176992620.

Shang Li, Xiaoou Li, Xiaodong Wang, and Jingchen Liu. Sequential hypothesis test with online usage-constrained sensor selection. *IEEE Transactions on Information Theory*, 65(7):4392–4410, 2019. doi: 10.1109/TIT.2019.2910730. URL https://doi.org/10.1109/TIT.2019.2910730.

Weizhi Li, Prad Kadambi, Pouria Saidi, Karthikeyan Natesan Ramamurthy, Gautam Dasarathy, and Visar Berisha. Active sequential two-sample testing. *Transactions on Machine Learning Research*, 2024. URL https://openreview.net/forum?id=EzPRgIq2Tk.

Yonglong Li and Vincent Y. F. Tan. Second-order asymptotics of sequential hypothesis testing. *IEEE Transactions on Information Theory*, 66(11):7222–7230, 2020. doi: 10.1109/TIT.2020.3006014. URL https://doi.org/10.1109/TIT.2020.3006014.

Jingyu Liu and Yanglei Song. Sequential multiple testing: A second-order asymptotic analysis. *arXiv preprint arXiv:2603.04685*, 2026. URL https://arxiv.org/abs/2603.04685.

M. B. Malyutov and I. I. Tsitovich. Second-order optimal sequential tests. In Anthony Atkinson, Barbara Bogacka, and Anatoly Zhigljavsky (eds.), *Optimum Design 2000*, pp. 67–78. Springer, 2001. doi: 10.1007/978-1-4757-3419-5_7. URL https://doi.org/10.1007/978-1-4757-3419-5_7.

Mohammad Naghshvar and Tara Javidi. Active sequential hypothesis testing. *The Annals of Statistics*, 41(6):2703–2738, 2013a. doi: 10.1214/13-AOS1144. URL https://doi.org/10.1214/13-AOS1144.

Mohammad Naghshvar and Tara Javidi. Sequentiality and adaptivity gains in active hypothesis testing. *IEEE Journal of Selected Topics in Signal Processing*, 7(5):768–782, 2013b. doi: 10.1109/JSTSP.2013.2261279. URL https://doi.org/10.1109/JSTSP.2013.2261279.

Sirin Nitinawarat and Venugopal V. Veeravalli. Controlled sensing for sequential multihypothesis testing with controlled markovian observations and non-uniform control cost. *Sequential Analysis*, 34(1):1–24, 2015. doi: 10.1080/07474946.2014.961864. URL https://doi.org/10.1080/07474946.2014.961864.

Sirin Nitinawarat, George K. Atia, and Venugopal V. Veeravalli. Controlled sensing for multihypothesis testing. *IEEE Transactions on Automatic Control*, 58(10):2451–2464, 2013. doi: 10.1109/TAC.2013.2261188. URL https://doi.org/10.1109/TAC.2013.2261188.

Gayathri R. Prabhu, Srikrishna Bhashyam, Aditya Gopalan, and Rajesh Sundaresan. Sequential multihypothesis testing in multi-armed bandit problems: An approach for asymptotic optimality. *IEEE Transactions on Information Theory*, 68(7):4790–4817, 2022. doi: 10.1109/TIT.2022.3159600. URL https://doi.org/10.1109/TIT.2022.3159600.

David Siegmund. *Sequential Analysis: Tests and Confidence Intervals*. Springer Series in Statistics. Springer, 1985. doi: 10.1007/978-1-4757-1862-1. URL https://doi.org/10.1007/978-1-4757-1862-1.

Nidhin Koshy Vaidhiyan and Rajesh Sundaresan. Active search with a cost for switching actions. In *2015 Information Theory and Applications Workshop (ITA)*, pp. 17–24, 2015. doi: 10.1109/ITA.2015.7308961. URL https://doi.org/10.1109/ITA.2015.7308961.

Venugopal V. Veeravalli, Georgios Fellouris, and George V. Moustakides. Quickest change detection with controlled sensing. *IEEE Journal on Selected Areas in Information Theory*, 5:1–11, 2024. doi: 10.1109/JSAIT.2024.3362324. URL https://doi.org/10.1109/JSAIT.2024.3362324.

George Vershinin, Asaf Cohen, and Omer Gurewitz. Active sequential hypothesis testing with non-homogeneous costs. In *2026 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2026. doi: 10.1109/ICASSP55912.2026.11463839. URL https://doi.org/10.1109/ICASSP55912.2026.11463839.

Abraham Wald and Jacob Wolfowitz. Optimum character of the sequential probability ratio test. *The Annals of Mathematical Statistics*, 19(3):326–339, 1948. doi: 10.1214/aoms/1177730197. URL https://doi.org/10.1214/aoms/1177730197.

Michael Woodroofe. *Nonlinear Renewal Theory in Sequential Analysis*, volume 39 of *CBMS-NSF Regional Conference Series in Applied Mathematics*. SIAM, 1982. doi: 10.1137/1.9781611970302. URL https://doi.org/10.1137/1.9781611970302.

A Proofs for Sections 3–5

A.1 Proof of Proposition 1

Fix $e \in E$. The function $q \mapsto \Delta_q(e)$ is affine and nonnegative on the convex set $\mathcal{D}_i^\star$. Suppose $\Delta_{q^\circ}(e) = 0$ at some $q^\circ \in \mathrm{i}(\mathcal{D}_i^\star)$. For any $q \in \mathcal{D}_i^\star$, the relative-interior property lets the segment from q through q° extend past q° while remaining in $\mathcal{D}_i^\star$. Hence $q^\circ = \alpha q + (1-\alpha)\widetilde{q}$ for some $\widetilde{q} \in \mathcal{D}_i^\star$ and $\alpha \in (0,1)$. Hence

$$0 = \Delta_{q^\circ}(e) = \alpha\Delta_q(e) + (1-\alpha)\Delta_{\widetilde{q}}(e),$$

and nonnegativity gives $\Delta_q(e) = 0$. The converse follows from the definition of E_i°.

For the second claim, let $w^\star$ be primal optimal and $q = (\lambda, h) \in \mathcal{D}_i^\star$. Flow conservation gives

$$\sum_{e \in E} w_e^\star \Delta_q(e) = C_i^\star - \lambda^\top \mathbf{D}_i w^\star. \tag{50}$$

The left side is nonnegative, while primal feasibility $\mathbf{D}_i w^\star \geq C_i^\star \mathbf{1}$ and $\lambda \in \Delta_r$ make the right side nonpositive. Both are therefore zero. Every edge in $\mathrm{supp}(w^\star)$ is tight for every dual optimizer and hence belongs to E_i°.

For the third claim, $q \mapsto \lambda_j$ is affine and nonnegative on $\mathcal{D}_i^\star$ for each j. If λ_j vanished at a relative-interior point, the same affine-zero argument would force it to vanish on the entire optimal dual set, contradicting the existence of a full-support optimum. Thus every relative-interior optimal multiplier has full support. Applying equation 50 with such a multiplier gives $(\lambda^\circ)^\top \mathbf{D}_i w^\star = C_i^\star$. Since every coordinate of $\mathbf{D}_i w^\star$ is at least $C_i^\star$ and every $\lambda_j^\circ > 0$, all coordinates must equal $C_i^\star$.

A.2 Critical cycle rank and relative interior

Proof of Corollary 2. Let $w^\star$ be any primal optimizer. By Proposition 1, $w^\star \in \mathcal{W}_i^\circ$ and $\mathbf{D}_i^\circ w^\star = C_i^\star \mathbf{1}$, hence $\bar{\mathbf{D}}_i w^\star = 0$. For $x \in \ker B_i^\circ$, set

$$x_0 := x - (\mathbf{1}^\top x) w^\star.$$

Then

$$x_0 \in \ker B_i^\circ \cap \mathbf{1}^\perp, \qquad \bar{\mathbf{D}}_i x_0 = \bar{\mathbf{D}}_i x,$$

which proves equation 15. Since $\mathbf{1}^\top w^\star = 1$, the linear functional $\mathbf{1}^\top$ is nonzero on $\ker B_i^\circ$, and therefore

$$\dim(\ker B_i^\circ \cap \mathbf{1}^\perp) = c_i - 1.$$

This gives equation 16. Finally, $\dim \mathsf{Q}_i = r - 1$, while $\bar{\mathcal{K}}_i^\circ \subseteq \bar{\mathcal{L}}_i^\circ$. If 0 lies in the interior of $\bar{\mathcal{K}}_i^\circ$ relative to Q_i, then $\bar{\mathcal{L}}_i^\circ = \mathsf{Q}_i$. Hence $r - 1 \leq c_i - 1$, proving equation 17. □

Proof of Proposition 3. Let v support $\bar{\mathcal{K}}_i^\circ$ at the origin, with the sign chosen so that $v^\top \mathbf{D}_i^\circ w \geq 0$ for every $w \in \mathcal{W}_i^\circ$. Min-cost circulation duality gives a potential $\psi : \mathcal{Z} \to \mathbb{R}$, normalized by $\psi(z_{\mathrm{ref}}) = 0$, such that

$$a_v(e) := v^\top d_i(e) - [\psi(z') - \psi(z)] \geq 0, \qquad e = (z, a, z') \in E_i^\circ.$$

Fix $q^\circ = (\lambda, h)$. For sufficiently small $\varepsilon > 0$,

$$\lambda_\varepsilon := \lambda - \varepsilon v, \qquad h_\varepsilon := h + \varepsilon\psi$$

has $\lambda_\varepsilon \in \Delta_r$. Its critical-edge slack is $\varepsilon a_v(e)$, while the finitely many positive noncritical slacks remain nonnegative. Thus $(\lambda_\varepsilon, h_\varepsilon)$ is another optimal dual solution. Every critical edge must therefore have $a_v(e) = 0$. Summing around a signed circulation cancels the potential, so v annihilates $\bar{\mathcal{L}}_i^\circ$.

Conversely, suppose that v annihilates this image. For any $w \in \mathcal{W}_i^\circ$, the difference from an equalizing optimizer $w^\star$ is a signed circulation. Since $v^\top \mathbf{D}_i^\circ w^\star = 0$, also $v^\top \mathbf{D}_i^\circ w = 0$. This proves equation 19; equation 18 and the rank criterion follow by finite-dimensional separation. □

A.3 Exact accounting

For the fixed $q^\circ = (\lambda, h)$, the increment of $\mathcal{M}_t$ in equation 22 has conditional mean

$$\lambda^\top d_i(E_t) + h(Z_t) - h(Z_{t-1}) - C_i^\star + \Delta(E_t) = 0.$$

The increments are uniformly bounded, so optional stopping at $T_b \wedge n$ gives

$$\mathbb{E}_i^\pi[\mathcal{M}_{T_b \wedge n}] = 0.$$

If $|\mathcal{M}_t - \mathcal{M}_{t-1}| \leq K$, then

$$|\mathcal{M}_{T_b} - \mathcal{M}_{T_b \wedge n}| \leq K(T_b - n)_+,$$

whose expectation tends to zero when $\mathbb{E}_i^\pi T_b < \infty$. Hence $\mathbb{E}_i^\pi[\mathcal{M}_{T_b}] = 0$. At stopping,

$$\lambda^\top S_{T_b} = \min_j S_{T_b}^j + I_{T_b}^\lambda = b + O_b + I_{T_b}^\lambda.$$

Substituting this identity into $\mathbb{E}_i^\pi[\mathcal{M}_{T_b}] = 0$ proves equation 23. Since $0 \leq O_b \leq L$ and h is bounded on the finite state space, equation 24 follows uniformly in b and over policies with finite expected stopping time.

A.4 Proof of Corollary 5

The lower bound follows from equation 23: the terms $I_{T_b}^\lambda, D_{T_b}, O_b$ are nonnegative, and the potential term is uniformly bounded.

For the upper bound, fix a primal optimizer $w^\star \in \mathcal{W}(G)$, so $\mathbf{D}_i w^\star \geq C_i^\star \mathbf{1}$. Write its finite cycle decomposition as

$$w^\star = \sum_{\ell=1}^{m} \alpha_\ell w^{c_\ell}, \qquad \alpha_\ell \geq 0, \quad \sum_\ell \alpha_\ell = 1,$$

where w^{c_ℓ} is the normalized edge-frequency vector of the directed cycle c_ℓ.

Set $H_b = \lceil \sqrt{b} \rceil$ and use the initial state as the root: $z_{\mathrm{r}} := Z_0$. For each cycle, fix root-to-cycle and cycle-to-root paths. During one macro-tour repeat c_ℓ

$$n_\ell := \left\lfloor \frac{\alpha_\ell H_b}{|c_\ell|} \right\rfloor$$

times, inserting the fixed connectors. Since the graph, cycle decomposition, and connectors are fixed,

$$L_b = H_b + O(1)$$

for the macro-tour length, and its mean evidence increment satisfies

$$\mathbb{E}_i[\Delta S_{\mathrm{mac}}] = H_b \mathbf{D}_i w^\star + O(1) \geq C_i^\star H_b \mathbf{1} - O(1).$$

Thus for some $K < \infty$, the mean rate at tour boundaries is at least

$$C_{i,b} := C_i^\star - \frac{K}{H_b} = C_i^\star - O(b^{-1/2})$$

in every coordinate.

Repeat the tour indefinitely and let $t_k = kL_b$. Over an $O(b)$ horizon, the $O(b^{-1/2})$ rate loss contributes only $O(\sqrt{b})$ missing mean evidence. Choose A large enough to absorb this loss, and let k_b be the first boundary with

$$t_{k_b} \geq b/C_i^\star + A\sqrt{b}.$$

For all alternatives, the mean evidence at t_{k_b+m} exceeds b by at least a constant multiple of $(m+1)H_b$. Because the edge sequence is deterministic, the likelihood increments are independent and bounded. Hoeffding's inequality and a union bound give constants $c, C > 0$ such that

$$\mathbb{P}_i(T_b > t_{k_b+m}) \leq C \exp\left\{-c\frac{(m+1)^2 H_b^2}{b + (m+1)H_b}\right\}. \tag{51}$$

Pathwise,

$$(T_b - t_{k_b})_+ \leq L_b \sum_{m \geq 0} \mathbf{1}\{T_b > t_{k_b+m}\}.$$

For $m \leq \sqrt{b}$, the exponent in equation 51 is at most $-c'(m+1)^2$; for $m > \sqrt{b}$, it is at most $-c'(m+1)\sqrt{b}$. The tail sum is therefore $O(1)$. Since $L_b = O(\sqrt{b})$ and $t_{k_b} = b/C_i^\star + O(\sqrt{b})$,

$$\mathbb{E}_i T_b \leq \frac{b}{C_i^\star} + O(\sqrt{b}).$$

A.5 Proof of Theorem 6(a)

Let

$$P_\lambda := I - \mathbf{1}\lambda^\top, \qquad \mathsf{H}_\lambda := \{x : \lambda^\top x = 0\}, \qquad R_t := P_\lambda S_t.$$

Since $\lambda^\top \mathbf{1} = 1$, P_λ identifies the quotient Q_i with H_λ, and

$$I_t^\lambda = -\min_j R_t^j \leq C_I \|R_t\|$$

for some $C_I < \infty$.

The interior condition implies that some $\eta > 0$ satisfies

$$\min_{w \in \mathcal{W}_i^\circ} \left\langle \frac{y}{\|y\|}, P_\lambda \mathbf{D}_i^\circ w \right\rangle \leq -\eta \qquad (y \neq 0).$$

A linear functional on $\mathcal{W}_i^\circ$ attains its minimum at an extreme point; each extreme point is a normalized directed-cycle circulation. Fix $z_{\mathrm{r}} = Z_0$. For each simple critical cycle, choose critical paths from z_{r} to the cycle and back. At the beginning of a block with current relative evidence $y \neq 0$, choose a minimizing cycle, travel to it, repeat it N times, and return to the root. Since the set of cycles and connectors is finite, their total connector drift is uniformly bounded. Taking N sufficiently large gives some $\delta_0 > 0$ such that the mean block increment $\mu(y)$ satisfies

$$\left\langle \frac{y}{\|y\|}, \mu(y) \right\rangle \leq -2\delta_0, \qquad y \neq 0. \tag{52}$$

At $y = 0$, choose any fixed critical cycle. All edges used by the policy are critical, so $D_t = 0$ identically.

Let $0 = \tau_0 < \tau_1 < \cdots$ be the root-to-root block boundaries and $Y_k := R_{\tau_k}$. There are constants $L_B, J_B < \infty$ such that

$$1 \leq \tau_{k+1} - \tau_k \leq L_B, \qquad \sup_{\tau_k \leq t \leq \tau_{k+1}} \|R_t - Y_k\| \leq J_B. \tag{53}$$

Let $\Xi_{k+1} := Y_{k+1} - Y_k$. On $\{Y_k \neq 0\}$, with $u_k := Y_k/\|Y_k\|$,

$$\mathbb{E}_i[\langle u_k, \Xi_{k+1} \rangle \mid \mathcal{F}_{\tau_k}] \leq -2\delta_0.$$

For large $\|y\|$ and $\|\xi\| \le J_B$,

$$\|y+\xi\| \le \|y\| + \left\langle \frac{y}{\|y\|}, \xi \right\rangle + C\frac{J_B^2}{\|y\|}.$$

It follows that the radial drift is at most $-\delta_0$ outside a fixed ball. The radial increments are bounded, so Hoeffding's lemma yields, for sufficiently small $\alpha > 0$, constants $\rho < 1, C < \infty$ such that

$$\mathbb{E}_i[e^{\alpha\|Y_{k+1}\|} \mid \mathcal{F}_{\tau_k}] \le \rho e^{\alpha\|Y_k\|} + C.$$

Iteration yields

$$\sup_{k\ge 0} \mathbb{E}_i[e^{\alpha\|Y_k\|}] < \infty. \tag{54}$$

Define

$$K_b := \min\{k \ge 1 : T_b \le \tau_k\}.$$

By equation 53,

$$I^\lambda_{T_b} \le C_I(\|Y_{K_b-1}\| + J_B).$$

Choose $K_0 > 1/C_i^\star$ and $m_b = \lceil K_0 b \rceil$. On $\{K_b \le m_b\}$,

$$I^\lambda_{T_b} \le C(1 + \max_{k<m_b} \|Y_k\|).$$

The exponential moment bound gives

$$\mathbb{P}_i(\max_{k<m} \|Y_k\| > x) \le Cme^{-\alpha x},$$

and therefore

$$\mathbb{E}_i \max_{k<m_b} \|Y_k\| = O(\log b). \tag{55}$$

To control the number of blocks before certification, define at block boundaries

$$Z_m := \lambda^\top S_{\tau_m} - C_i^\star \tau_m$$

is a martingale with uniformly bounded increments, because every block returns to the same root and uses only zero-slack edges. If $K_b > m$, then $\lambda^\top S_{\tau_m} < b + I^\lambda_{\tau_m}$. Since $\tau_m \ge m$, with $x_m = (C_i^\star m - b)/2$,

$$\{K_b > m\} \subseteq \{Z_m \le -x_m\} \cup \{I^\lambda_{\tau_m} > x_m\}.$$

For $m \ge K_0 b$, $x_m \ge c_0 m$. Azuma–Hoeffding and equation 54 give

$$\mathbb{P}_i(K_b > m) \le Ce^{-cm}, \qquad m \ge K_0 b. \tag{56}$$

Thus $\mathbb{E}_i K_b < \infty$, and $T_b \le \tau_{K_b} \le L_B K_b$ implies $\mathbb{E}_i T_b < \infty$.

For the late-block contribution, write

$$\begin{aligned}&\mathbb{E}_i[I^\lambda_{T_b}\mathbf{1}\{K_b > m_b\}] \\ &\quad\le C\sum_{k>m_b} \mathbb{E}_i[(1+\|Y_{k-1}\|)\mathbf{1}\{K_b = k\}].\end{aligned}$$

Cauchy–Schwarz, equation 54, and equation 56 make the sum $O(e^{-c'b})$. Together with equation 55,

$$\mathbb{E}_i[I^\lambda_{T_b}] = O(\log b), \qquad D_{T_b} = 0.$$

The exact accounting theorem now applies and proves part (a).

A.6 Proof of Theorem 6(b)

Fix $v \in \mathcal{U}_i^\circ \setminus \{0\}$ with $\nu_0 := \underline{\nu}_i(v) > 0$. Since $\mathbf{D}_i^{\circ\top} v \in \mathrm{im}(B_i^{\circ\top})$, there is a potential $\psi : \mathcal{Z} \to \mathbb{R}$ such that

$$v^\top d_i(e) = \psi(z') - \psi(z), \qquad e = (z, a, z') \in E_i^\circ. \tag{57}$$

For a noncritical edge $e = (z, a, z')$, set

$$\widetilde{a}_v(e) := v^\top d_i(e) - [\psi(z') - \psi(z)].$$

Finiteness of the edge set and strict positivity of the noncritical slacks give

$$\left| \sum_{s=1}^{t} \widetilde{a}_v(E_s) \mathbf{1}\{E_s \notin E_i^\circ\} \right| \le \rho_v D_t \tag{58}$$

for some $\rho_v < \infty$.

Let

$$M_t^v := \sum_{s=1}^{t} v^\top [X_{E_s}(Y_s) - d_i(E_s)].$$

Equations equation 57 and equation 58 imply

$$\left| v^\top S_t - M_t^v - [\psi(Z_t) - \psi(Z_0)] \right| \le \rho_v D_t. \tag{59}$$

Moreover, with $m_t := \min_j S_t^j$ and $\lambda_{\min} := \min_j \lambda_j > 0$,

$$|v^\top S_t| = |v^\top (S_t - m_t \mathbf{1})| \le \frac{\|v\|_1}{\lambda_{\min}} I_t^\lambda. \tag{60}$$

Suppose, toward a contradiction, that there are $b_n \to \infty$ and finite-mean causal policies π_n such that

$$\mathbb{E}_i^{\pi_n}[I_{T_{b_n}}^\lambda] + \mathbb{E}_i^{\pi_n}[D_{T_{b_n}}] = o(\sqrt{b_n}). \tag{61}$$

If noncritical edges exist, let

$$\Delta_{\min} := \min_{e \notin E_i^\circ} \Delta(e) > 0.$$

Then $N_{\mathrm{nc}}(t) \le D_t / \Delta_{\min}$; if every edge is critical, set $N_{\mathrm{nc}}(t) = 0$. Hence

$$\mathbb{E}_i^{\pi_n} N_{\mathrm{nc}}(T_{b_n}) = o(\sqrt{b_n}). \tag{62}$$

Exact accounting also yields

$$\mathbb{E}_i^{\pi_n} T_{b_n} = \frac{b_n}{C_i^\star} + o(\sqrt{b_n}). \tag{63}$$

We next lower-bound the quadratic variation of M^v. Split an arbitrary trajectory into maximal critical segments. Within each segment, delete directed cycles until a simple path remains. Every deleted cycle, normalized by its length, belongs to $\mathcal{W}_i^\circ$, so its average v-variance is at least ν_0. Each residual simple path has at most $|\mathcal{Z}| - 1$ edges. It follows that, for some $K_v < \infty$,

$$\langle M^v \rangle_t \ge \nu_0 t - K_v \big(N_{\mathrm{nc}}(t) + 1\big). \tag{64}$$

For the remainder of the proof suppress n and write b, π, T_b. Extend the process after T_b, only for analysis, by following critical edges forever. Such a continuation exists because G_i° is strongly connected, and it adds no noncritical edge. Fix $A > 0$ and define

$$n_b := \frac{b}{C_i^\star}, \qquad t_b^- := \lfloor n_b - A\sqrt{b} \rfloor, \qquad t_b^+ := \lceil n_b + A\sqrt{b} \rceil.$$

Equations equation 62 and equation 64 give, uniformly along the contradicting sequence,

$$\mathbb{E}_i \langle M^v \rangle_{t_b^-} \ge c_1 b$$

for all large b. The increments of M^v are uniformly bounded, so Burkholder's inequality gives

$$\mathbb{E}_i[(M_{t_b^-}^v)^2] \ge c_1 b, \qquad \mathbb{E}_i[(M_{t_b^-}^v)^4] \le c_2 b^2.$$

Paley–Zygmund applied to $(M^v_{t_b^-})^2$ therefore yields constants $c_3, p_0 > 0$ such that

$$\mathbb{P}_i(|M^v_{t_b^-}| \geq c_3\sqrt{b}) \geq p_0. \tag{65}$$

It remains to show that this fluctuation survives until certification. If $1 \leq x \leq n_b/\sqrt{b}$ and $T_b \leq n_b - x\sqrt{b}$, then for some $s \leq n_b - x\sqrt{b}$, $\lambda^\top S_s \geq b$. The dual decomposition equation 22 and $D_s \geq 0$ imply

$$\mathcal{M}_s \geq b - C_i^\star s - \mathrm{span}(h) \geq C_i^\star x\sqrt{b} - \mathrm{span}(h).$$

Maximal Azuma–Hoeffding gives

$$\mathbb{P}_i(T_b \leq n_b - x\sqrt{b}) \leq Ce^{-cx^2}. \tag{66}$$

After integration, $\mathbb{E}_i[(n_b - T_b)_+] = O(\sqrt{b})$. Combining this with equation 63 yields $\mathbb{E}_i[(T_b - n_b)_+] = O(\sqrt{b})$. Thus

$$\mathbb{P}_i(T_b < t_b^-) \leq Ce^{-cA^2}, \qquad \mathbb{P}_i(T_b > t_b^+) \leq C/A + o(1).$$

Choose A so that both terms are below $p_0/8$.

The interval $[t_b^-, t_b^+]$ has $O(\sqrt{b})$ steps. Doob's L^2 maximal inequality and the bounded conditional variances give

$$\mathbb{P}_i\left(\max_{t_b^- \leq t \leq t_b^+} |M^v_t - M^v_{t_b^-}| \geq \frac{c_3}{2}\sqrt{b}\right) = O(b^{-1/2}). \tag{67}$$

Together with equation 65 and the two location bounds, this implies

$$\liminf_{b\to\infty} \mathbb{P}_i\left(|M^v_{T_b}| \geq \frac{c_3}{2}\sqrt{b}\right) > 0. \tag{68}$$

On the other hand, equation 59 and equation 60 give

$$|M^v_{T_b}| \leq \frac{\|v\|_1}{\lambda_{\min}} I^\lambda_{T_b} + \rho_v D_{T_b} + \mathrm{span}(\psi).$$

Its expectation is $o(\sqrt{b})$ by equation 61, which contradicts equation 68 by Markov's inequality. Therefore equation 32 holds. Exact accounting and Corollary 5 give the matching lower and upper bounds in Theorem 6(b).

B Sharp Flex/Lock Certification Asymptotics

We prove Theorem 8 under H_0 in the Lock system. Each access state has a unique admissible action, so the experiment sequence is deterministic up to the fixed initial phase.

For $j = 1, \ldots, r$, let $N_j(t)$ denote the number of executions of a_j by physical time t. The deterministic rotor satisfies

$$\left|N_j(t) - \frac{t}{r}\right| \leq 1 \qquad \text{for every } t \geq 0,\ j = 1, \ldots, r. \tag{69}$$

Let

$$\xi := \log \frac{d\,\mathrm{Bern}(p)}{d\,\mathrm{Bern}(q)}(Y), \qquad Y \sim \mathrm{Bern}(p).$$

Then $\mathbb{E}_0[\xi] = d$, $\mathrm{Var}_0(\xi) = \sigma^2$, and ξ is bounded. For each j, let $(\xi_{j,m})_{m\geq 1}$ be the successive increments observed when a_j is executed. Under H_0, these sequences are independent across j and i.i.d. within each coordinate, so

$$S^j_t = \sum_{m=1}^{N_j(t)} \xi_{j,m}. \tag{70}$$

B.1 Baseline Gaussian fluctuations

Set

$$n_b := \left\lfloor \frac{rb}{d} \right\rfloor. \tag{71}$$

By equation 69, $N_j(n_b) = b/d + O(1)$ and $dN_j(n_b) - b = O(1)$, uniformly in j.

Lemma 11. *As $b \to \infty$,*

$$\left(\frac{S^j_{n_b} - b}{\sqrt{b}} \right)_{j=1}^{r} \Longrightarrow \frac{\sigma}{\sqrt{d}}(Z_1, \dots, Z_r), \tag{72}$$

where $Z_1, \dots, Z_r$ are independent standard Gaussian variables.

Proof. For each j,

$$\frac{S^j_{n_b} - b}{\sqrt{b}} = \frac{\sum_{m=1}^{N_j(n_b)} (\xi_{j,m} - d)}{\sqrt{b}} + \frac{dN_j(n_b) - b}{\sqrt{b}}.$$

The second term is $o(1)$, while $N_j(n_b)/b \to 1/d$. The central limit theorem gives each coordinate limit, and independence of the coordinate streams gives the joint convergence. □

Define

$$x^\star_b := -\frac{r}{d} \min_{1 \le j \le r} \frac{S^j_{n_b} - b}{\sqrt{b}}. \tag{73}$$

Lemma 11 and the continuous mapping theorem give

$$x^\star_b \Longrightarrow -\frac{r\sigma}{d^{3/2}} \min_{1 \le j \le r} Z_j. \tag{74}$$

In particular, $(x^\star_b)$ is tight.

B.2 Local uniform linearization

Lemma 12. *For every fixed $K < \infty$,*

$$\max_{1 \le j \le r} \sup_{|x| \le K} \left| \frac{S^j_{\lfloor n_b + x\sqrt{b} \rfloor} - b}{\sqrt{b}} - \left[\frac{S^j_{n_b} - b}{\sqrt{b}} + \frac{d}{r} x \right] \right| \xrightarrow{\mathbb{P}_0} 0. \tag{75}$$

Proof. Fix K. Uniformly over $|x| \le K$ and j,

$$N_j(\lfloor n_b + x\sqrt{b} \rfloor) - N_j(n_b) = \frac{x\sqrt{b}}{r} + O(1). \tag{76}$$

The deterministic mean contribution divided by $\sqrt{b}$ is therefore $(d/r)x + o(1)$, uniformly over the same range.

Across the entire local window, each coordinate gains or loses at most $C_K\sqrt{b}$ observations. Since $\xi - d$ is bounded, a maximal Hoeffding bound gives, for every $\varepsilon > 0$,

$$\mathbb{P}_0 \left(\max_{1 \le j \le r} \sup_{|x| \le K} \left| \sum_{\ell \in \mathcal{I}_{b,j}(x)} (\xi_{j,\ell} - d) \right| > \varepsilon\sqrt{b} \right) \le C \exp\{-c\varepsilon^2\sqrt{b}\},$$

where $\mathcal{I}_{b,j}(x)$ indexes the observations added for $x \ge 0$ and removed for $x < 0$. This proves the claim. □

B.3 First-crossing approximation

Define

$$X_b := \frac{T_{b,\mathrm{Lock}} - n_b}{\sqrt{b}}. \tag{77}$$

Because the likelihood coordinates are not monotone, a local linearization alone does not identify the first crossing. We first rule out crossings far before the local window.

Lemma 13. *There exist constants $c, C > 0$ such that, for every $x \geq 1$ and all sufficiently large b,*

$$\mathbb{P}_0\left(T_{b,\mathrm{Lock}} \leq n_b - x\sqrt{b}\right) \leq Ce^{-cx^2}. \tag{78}$$

Proof. If $x > n_b/\sqrt{b}$, the event is empty. Thus only $1 \leq x \leq n_b/\sqrt{b}$ needs to be considered. Let

$$\overline{S}_t := \frac{1}{r}\sum_{j=1}^{r} S_t^j.$$

Exactly one coordinate is updated at each physical time, so

$$\overline{S}_t - \frac{d}{r}t$$

is a martingale with independent, uniformly bounded increments. If $T_{b,\mathrm{Lock}} \leq t$, then for some $s \leq t$, $\overline{S}_s \geq b$. With $t = \lfloor n_b - x\sqrt{b} \rfloor$,

$$\max_{s \leq t}\left(\overline{S}_s - \frac{d}{r}s\right) \geq \frac{d}{r}x\sqrt{b} - O(1).$$

The maximal Hoeffding inequality proves the bound. □

Lemma 14. *As $b \to \infty$,*

$$X_b - x_b^\star \xrightarrow{\mathbb{P}_0} 0. \tag{79}$$

Proof. Fix $\varepsilon > 0$. Tightness of $x_b^\star$ allows K to be chosen so that

$$\limsup_{b\to\infty} \mathbb{P}_0(|x_b^\star| > K/2) < \varepsilon.$$

Increasing K if needed, Lemma 13 also gives

$$\limsup_{b\to\infty} \mathbb{P}_0(X_b < -K) < \varepsilon.$$

Let $c_\star = d/r$, and consider the event on which $|x_b^\star| \leq K/2$ and the uniform error in Lemma 12 on $[-K, K]$ is at most $c_\star\varepsilon/4$. This event has probability $1 - 2\varepsilon - o(1)$.

Let $j_b^\star$ attain the minimum in equation 73. Then

$$\frac{S_{n_b}^{j_b^\star} - b}{\sqrt{b}} + c_\star x_b^\star = 0.$$

For any $x \leq x_b^\star - \varepsilon$ with $|x| \leq K$, Lemma 12 gives

$$\frac{S_{\lfloor n_b + x\sqrt{b} \rfloor}^{j_b^\star} - b}{\sqrt{b}} < 0$$

for all sufficiently large b. Hence certification cannot occur in the local window before $x_b^\star - \varepsilon$, and the early-crossing event controls times before that window. Thus

$$X_b \geq x_b^\star - \varepsilon - o(1).$$

At $x = x_b^\star + \varepsilon$, every coordinate j satisfies

$$\frac{S_{n_b}^{j} - b}{\sqrt{b}} + c_\star x \geq c_\star\varepsilon.$$

The same uniform approximation gives $S_{\lfloor n_b + (x_b^\star + \varepsilon)\sqrt{b} \rfloor}^{j} \geq b$ for every j, whence

$$X_b \leq x_b^\star + \varepsilon + o(1).$$

Since $\varepsilon > 0$ is arbitrary, the result follows. □

Lemma 14, equation 74, and Slutsky's theorem yield

$$\frac{T_{b,\mathrm{Lock}} - \frac{rb}{d}}{\sqrt{b}} \Longrightarrow -\frac{r\sigma}{d^{3/2}} \min_{1\le j\le r} Z_j, \tag{80}$$

because $n_b - rb/d = O(1)$. This proves equation 40.

B.4 Uniform integrability and the mean expansion

Lemma 15. *The family $(X_b)_{b\ge 1}$ is uniformly integrable. In fact, for some $b_0 < \infty$,*

$$\sup_{b\ge b_0} \mathbb{E}_0[|X_b|^2] < \infty. \tag{81}$$

Proof. Lemma 13 controls the negative tail:

$$\mathbb{P}_0(X_b < -x) \le Ce^{-cx^2}, \qquad x \ge 1.$$

For the positive tail, let $t_b(x) = \lfloor n_b + x\sqrt{b} \rfloor$. If $X_b > x$, then at least one coordinate is below b at $t_b(x)$. Hence

$$\mathbb{P}_0(X_b > x) \le \sum_{j=1}^{r} \mathbb{P}_0(S^j_{t_b(x)} < b). \tag{82}$$

By equation 69,

$$dN_j(t_b(x)) - b = \frac{d}{r} x\sqrt{b} + O(1).$$

Hoeffding's inequality applied to equation 70 gives

$$\mathbb{P}_0(X_b > x) \le C \exp\left\{ -c \frac{x^2}{1 + x/\sqrt{b}} \right\}, \qquad x \ge 1. \tag{83}$$

Using

$$\mathbb{E}[(X_b^+)^2] = 2\int_0^\infty x\, \mathbb{P}(X_b > x)\, dx,$$

split the integral at $x = \sqrt{b}$. The first part is controlled by $Ce^{-c'x^2}$, while for $x > \sqrt{b}$,

$$\frac{x^2}{1 + x/\sqrt{b}} \ge c' x\sqrt{b}.$$

The negative tail is integrable by the preceding Gaussian bound. Hence equation 81 holds, and (X_b) is uniformly integrable. □

By equation 80 and Lemma 15,

$$\frac{\mathbb{E}_0[T_{b,\mathrm{Lock}}] - \frac{rb}{d}}{\sqrt{b}} \longrightarrow -\frac{r\sigma}{d^{3/2}} \mathbb{E}\left[\min_j Z_j\right] = \kappa_r.$$

Therefore

$$\mathbb{E}_0[T_{b,\mathrm{Lock}}] = \frac{rb}{d} + \kappa_r\sqrt{b} + o(\sqrt{b}). \tag{84}$$

Lock has no control choice, so this is $V^{\mathrm{cert}}_{0,\mathrm{Lock}}(b)$.

For Flex, expose independent increment streams $(\xi_{j,m})_{m\ge 1}$, one for each experiment, and let $\tau_b^{(j)}$ be the first index at which the j-th partial sum reaches b. If a policy certifies at time T_b, then $N_j(T_b) \ge \tau_b^{(j)}$ for every j, and therefore

$$T_b = \sum_{j=1}^{r} N_j(T_b) \ge \sum_{j=1}^{r} \tau_b^{(j)}.$$

Flex attains equality by sampling the coordinates successively and never returning to a coordinate after its first crossing. Hence $V^{\mathrm{cert}}_{0,\mathrm{Flex}}(b) = r\mathbb{E}_0\tau_b$. If L_+ bounds the positive part of one increment, the overshoot lies in $[0, L_+]$. Wald's identity gives

$$b \le d\mathbb{E}_0\tau_b \le b + L_+,$$

which proves equation 41. Since a bounded remainder is $o(\sqrt{b})$, equation 43 follows. This completes the proof of Theorem 8.

C Proof of the Routing Separation

Both routing graphs have four edges and three vertices and are strongly connected, so their circulation spaces have dimension two. In the Split system, every stationary flow is a mixture of the two root-to-root cycles. If the (a_1, a_1) cycle has mixture weight u, the aggregate action frequencies are $(u, 1-u)$, and the first-order objective is $d\min\{u, 1-u\}$. Its unique maximizer is $u = 1/2$, with value $d/2$. The equal mixture has full edge support, so every edge is critical. Varying u generates the full one-dimensional relative-likelihood space.

In the Coupled system, every root-to-root cycle contains one use of each action. Hence every stationary flow has aggregate action frequencies $(1/2, 1/2)$, value $d/2$, and zero relative drift. A full-support stationary flow again exists, so every edge is critical, but the circulation image is $\{0\}$. This proves equation 44 and all first-order claims in Proposition 9.

For the Split upper bound, inspect the process only at root returns. Repeating one cycle gives a two-increment random walk of mean $2d$. Run the first cycle type until coordinate one is at least b at a return to the root, then do the same for coordinate two. Each batch overshoot is bounded, so Wald's identity gives total expected physical time $2b/d + O(1)$. The universal lower bound supplies the matching statement equation 45.

For Coupled, expose independent streams $(\xi_{j,m})_{m\geq 1}$ for the two coordinates. Regardless of the root decisions, each completed round consumes one increment from each stream, and at an intermediate time the two sample counts differ by at most one. Relative to the deterministic rotor used in Appendix B, every coordinate sum at a fixed physical time therefore differs by at most one bounded increment. In particular, at $n_b = \lfloor 2b/d \rfloor$,

$$\left(\frac{S^j_{n_b} - b}{\sqrt{b}}\right)^2_{j=1} \Longrightarrow \frac{\sigma}{\sqrt{d}}(Z_1, Z_2)$$

uniformly over causal routing policies. The same bounded count discrepancy preserves the local uniform linearization, with slope $d/2$. The early crossing bound uses only the averaged evidence martingale and is likewise policy-uniform. Writing $x^\star_b$ for the local Gaussian root defined as in equation 73, these estimates give, for every $\varepsilon > 0$,

$$\sup_\pi \mathbb{P}^\pi_0\left(\left|\frac{T^\pi_b - n_b}{\sqrt{b}} - x^\star_b\right| > \varepsilon\right) \longrightarrow 0.$$

The tail estimates in Lemma 15 also depend only on $|N_j(t) - t/2| \leq 1$, and therefore give

$$\sup_{b\geq b_0}\sup_\pi \mathbb{E}^\pi_0\left[\left|\frac{T^\pi_b - n_b}{\sqrt{b}}\right|^2\right] < \infty.$$

Uniform integrability now yields, uniformly over causal routing policies,

$$\frac{T_b - 2b/d}{\sqrt{b}} \Longrightarrow -\frac{2\sigma}{d^{3/2}}\min\{Z_1, Z_2\}$$

with uniform convergence of the corresponding means. Taking the infimum over policies proves equation 46.

D Fixed-Confidence Translation

We prove Corollary 10. Throughout,

$$B_\delta := \log\frac{1}{\delta}, \qquad \beta_\delta := B_\delta + \log r,$$

and $\delta < 1/2$.

D.1 First-order lower bound

Let $(\pi, \tau, \widehat{\theta})$ be any δ-correct procedure under either access system. By definition, $\tau < \infty$ almost surely under every hypothesis and $\widehat{\theta}$ is $\mathcal{F}_\tau$-measurable. The lower bound is trivial if $\mathbb{E}_0[\tau] = \infty$, so assume $\mathbb{E}_0[\tau] < \infty$.

For each $j = 1, \dots, r$,

$$S_t^j = \log \frac{d\mathbb{P}_0|_{\mathcal{F}_t}}{d\mathbb{P}_j|_{\mathcal{F}_t}}.$$

Data processing through the terminal decision gives

$$\mathbb{E}_0[S_\tau^j] \geq \mathrm{kl}\left(\mathbb{P}_0(\widehat{\theta} = 0), \mathbb{P}_j(\widehat{\theta} = 0)\right) \geq \mathrm{kl}(1-\delta, \delta), \tag{85}$$

where

$$\mathrm{kl}(u, v) := u \log \frac{u}{v} + (1-u) \log \frac{1-u}{1-v}.$$

Moreover,

$$\mathrm{kl}(1-\delta, \delta) = B_\delta - o(1).$$

In the Flex/Lock family,

$$\mathbb{E}_0[S_\tau^j] = d\, \mathbb{E}_0[N_j(\tau)]. \tag{86}$$

Since $\sum_j N_j(\tau) = \tau$, summing equation 85 gives

$$\mathbb{E}_0[\tau] \geq \frac{r}{d} \mathrm{kl}(1-\delta, \delta) = \frac{r}{d} B_\delta - o(1). \tag{87}$$

D.2 Near-certification for Lock

Fix $a > 2$ and let

$$\ell_\delta := B_\delta - a \log B_\delta. \tag{88}$$

Define

$$\mathcal{G}_\delta := \left\{ \widehat{\theta} = 0, \quad \min_{1 \leq j \leq r} S_\tau^j \geq \ell_\delta \right\}. \tag{89}$$

Lemma 16. *For every δ-correct Lock procedure,*

$$\mathbb{P}_0(\mathcal{G}_\delta^c) \leq \delta + r B_\delta^{-a}. \tag{90}$$

On $\mathcal{G}_\delta$,

$$T_{\ell_\delta, \mathrm{Lock}} \leq \tau. \tag{91}$$

Proof. For each $j \geq 1$, the stopped likelihood-ratio identity gives

$$\begin{aligned} \mathbb{P}_0(\widehat{\theta} = 0,\, S_\tau^j < \ell_\delta) &= \mathbb{E}_j[e^{S_\tau^j} \mathbf{1}\{\widehat{\theta} = 0,\, S_\tau^j < \ell_\delta\}] \\ &\leq e^{\ell_\delta} \mathbb{P}_j(\widehat{\theta} = 0) \leq B_\delta^{-a}. \end{aligned}$$

A union bound over j, together with $\mathbb{P}_0(\widehat{\theta} \neq 0) \leq \delta$, proves equation 90. On $\mathcal{G}_\delta$, all coordinates exceed ℓ_δ at time τ, so $T_{\ell_\delta, \mathrm{Lock}} \leq \tau$. □

Lemma 15 implies

$$\mathbb{E}_0[T_{b,\mathrm{Lock}}^2] = O(b^2). \tag{92}$$

Therefore, by Cauchy–Schwarz and Lemma 16,

$$\begin{aligned} \mathbb{E}_0[\tau] &\geq \mathbb{E}_0[T_{\ell_\delta, \mathrm{Lock}} \mathbf{1}_{\mathcal{G}_\delta}] \\ &\geq \mathbb{E}_0[T_{\ell_\delta, \mathrm{Lock}}] - \left(\mathbb{E}_0[T_{\ell_\delta, \mathrm{Lock}}^2]\right)^{1/2} \mathbb{P}_0(\mathcal{G}_\delta^c)^{1/2} \\ &= \mathbb{E}_0[T_{\ell_\delta, \mathrm{Lock}}] - o(\sqrt{B_\delta}). \end{aligned} \tag{93}$$

Using Theorem 8 at $b = \ell_\delta$, together with $B_\delta - \ell_\delta = O(\log B_\delta) = o(\sqrt{B_\delta})$, yields uniformly over all δ-correct Lock procedures

$$\mathbb{E}_0[\tau] \geq \frac{r}{d} B_\delta + \kappa_r \sqrt{B_\delta} + o(\sqrt{B_\delta}). \tag{94}$$

D.3 LIKELIHOOD-RATIO STOPPING AND FALLBACK

For $k \neq j$, let

$$\Lambda_t^{k,j} := \log \frac{d\mathbb{P}_k|_{\mathcal{F}_t}}{d\mathbb{P}_j|_{\mathcal{F}_t}}. \tag{95}$$

Define

$$\tau_\beta := \inf\left\{ t \geq 0 : \exists k \in \Theta \text{ such that } \min_{j \neq k} \Lambda_t^{k,j} \geq \beta \right\}, \tag{96}$$

and output any such k.

For any predictable sampling policy, this rule is δ-correct whenever $\tau_{\beta_\delta} < \infty$ almost surely under every hypothesis. Under the true hypothesis θ, $\exp\{\Lambda_t^{k,\theta}\}$ is a nonnegative martingale for each $k \neq \theta$; Ville's inequality and a union bound give

$$\mathbb{P}_\theta(\widehat{\theta} \neq \theta) \leq r e^{-\beta_\delta} = \delta.$$

For termination we use deterministic round-robin sampling as a fallback. Define

$$\gamma := \min_{\theta \neq k} \frac{1}{r} \sum_{j=1}^{r} D_{\mathrm{KL}}(P_\theta^{a_j} \| P_k^{a_j}). \tag{97}$$

For the Flex/Lock family, $p \neq q$ implies $\gamma > 0$. Pairwise log-likelihood increments are uniformly bounded over this finite Bernoulli family.

Fallback bound. Fix $C_0 < \infty$. Suppose round-robin sampling starts from a history with

$$\Lambda^{\theta,k} \geq -C_0 B_\delta \qquad \text{for every } \theta \neq k.$$

Under every true θ, the expected additional time to τ_{β_δ} is then $O(B_\delta)$, uniformly over such histories, and the stopping time is almost surely finite. Inspect the process at round-robin cycle boundaries. For each $k \neq \theta$, one cycle gives a bounded independent increment to $\Lambda^{\theta,k}$ with mean at least $r\gamma$. After $m_\delta = \lceil C_1 B_\delta \rceil$ cycles, for C_1 large enough, the mean is above β_δ by $\Omega(B_\delta)$; after another s cycles the margin is $\Omega(B_\delta + s)$. Hoeffding's inequality and a union bound therefore give

$$\mathbb{P}_\theta(\tau_{\beta_\delta} > r(m_\delta + s)) \leq C e^{-c(B_\delta + s)} \leq C e^{-cs}.$$

Summing the tail proves the claim.

D.4 LOCK ACHIEVABILITY

Under H_0 in Lock,

$$\Lambda_t^{0,j} = S_t^j, \qquad j = 1, \ldots, r.$$

Hence

$$\tau_{\beta_\delta} \leq T_{\beta_\delta, \mathrm{Lock}}. \tag{98}$$

The forced Lock schedule is round robin up to phase, so the preceding bound guarantees almost-sure termination. The global LR rule is therefore δ-correct. Since $\beta_\delta = B_\delta + O(1)$, Theorem 8 gives

$$V_{0,\mathrm{Lock}}^{\mathrm{fc}}(\delta) \leq \frac{r}{d} B_\delta + \kappa_r \sqrt{B_\delta} + o(\sqrt{B_\delta}). \tag{99}$$

Together with equation 94, this proves equation 48.

D.5 FLEX ACHIEVABILITY

Monitor the global likelihood-ratio rule from time zero. In Flex, run the coordinatewise policy from the proof of Theorem 8: sample a_1 until $\Lambda^{0,1}$ first reaches β_δ, freeze that coordinate, and continue through a_r. If this certification procedure finishes, the global rule has stopped and declared 0. Under H_0, its uncapped duration is

$$T_{\beta_\delta}^{\mathrm{seq}} = \sum_{j=1}^{r} \tau_{\beta_\delta}^{(j)},$$

and equation 41 gives

$$\mathbb{E}_0[T^{\mathrm{seq}}_{\beta_\delta}] = \frac{r}{d}B_\delta + O(1), \tag{100}$$

because $\beta_\delta = B_\delta + \log r$.

Cap this phase at ΓB_δ physical steps and switch permanently to round robin if the global rule has not stopped. A bounded positive-drift random walk has an exponential upper tail for its passage time; a union bound over the r coordinates therefore gives, for sufficiently large Γ,

$$\mathbb{P}_0\Big(T^{\mathrm{seq}}_{\beta_\delta} > \Gamma B_\delta\Big) \le Ce^{-cB_\delta}. \tag{101}$$

Before fallback, only $O(B_\delta)$ observations have been collected, so all pairwise log-likelihood ratios are bounded below by $-C_0B_\delta$. The fallback bound gives conditional expected additional cost $O(B_\delta)$, uniformly over the fallback history. Thus

$$\mathbb{E}_0[\tau_{\beta_\delta}] \le \mathbb{E}_0[T^{\mathrm{seq}}_{\beta_\delta}] + CB_\delta\mathbb{P}_0(T^{\mathrm{seq}}_{\beta_\delta} > \Gamma B_\delta) = \frac{r}{d}B_\delta + O(1). \tag{102}$$

The cap and fallback guarantee almost-sure termination under every hypothesis, while the global rule guarantees δ-correctness. Together with equation 87, this proves equation 47. Combining it with equation 48 proves equation 49.

E Numerical Evidence for Both Separations

The simulations mirror the two constructions rather than introducing a new benchmark. All runs use the Bernoulli model of Section 6.1 with

$$p = 0.3, \qquad q = 0.6, \qquad d = 0.1837868974, \qquad \sigma = 0.5740881131.$$

For Flex/Lock we set $r = 4$, which gives $\kappa_4 = 30.0013295$. The routing pair has $r = 2$ and $\kappa_2 = 8.2217032$.

E.1 Policies and protocol

The Flex curve uses the optimal coordinatewise-freeze policy from Theorem 8: sample a_j until its likelihood coordinate first reaches b, then never sample that coordinate again. Its stopping time is exactly $\sum_j \tau_b^{(j)}$. Lock follows its forced rotor.

For Split, at each return to z_0 we choose a currently weakest coordinate and traverse the corresponding two-step cycle, either (a_1, a_1) or (a_2, a_2). For Coupled, the weaker coordinate is sampled first, but the second step must sample the other coordinate. Hence every completed Coupled round contains one observation from each stream, regardless of the root decision.

Certification is checked after every physical observation. We use

$$b \in \{20, 40, 80, 160, 320, 640, 1280\},$$

20,000 independent trajectories per simulated policy, and random seed 20260809. Pointwise intervals use 1.96 Monte Carlo standard errors. The same increment streams are reused across thresholds within a trajectory.

E.2 Finite-threshold results

For Flex and Lock define $\Delta_b = \widehat{\mathbb{E}}_0[T_b] - rb/d$; for Split and Coupled replace rb/d by $2b/d$. Table 1 reports the estimated excess.

A descriptive weighted fit

$$\widehat{\mathbb{E}}_0[T_b] - \frac{b}{C^\star} = \widehat{\alpha}\sqrt{b} + \widehat{\beta}$$

over the four largest thresholds gives the coefficients in Table 2. The fitted diffusion coefficient is essentially zero for both restoring systems and agrees with the explicit constant for both deficient systems.

Table 1: Estimated excess stopping time; Monte Carlo standard errors are in parentheses. Flex denotes the optimal freeze policy, and Split uses the weakest-coordinate restoring policy.

b	Flex	Lock	Split	Coupled
20	5.80 (0.47)	149.29 (0.89)	3.24 (0.33)	38.58 (0.46)
40	6.08 (0.66)	205.97 (1.15)	3.57 (0.46)	54.32 (0.62)
80	5.45 (0.92)	284.80 (1.53)	3.45 (0.65)	76.93 (0.84)
160	5.32 (1.31)	398.35 (2.07)	4.25 (0.92)	107.28 (1.16)
320	5.89 (1.84)	555.59 (2.82)	3.98 (1.30)	148.79 (1.60)
640	6.64 (2.60)	772.54 (3.85)	4.61 (1.83)	209.80 (2.22)
1280	8.05 (3.70)	1096.78 (5.38)	5.13 (2.61)	298.46 (3.11)

Table 2: Descriptive weighted fits of the excess to $\alpha\sqrt{b} + \beta$. Thresholds share trajectories, so the fit is a scaling diagnostic rather than an independent inferential test.

System or policy	fitted $\widehat{\alpha}$	theory α
Flex (freeze)	0.11	0.00
Flex (block)	-0.03	0.00
Flex (round robin)	30.01	30.00
Lock	30.01	30.00
Split	0.03	0.00
Coupled	8.21	8.22

Policy ablation. The access graph alone does not produce the bounded correction. In Flex, the coordinatewise-freeze and weakest-coordinate block policies both have bounded excess, whereas nonadaptive round robin has exactly the Lock law under H_0. Figure 3 isolates this distinction.

Distributional check. At $b = 1280$, the Kolmogorov–Smirnov distances between the empirical centered hitting times and their Gaussian-min limits are 0.0168 for Lock and 0.0108 for Coupled.

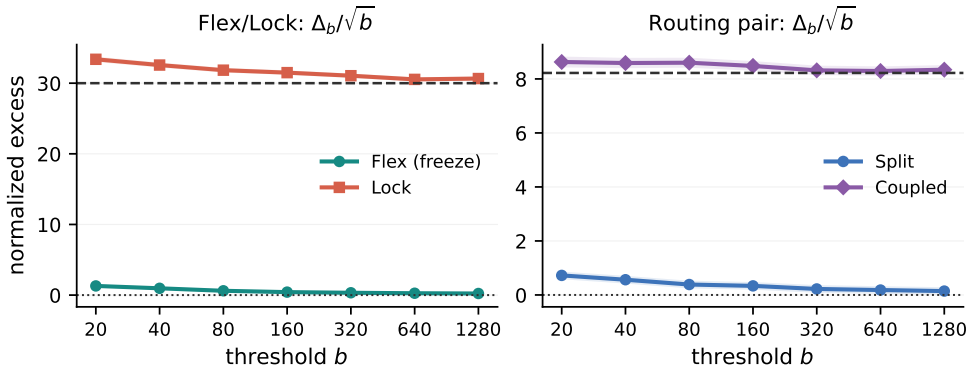

Figure 2: Excess normalized by $\sqrt{b}$. The Flex and Split curves approach zero; Lock and Coupled approach κ_4 and κ_2, respectively. Shaded bands are pointwise 95% Monte Carlo intervals.

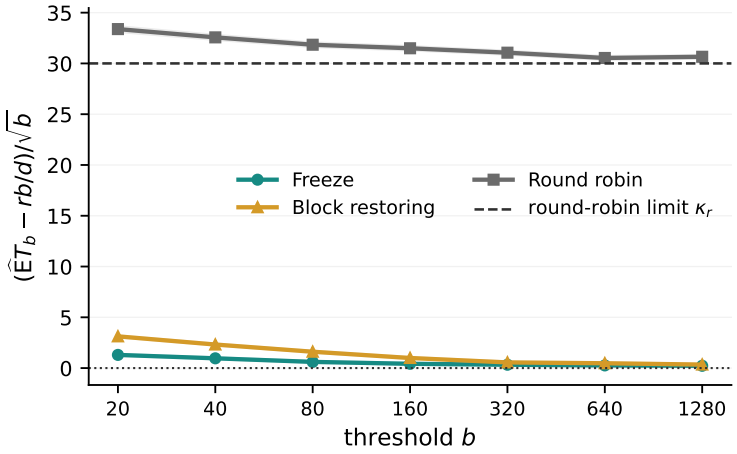

Figure 3: Flex policy ablation. Restoring policies have normalized excess tending to zero; round robin retains the diffusion correction κ_4. This separates usable circulation freedom from mere availability of the Flex graph.

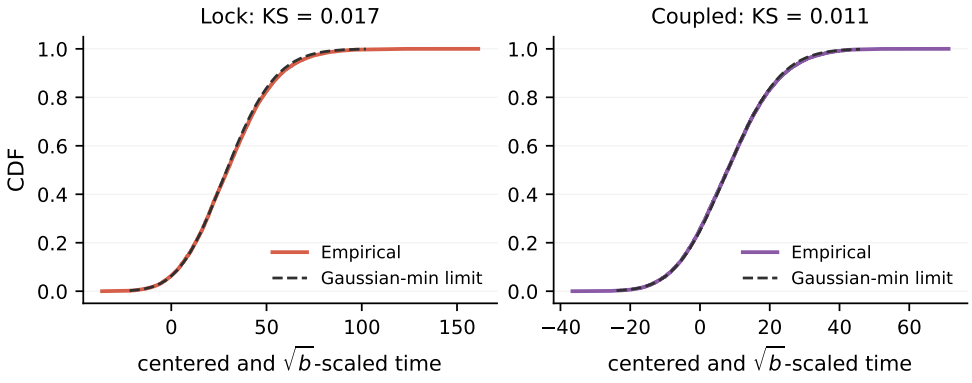

Figure 4: Empirical centered and $\sqrt{b}$-scaled stopping-time distributions at $b = 1280$, compared with the Gaussian-min limits in Theorem 8 and Proposition 9.

F Detailed Related Work and Positioning

Sequential design and active testing. The classical starting points are the optimality of the sequential probability ratio test (Wald & Wolfowitz, 1948) and Chernoff's adaptive choice of experiments (Chernoff, 1959). Modern active-testing formulations characterize sequentiality and adaptivity gains, controlled sensing, and problem-specific action selection (Naghshvar & Javidi, 2013a;b; Nitinawarat et al., 2013; Cohen & Zhao, 2015). Composite alternatives and nonparametric two-sample questions enlarge the statistical model (Deshmukh et al., 2021; Li et al., 2024), while error-exponent regions can couple several action policies (Hsu & Wang, 2025). These works establish the importance of adaptive sampling, but their leading-order objects are still experiment proportions or action policies evaluated through aggregate information rates. Our question begins after that rate has been matched: which evidence imbalances can be repaired while remaining on the first-order optimal face?

Fixed-confidence identification and allocation tracking. Information-theoretic lower bounds for best-arm identification lead to instance-dependent allocation problems (Kaufmann et al., 2016); Track-and-Stop turns the optimal allocation into an asymptotically optimal fixed-confidence procedure (Garivier & Kaufmann, 2016). Sequential multi-hypothesis testing in bandit models similarly combines likelihood-based stopping with adaptive sampling and switching costs (Prabhu et al., 2022). This literature explains why a policy should track optimal proportions. Dynamic access adds a distinct question: even if the target proportions are feasible and already known, can the path realize the local deviations needed to counter stochastic imbalance? The critical circulation image is the tangent-space analogue of a target allocation.

Controlled Markov structure and constrained sensing. Occupation measures and stationary control laws are central in adaptive control of Markov chains (Graves & Lai, 1997). Controlled Markov observations introduce state-dependent data and nonuniform control costs (Nitinawarat & Veeravalli, 2015); restless-bandit identification makes state–action flow conservation explicit in the first-order lower bound (Karthik et al., 2024). Other constraints price switching (Vaidhiyan & Sundaresan, 2015), restrict cumulative sensor use (Li et al., 2019), assign heterogeneous action costs (Vershinin et al., 2026), or expose time-varying source sets (Hsu & Wang, 2025). Controlled sensing also appears in quickest change detection (Veeravalli et al., 2024). These models motivate our stationary-flow benchmark. The distinction is that our constraint is endogenous reachability: choosing an edge changes the next feasible action set. At second order, the relevant degrees of freedom are therefore not all stationary flows, but signed circulations supported on the universally zero-slack face.

Second-order sequential asymptotics. Renewal and boundary-crossing methods quantify overshoot and corrected Brownian approximations in classical sequential analysis (Woodroofe, 1982; Siegmund, 1985). For controlled sequential design, second-order efficiency and controlled-random-walk effects were studied in earlier work (Keener, 1984; Lalley & Lorden, 1986); related lower bounds and constructions appear in Malyutov & Tsitovich (2001). For multihypothesis tests, the MSPRT program establishes first-order optimality and accurate expected sample-size expansions (Dragalin et al., 1999; 2000). More recent refinements address binary sequential testing and multiple independent streams (Li & Tan, 2020; Liu & Song, 2026). Those results identify overshoot, stopping-boundary, or multidimensional fluctuation corrections for a fixed action geometry. Our contribution is logically prior: under dynamic access, the critical circulation geometry decides whether the first unbounded correction can be controlled at all, and hence whether its scale is logarithmic or diffusive.

G Scope, Limitations, and Interpretation

What the dichotomy classifies. Fix a true hypothesis and its first-order optimal dual face. Under a full-support optimal multiplier and a strongly connected critical graph, the rank test classifies the first growing correction to the certification benchmark: full relative rank gives an $O(\log b)$ upper bound, while a noisy missing contrast forces $\Theta(\sqrt{b})$ excess. The Flex/Lock and Split/Coupled pairs sharpen these orders to exact $O(1)$ or explicit $\sqrt{b}$ expansions. The theorem is therefore a structural classifier, not a claim that every full-rank model has the same $O(1)$ constant or that every rank-deficient model shares one diffusion coefficient.

Table 3: The critical circulation image connects, but is not reducible to, the standard objects in the adjacent literatures.

Line of work	Primary first-order object	Distinction addressed here
Active sequential design and pure exploration	Experiment frequencies or target allocations	Execution order and local corrective motion are absent from the summary.
Controlled Markov and constrained sensing	State–action occupation measures, budgets, or accumulated costs	Only zero-slack signed circulations are free at second order.
Classical second-order sequential analysis	Overshoot or boundary-crossing correction under a fixed action geometry	The access geometry selects the correction regime before the constant is computed.

Why the assumptions are structural. Full support ensures that every alternative coordinate participates in the relative-evidence geometry. If the optimal multiplier lies on the boundary, the correct object is stratified by the active alternative set; a global rank computed in all $r - 1$ directions can then be misleading. Strong connectivity ensures that stationary perturbations are represented by circulations on a single recurrent component. With several critical components, the initial state, component selection, and the cost of moving between components enter in addition to within-component rank. Finally, positive variance in a missing contrast is what converts lack of mean control into a diffusion-scale lower bound. A variance-degenerate deficit is not evidence against the geometry; it signals that the next nonzero stochastic order, rather than $\sqrt{b}$, must be analyzed.

Three distinctions that should not be collapsed. First, the full access graph is not the critical graph: a noncritical edge may repair imbalance, but the exact accounting identity charges its positive slack. Second, critical cycle rank is not circulation-image rank: cycles can be statistically redundant after projection through the likelihood drift map, as the routing pair demonstrates. Third, available feedback is not used feedback: Flex admits restoring cycles, yet round robin reproduces the Lock diffusion term. These are respectively an optimization distinction, a geometry distinction, and an algorithmic distinction. Conflating any pair erases the mechanism responsible for the separation.

What is expected to remain stable. The exact excess identity and the definition of the critical circulation image do not depend on the Bernoulli examples. They extend to other finite access graphs and observation families whenever the stopping and optional-stopping conditions hold. What may change under unbounded increments, nonstationary observation laws, multiple recurrent components, or boundary dual solutions is the concentration argument converting the geometry into a particular rate. This separates a stable diagnostic layer—dual slack plus projected circulations—from model-specific probability estimates. Extending the rate classification to those regimes is the natural next step, rather than a different explanation of the examples proved here.