Cube Rounding for Matrix Sketching: Complete-Path Risk and Efficient Gram-Aware Updates

Anonymous authors
Paper under double-blind review

Abstract

Cube sampling preserves prescribed inclusion probabilities, yet minimizing the variance of individual rounding steps need not minimize the error of the final matrix sketch. We study this gap through the complete flight-and-landing process and develop an efficient implementation of a Gram-aware sampling rule. For a specified class of legal actions and landing procedures, we show how residual balance infeasibility forces terminal mutual exclusion, yielding policy-invariant covariance structure and conditions for exact optimality and quantitative near-optimality. Conversely, an explicit six-unit construction exhibits a unique locally optimal update with strictly positive complete-path regret, revealing the role of subsequent branching and landing. For rank-one positive semidefinite contributions, the resulting risk analysis connects Gram approximation to PCA reconstruction and subspace error. We then introduce Incremental Gram-Cube (IGC), which combines cross-step Gram reuse, an endpoint nullspace basis, and direct two-dimensional direction evaluation to accelerate a directional-energy surrogate on four-unit buffers. Under matched tie conventions, the implementation preserves the sampling rule in exact arithmetic. On 24 development cases spanning synthetic populations, real-data subsets, and larger synthetic inputs, IGC reduces end-to-end sampling-and-estimation time by 7.30–12.20% across the three groups relative to an already cached implementation of the same rule. All 9,216 paired evaluation draws produce identical selected indices and Frobenius errors. These results characterize when local Cube decisions suffice and demonstrate how matrix-aware updates can be computed with less overhead.

1 Introduction

Balanced randomized rounding converts fractional inclusion probabilities into a sample while preserving specified linear structure. Cube flight moves in the balance nullspace; landing resolves the remaining fractional coordinates (Deville & Tillé, 2004). For matrix sketching, the same local move changes immediate variance, future active sets, available continuations and landing error. Local legality or low immediate variance alone therefore does not determine the quality of the final estimator.

We ask when immediate-variance minimization is sufficient for the complete path, when it fails, and how a specified matrix-aware update can be evaluated with less work. The structural bridge is $J = \langle G, \Sigma \rangle$: policy choices determine the terminal covariance Σ, while computation can exploit the fixed contribution Gram matrix G. These are complementary questions about the risk and cost of an unbiased matrix sketch.

Our first contribution links residual balance infeasibility to terminal mutual exclusion and policy-invariant covariance, giving exact and quantitative near-optimality guarantees. A separate six-unit construction has a unique myopic root with positive complete-path regret because its branches and landing costs differ. Further strict-open and normal-cone certificates describe general matrix geometry; their scope is stated explicitly. The second contribution specializes complete risk to rank-one PSD Gram sketches and PCA, and experimentally distinguishes improvements in exact Gram risk from sampled downstream effects.

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.

The third contribution is Incremental Gram-Cube (IGC), a four-unit implementation of a specified Gram-aware direction rule. Reusing surviving Gram entries, choosing an endpoint nullspace basis and directly evaluating a two-dimensional direction remove repeated work. The added algebraic reductions save 7.30–12.20% of complete-draw time relative to an already cached implementation in three development groups. This energy-surrogate rule differs from the all-triple myopic policy in the exact constructions; its acceleration does not confer their action-class-specific optimality certificates. Likewise, the older lookahead risk gains and new timing gains were not measured jointly for one controller.

2 Related Work

Complete Cube analyses already retain flight and landing martingale increments (Breidt & Chauvet, 2011); controlled-martingale direction heuristics appear in concurrent work (Bousabaa & Sirolli, 2026). Ordered pivotal sampling has fixed-order zero joint-inclusion characterizations (Chauvet, 2012); our residual-derived protected block is shared by every legal policy under a declared landing interface. Fixed-marginal support LPs (Hasler et al., 2026) are outer design benchmarks, not Cube-path reachability proofs. Inequality-balanced sampling studies category bounds, matrix rounding and spreading (Tripet & Tillé, 2026); we retain fixed constraints and analyze complete-risk policy structure.

Dependent rounding and balanced design address other marginal or discrepancy guarantees (Gandhi et al., 2006; Harshaw et al., 2024). Greedy matching failures (Reingold & Tarjan, 1981; Frieze et al., 1990) motivate only our older half-integral context; the headline example uses stochastic triple circuits, branch-sensitive continuation and landing. Planning certificates (McMahan et al., 2005; Ross et al., 2008; Brown et al., 2010) and randomized spectral approximation (Drineas & Mahoney, 2016) are established tools, as is the population-gap PCA perturbation theorem (Yu et al., 2015). The contribution is Cube-specific terminal-correlation structure and its exact risk link to a rank-one PSD Gram-sketch application, rather than these tools.

Fast implementations and scope of the additional contribution. Chauvet & Tillé (2006) already reduced Cube execution from quadratic to linear dependence on population size for fixed balance dimension. Our four-unit implementation adds a matrix-aware direction computation within a small buffer; it is not the first fast Cube method. Gram reuse and a closed two-dimensional eigensolver are computational devices, not evidence of a new terminal-risk guarantee. The measured reference already includes Gram reuse, and the public Cube comparison is reported even though it remains faster.

3 Registered Complete-Path Cube Control

Let $p_t \in [0,1]^N$ be the inclusion-probability vector and let $C^\top p_t$ be the registered balance. A legal action chooses a direction u_t supported on fractional coordinates with $C^\top u_t = 0$. The two maximal unit-box steps $\alpha_t^+, \alpha_t^- > 0$ are selected with probabilities

$$\rho_t^+ = \frac{\alpha_t^-}{\alpha_t^+ + \alpha_t^-}, \qquad \rho_t^- = \frac{\alpha_t^+}{\alpha_t^+ + \alpha_t^-},$$

so $\mathbb{E}[p_{t+1} \mid \mathcal{F}_t] = p_t$. The registered state contains every variable used by the finite action generator, stopping rule, and mean-preserving landing kernel. A policy in Π_L is nonanticipating and uses only this registered interface.

The success theorems use a registered exact-mass pairwise landing kernel. It leaves integral coordinates unchanged; when the only fractional coordinates are r, s with $p_r + p_s = 1$, it outputs $(I_r, I_s) = (1, 0)$ with probability p_r and $(0, 1)$ with probability p_s. This kernel is mean preserving and preserves the integer sample size pathwise. Mean preservation alone is not sufficient for the mutual- exclusion theorem.

For matrix contributions Z_i, set $M_t = \sum_i p_{t,i} Z_i$, $\widehat{M} = \sum_i I_i Z_i$, and $J(\pi) = \mathbb{E}_\pi \left\| \widehat{M} - M_0 \right\|_F^2$.

Proposition 1 (Complete flight-plus-landing objective). For every registered policy, the stopped flight increments and the final landing increment are orthogonal martingale differences. Hence

$$J(\pi) = \sum_t \mathbb{E}\left\|M_{(t+1)\wedge\tau} - M_{t\wedge\tau}\right\|_F^2 + \mathbb{E}\left\|\widehat{M} - M_\tau\right\|_F^2,$$

and $\mathbb{E}[\widehat{M}] = M_0$.

This known decomposition identifies the exact finite-horizon Bellman cost. For the registered action set $\mathcal{A}_L(s)$,

$$V^*_{\mathcal{A}_L}(s) = \min_{a\in\mathcal{A}_L(s)}\{c(s,a) + \mathbb{E}[V^*_{\mathcal{A}_L}(S') \mid s,a]\},$$

with landing MSE at terminal states. The myopic policy μ minimizes only $c(s,a)$.

Let τ be the bounded flight stopping time and let $V^*(S_\tau)$ equal the expected value of the fixed registered landing kernel. For $Q^*(s,a) = c(s,a) + \mathbb{E}[V^*(S') \mid s,a]$, define the optimal advantage $A^*(s,a) = Q^*(s,a) - V^*(s) \geq 0$. Finite-horizon telescoping gives the standard performance-difference identity

$$J(\mu; s_0) - V^*(s_0) = \mathbb{E}_\mu\left[\sum_{t<\tau} A^*(S_t, \mu(S_t))\right]. \tag{1}$$

Consequently, myopic is complete-path optimal exactly when its action is Q^*-optimal almost surely under its state–time occupancy measure. On the finite registered tree this means every state–time pair reached with positive probability. Bellman optimality at the root alone is only necessary. The identity provides an exact check of complete-path optimality; it is not itself a novelty claim.

For the scalar-auxiliary two-row implementation, pairwise-distinct active scores imply that every support-minimal legal circuit has size three. Thus enumerating every active triple is the complete legal circuit set. A buffer of size L is complete only when it contains all fractional coordinates; otherwise every claim is relative to $\mathcal{A}_L$. The strict six-coordinate point uses every triple and hence has zero restriction gap to the standard complete circuit class; the finite circuit class is not equated with the infinite set of all nonminimal nullspace directions.

4 The Success Side

Assume integral mass $m = \sum_i p_i$, scalar auxiliary balance $b = \sum_i p_i x_i$, and the registered exact-mass pairwise landing kernel above, invoked when at most two coordinates remain fractional. Define

$$\mathcal{P}^{m-1}_{-i} = \left\{\sum_{j\neq i} q_j x_j : 0 \leq q_j \leq 1,\ \sum_{j\neq i} q_j = m-1\right\}, \quad F = \{i : b - x_i \notin \mathcal{P}^{m-1}_{-i}\}.$$

Theorem 2 (Residual-feasibility exact success). Every registered terminal selection satisfies $\sum_{i\in F} I_i \leq 1$ almost surely. If $Z_i = 0$ for $i \notin F$, then every registered policy has the same complete MSE,

$$J(\pi) = \sum_{i\in F} p_i \left\|Z_i\right\|_F^2 - \left\|\sum_{i\in F} p_i Z_i\right\|_F^2.$$

Consequently myopic is complete-path optimal on this stratum.

The proof is geometric: if a protected coordinate reached one during flight, the remaining coordinates would provide the forbidden residual completion. Pairwise landing can then select at most one protected coordinate, fixing its entire covariance block.

Exact sample mass makes risk invariant under $Z_i \mapsto Z_i - A$. Thus common residual contributions $Z_i = A$ also give exact policy equivalence; a translation-centered certificate can

strictly tighten the bound below (Appendix K). This uses standard fixed-mass contribution-difference symmetry (Breidt & Chauvet, 2011), not a centered-feature strict-success claim. For nonzero deviations define

$$J_F = \sum_{i\in F} p_i(1-p_i)\,\|Z_i\|_F^2 - 2\sum_{\substack{i<j\\ i,j\in F}} p_i p_j\,\langle Z_i, Z_j\rangle\,,$$

$$\overline{R} = \sum_{i\in R} p_i(1-p_i)\,\|Z_i\|_F^2 + 2\sum_{\substack{i<j\\ i,j\in R}} \beta_{ij}|\,\langle Z_i, Z_j\rangle\,|,$$

where β_{ij} is the largest absolute covariance compatible with the two Bernoulli marginals. For $w_{ij} = \min(p_i, p_j) - \max(0, p_i + p_j - 1)$, set

$$B_{\mathrm{Fr}} = 2\sum_{i<j:\{i,j\}\not\subseteq F} w_{ij}|\,\langle Z_i, Z_j\rangle\,|,$$

and, with $x_+ = \max(x, 0)$, $L_E = (\sqrt{J_F} - \sqrt{\overline{R}})_+^2$ and $U_E = (\sqrt{J_F} + \sqrt{\overline{R}})^2$.

Theorem 3 (Full-dimensional near success). Let Π_L be nonempty, and let $J^*_{\Pi_L} = \inf_{\pi\in\Pi_L} J(\pi)$. For every $\pi \in \Pi_L$, $L_E \le J(\pi) \le U_E$. Therefore

$$J(\mu) - J^*_{\Pi_L} \le B := \min\{B_{\mathrm{Fr}}, U_E - L_E\}.$$

If $L_E > 0$, then

$$\frac{J(\mu)}{J^*_{\Pi_L}} \le 1 + \frac{B}{L_E}.$$

If some $i \in F$ has $0 < p_i < 1$, then for every $\varepsilon > 0$ the conditions $L_E > 0$ and $B/L_E < \varepsilon$ define a nonempty open set in the full contribution space. If additionally $R \neq \varnothing$, this open set contains points where every residual contribution is nonzero and is not confined to the positive-codimension residual-zero stratum.

Corollary 4 (Action-class robustness of the energy guarantee). If $\Pi_L \subseteq \Pi'$ and every policy in Π' returns a binary terminal vector and preserves the same terminal marginals and mutual exclusion on F, then the preceding interval, additive bound, and ratio remain valid with $J^*_{\Pi'}$ in place of $J^*_{\Pi_L}$. Thus the energy theorem extends beyond triple circuits whenever its terminal invariants survive; the strict-open and normal-cone certificates remain class-relative and must be recomputed after expansion.

Further exact-success constructions and a covariance-based three-way certificate are given in Appendix A. The strict-open construction applies to general matrix contributions, not to a realizable rank-one PSD feature geometry.

5 The Failure Side

Theorem 5 (Branch- and landing-sensitive separation). There is a six-unit non-half-integral two-constraint Cube instance with every active triple registered such that (i) the unique immediate-variance minimizer has forced-root regret $369406/19173 > 0$; (ii) its two pivotal branches require different optimal continuations; and (iii) removing landing value changes the optimal root. Exact dynamic programming gives

$$J_{\mathrm{myopic}} = \frac{1113001}{4410}, \qquad J^* = \frac{298063}{1743}, \qquad J_{\mathrm{myopic}} - J^* = \frac{29785853}{366030} > 0.$$

The base state is

$$p = (1/2, 1/2, 1/4, 1/4, 1/4, 1/4), \quad x = (3, 1, -3, 2, -2, 0), \quad Z = (2, -4, 12, -16, 20, -24).$$

The unique myopic root is triple $(2, 4, 5)$, whereas the complete-path root is $(1, 2, 6)$. This isolates the intended mechanism without relying on action ties. Changing only Z_6 to $-24+t$ preserves that unique myopic root for $|t| \le 1/40$ and gives forced-root regret greater than $8959/800$, by a uniform perturbation argument (Appendix K). This is a conservative analytic failure interval, not a maximal phase boundary.

6 Consequences for Gram Sketching and PCA

The complete-path objective also covers a standard machine-learning statistic without changing the Cube process. Let $X \in \mathbb{R}^{n\times d}$ be a fixed training-centered matrix with rows $x_i^\top$, and let $p_i > 0$ be the registered inclusion probabilities. Define

$$\Gamma = \frac{1}{n}X^\top X, \qquad Z_i = \frac{x_i x_i^\top}{np_i}, \qquad \widehat{\Gamma} = \sum_{i=1}^{n} I_i Z_i. \tag{2}$$

Corollary 6 (Gram sketch specialization). For every registered mean-preserving flight-plus-landing policy,

$$\mathbb{E}[\widehat{\Gamma} \mid X] = \Gamma \quad \text{and} \quad J_\Gamma(\pi; X) = \mathbb{E}_\pi[\|\widehat{\Gamma} - \Gamma\|_F^2 \mid X].$$

The contribution Gram matrix used by the success–failure theory is

$$G_{ij} = \langle Z_i, Z_j\rangle_F = \frac{(x_i^\top x_j)^2}{n^2 p_i p_j}.$$

Hence every contribution is positive semidefinite and has rank at most one.

The corollary substitutes rank-one Gram contributions into the complete-path objective; it does not introduce another sampling algorithm. The restriction is still substantive because real feature rows impose a nonnegative squared-inner-product geometry. We ask whether the same success and failure mechanisms survive within that geometry.

Figure 1 in the appendix summarizes the link.

Blockwise lift. For the real-data experiment, a deterministic 192-row window is divided into 32 disjoint six-row blocks. Its preprocessing, partition, probability slots, and policies are fixed before Cube randomization. Writing $E_b = \widehat{\Gamma}_b - \Gamma_b$, each block is conditionally unbiased. Independent block randomization therefore gives

$$\mathbb{E}\left\|\sum_b E_b\right\|_F^2 = \sum_b \mathbb{E}\|E_b\|_F^2, \tag{3}$$

because all cross-block inner products have zero conditional expectation. The target in (3) is the registered window Gram, not the full-training Gram.

Let P_k denote a rank-k projector maximizing $\langle P, \Gamma\rangle_F$, and let $\widehat{P}_k$ maximize the same objective with $\widehat{\Gamma}$. The average reconstruction loss is $\mathcal{L}_\Gamma(P) = \mathrm{tr}((I-P)\Gamma)$.

Corollary 7 (PCA downstream consequence). For every registered policy,

$$\mathbb{E}[\mathcal{L}_\Gamma(\widehat{P}_k) - \mathcal{L}_\Gamma(P_k) \mid X] \le \sqrt{2k J_\Gamma(\pi; X)}. \tag{4}$$

If the population eigengap $\delta_k = \lambda_k(\Gamma) - \lambda_{k+1}(\Gamma)$ is positive, then

$$\mathbb{E}\|\sin\Theta(\widehat{P}_k, P_k)\|_F^2 \le \frac{4J_\Gamma(\pi; X)}{\delta_k^2}. \tag{5}$$

The standard perturbation argument is given in Appendix M.

These inequalities explain why complete-path Gram risk matters for PCA: smaller risk tightens valid reconstruction and subspace guarantees. They do not imply that every randomized realization or downstream classifier must improve when expected Gram risk decreases. The experiments therefore report null and adverse downstream outcomes as well as improvements.

7 Efficient Evaluation of a Gram-Aware Cube Rule

The preceding theory separates the legality of a Cube update from the quality of its complete path. We now address a complementary computational question: can a specified Gram-aware direction rule be evaluated more cheaply without deliberately replacing its matrix

geometry by a proxy? We call the resulting implementation Incremental Gram-Cube (IGC). The name describes a specialized implementation, not a new class of balanced-sampling guarantees. Small active sets and fast Cube flight are established ideas (Chauvet & Tillé, 2006).

Estimator and local rule. Write $\pi_i = p_{0,i}$ for the fixed initial inclusion probabilities, $z_i \in \mathbb{R}^d$ for feature rows, and $Z_i = z_i z_i^\top / \pi_i$. The target is $S = \sum_i z_i z_i^\top$, with unbiased estimator $\widehat{S} = \sum_i I_i Z_i$. Dividing both by N gives the normalized Gram convention used in the preceding PCA section. In particular, the normalization does not change a risk ratio. The contribution Gram entries are

$$G_{ij} = \langle Z_i, Z_j \rangle_F = (y_i^\top y_j)^2, \qquad y_i = z_i / \sqrt{\pi_i}. \tag{6}$$

These entries depend on the initial probabilities, not on the evolving fractional state p_t.

IGC processes a FIFO buffer of four fractional units, initially ordered by $t_i = \|z_i\|_2^2 / \pi_i$ with an index tie rule. The two flight constraints are sample mass and a fixed scalar auxiliary variable x. On the current buffer, it selects a nonzero u satisfying $\mathbf{1}^\top u = x^\top u = 0$ to minimize the directional energy

$$R(u) = \frac{u^\top G_W u}{u^\top u}. \tag{7}$$

This is a surrogate direction rule. The actual conditional variance of the maximal two-endpoint update is

$$c(p_t, u) = \alpha^+(p_t, u)\alpha^-(p_t, u)\, u^\top G_W u. \tag{8}$$

The step-length factor depends on the direction; minimizing Eq. (7) therefore does not generally minimize Eq. (8). IGC is not identified with the all-triple myopic policy μ in our exact constructions. It uses a different, specified action rule; those action-class-specific optimality certificates are not transferred to it.

Reuse of surviving Gram entries. We retain the previous four unit identifiers and their 4×4 Gram matrix. After retirement and refill, each entry whose two identifiers survive is copied, and every other off-diagonal entry is evaluated using Eq. (6). Diagonals are t_i^2. If r units survive between two full buffers, only $6 - \binom{r}{2}$ new feature inner products are needed. For the common transition with three survivors this is three rather than six. This is an operation count, not a claim of a twofold runtime speedup. Both the cached reference and our accelerated implementation use this reuse.

An endpoint basis and direct direction evaluation. Sort the four buffer identifiers by their strictly increasing auxiliary values, and apply the reference's local affine normalization to obtain $s_0 < s_1 < s_2 < s_3$. With $h = s_0 - s_3$, use

$$B = \begin{pmatrix} (s_3 - s_1)/h & (s_3 - s_2)/h \\ 1 & 0 \\ 0 & 1 \\ (s_1 - s_0)/h & (s_2 - s_0)/h \end{pmatrix}, \qquad M = B^\top G_W B, \quad H = B^\top B. \tag{9}$$

Thus $u = Bv$ and the optimization is the 2×2 generalized Rayleigh problem $\min_{v \neq 0} v^\top M v / (v^\top H v)$. Endpoints already maximize auxiliary separation, so an explicit search over six pivot pairs and a duplicate basis check can be removed. The final direction constraint check is retained. Local normalization is retained as well.

Let m_{ab} and h_{ab} denote entries of M and H, and set

$$\Delta = h_{11}m_{00} - h_{00}m_{11}, \quad b = h_{11}m_{01} - h_{01}m_{11}, \quad c = h_{00}m_{01} - h_{01}m_{00}, \quad D = \sqrt{\Delta^2 + 4bc}.$$

A minimum-eigenvalue direction is

$$v = \begin{cases} (-2b, \Delta + D)^\top, & \Delta \geq 0, \\ (\Delta - D, 2c)^\top, & \Delta < 0. \end{cases} \tag{10}$$

If this vector vanishes, M is proportional to H in exact arithmetic and we use the reference default $v = (1, 0)^\top$. Positive matrix rescaling does not affect the direction. The implementation normalizes G_W and M, checks the positive determinant of H, and clips a negative

roundoff discriminant to zero. It then applies the reference direction orientation, maximal steps, branch probabilities and tail routines. These elementary algebraic reductions are not claimed as new spectral theory.

Proposition 8 (Exact-arithmetic computation equivalence). For strictly ordered auxiliary scores, Eq. (9) has full column rank and spans the mass–auxiliary nullspace. The cached entries are the exact contribution Gram entries. Equation (10) minimizes Eq. (7). Consequently, on paths for which the reference and IGC use the same minimizer tie convention, the same queue, maximal-step updates and landing kernel induce the same terminal law in exact arithmetic.

Appendix C gives the proof and complete algorithm. The claim is conditional on the declared rule and tie semantics; it is neither a universal floating-point identity nor a terminal-risk improvement theorem. The implementation uses $O(Nd + N \log N)$ preprocessing and flight work, $O(kd^2)$ to construct the returned dense estimator, and $O(Nd + N + d^2)$ storage including transformed features and output. The optimizations reduce constants within this complexity. They do not improve the population-size order of classical fast Cube.

8 Experiments

The experiments test the scope of the theorems and report real-data gains, null results, and failures under the same fixed protocols.

Policies and statistical protocol. Let V_μ be the complete-path value of the myopic policy, including the registered landing kernel. For $h \geq 1$, define the rollout scores

$$\widetilde{V}_0(s) = V_\mu(s), \qquad \widetilde{V}_h(s) = \min_{a \in \mathcal{A}_L(s)} \left\{ c(s,a) + \mathbb{E}[\widetilde{V}_{h-1}(S') \mid s, a] \right\}. \tag{11}$$

Policy Hh takes an action attaining the minimum in Eq. (11) and recomputes the same h-step problem after every realized branch. Thus H2 and H3 are receding-horizon two- and three-flight-step policies with a myopic complete-path tail; they are not depth labels for one fixed root tree. The exact policy uses the full Bellman value. In this six-row, all-triple setting, every maximal flight fixes at least one coordinate. After three flights at most three fractional coordinates remain, so the only legal triple has no policy choice; hence H3 is structurally an exact-depth reference. Exact-tail rollout noninferiority is the standard policy-improvement property, not a new empirical law. Equal values need not imply identical tied actions or sampled endpoints.

The outer statistical unit is one of 20 fixed label-independent partitions, each with 32 six-row blocks and 32 common-random-number trajectories per policy. Trajectories are averaged within partition before policy differences are paired to myopic. Intervals use 10,000 percentile cluster-bootstrap resamples. The planned non-oracle exact-risk family is the eight H2/H3 versus myopic contrasts across four datasets; exact Bellman is an oracle audit. Other endpoints are reported separately. No multiplicity correction was registered, so intervals are dataset-wise rather than simultaneous and no family-wise error claim is made. These intervals are conditional on the fixed data and preprocessing: they describe partition/window and sketch randomness, not uncertainty from resampling new training or test populations.

Gram/PCA application study. The study first varies residual scale, protected–residual feature correlation, and anisotropy in 72 exact rank-one PSD cases. A fixed blockwise pipeline then evaluates Digits (development), UCI HAR and Fashion-MNIST (validation), and UCI Letter Recognition (untouched external) (Anguita et al., 2013; Xiao et al., 2017; Slate, 1991). The primary endpoint is complete-path Gram MSE; sampled Gram, PCA and ridge are propagation endpoints. Original operational floating ties and exact-tie supplements are distinct (Appendix I).

The phase comprises 72 executions but 63 distinct row geometries; the 12 zero-residual cells represent only three geometries because their correlation parameter is inactive. This is a frozen deterministic replication of a development-rehearsed family, not an independent holdout or prevalence estimate. The phase has zero structural, marginal, MSE-identity,

Table 1: Frozen exact block-Gram risk. Entries are policy-minus-myopic differences with dataset-wise 95% paired partition-cluster intervals. Negative values favor lookahead.

Dataset	H2 ΔJ_Γ [95% CI]	H3 ΔJ_Γ [95% CI]	H3 relative
Digits	-0.771 $[-0.827, -0.714]$	-0.782 $[-0.839, -0.724]$	-0.23%
UCI HAR	-1.350 $[-1.428, -1.269]$	-1.369 $[-1.450, -1.286]$	-2.44%
Fashion	-1.941 $[-2.111, -1.797]$	-1.972 $[-2.134, -1.829]$	-2.59%
Letter	-0.123 $[-0.133, -0.113]$	-0.125 $[-0.135, -0.115]$	-3.35%

success-theorem, or expectation-level PCA-bound violations: 12 zero-residual cases are exact theorem successes and all other 60 cases are strict failures, with maximum relative reduction $(J_\mu - J^*)/J_\mu = 25.33\%$ (Appendix D).

Table 1 gives all eight planned non-oracle primary-risk contrasts. Both H2 and H3 reduce exact additive block Gram risk on Digits, UCI HAR, Fashion-MNIST, and the untouched Letter evaluation; every reported interval lies below zero. As required by the exact-depth structure, H3 matches the exact Bellman value on every registered six-row block. Relative H3 reductions range from 0.23% on Digits to 3.35% on Letter; the corresponding myopic mean risks are 343.741, 56.214, 76.237, and 3.732. Most realized Gram, PCA, projector, and ridge intervals cross zero; the lone ridge exception is a tiny development-only Digits H2 gain ($+0.00117\,[0.00002, 0.00237]$). H2 and H3 also use roughly 34–35× and 41–42× the offline complete-policy evaluation work proxy of myopic (flight-cost cache misses), not online trajectory latency. Selected endpoint plots appear in Appendix D.

Additional mechanism, action-class, and real-state audits appear in Appendices B–G.

8.1 Complete-Draw Runtime of Incremental Gram-Cube

This experiment is separate from the preceding blockwise planning study. Its statistical target is the same full Frobenius error, but its datasets, action rule, hardware and timing endpoint differ. We do not combine the planning risk reductions with the following runtime reductions as if they were achieved by one evaluated policy.

Comparators and measurement. The principal reference, Cached Gram-Cube (implementation 88), already uses exact cross-step Gram reuse. Endpoint-only (182) changes the basis construction; IGC (186) additionally applies Eq. (10). All share the initial probabilities, FIFO rule, full contribution geometry, tails and dense Horvitz–Thompson output. A public Cube C-core adapter from sondage 0.9.1 (5) is a different-policy speed reference. It adapts the public C core's allocation, random-number and error interfaces; it is not a full R-package invocation or a claim of matching R's seeded stream. Bilinear and plane-projection approximations receive the same computational optimizations (187 and 188), avoiding an asymmetric implementation comparison.

We use 24 previously exposed development cases: eight synthetic cases, eight real-data subsets, and eight larger synthetic cases. Each group contains paired equal and unequal inclusion designs. Real inputs are subsets of Iris, Wine, Breast Cancer and Digits; they are not the four blockwise datasets above. Runs use one MacBook CPU numerical process and single-thread settings. Compilation uses C++17, -O3, -fno-fast-math and -ffp-contract=off. Every timed draw includes input checks, preparation, sorting, sampling, tails, dense estimator construction, output consumption and destruction. After warmup, five timing batches per case rotate method order. A group ratio is the sum of candidate batch seconds divided by the sum of reference batch seconds, not an average of selected speedup ratios.

Observed acceleration and sample agreement. Table 2 reports the complete implementation gain in one run; percentages are not multiplied across different experiments. All 24 case-wise median time ratios for IGC versus the cached reference are below one. In the same run, all 9,216 paired quality draws have identical selected indices and full Frobenius errors. This is finite-input agreement, not proof that floating-point trajectories are identical for every input. Endpoint construction supplies most of the measured gain. The direct direction

Table 2: Same-run complete-draw timing on new seeds (V103). All inputs were previously exposed development data. Lower time ratios are better; these are descriptive measurements without a cross-machine claim.

Group	Endpoint/ref.	IGC/ref.	Time saved	Speedup
Synthetic	0.8938	0.8780	12.20%	1.139×
Real subsets	0.9035	0.8939	10.61%	1.119×
Larger synthetic	0.9321	0.9270	7.30%	1.079×

formula adds group-level reductions of 1.77%, 1.07% and 0.55% over endpoint-only in this run; an earlier seed run gave 2.50%, 2.21% and 0.31%. Some individual cases are slower, so this second component is a small, workload-dependent optimization. The earlier endpoint-only trials also found a roughly 9.7–10.4% reduction on the real-subset group.

External and approximate comparators. IGC takes 1.854, 1.730 and 1.611 times the public Cube reference's complete draw time on the three groups. Thus the within-rule acceleration is not speed dominance over public Cube. IGC also takes 1.071 times the bilinear and 1.117 times the plane approximation's time on real subsets. These approximations have conditional risk tradeoffs: for example, IGC has lower paired risk than the plane approximation on equal-probability Wine and Digits, but no universal ordering follows. All these inputs were used during iterative method development. The experiment establishes a specialized computational contribution and exposes the remaining speed–risk tradeoff; it does not establish state-of-the-art dominance.

9 Limitations

The theory is relative to a specified action and landing class. The energy bound extends only when the terminal marginal and mutual-exclusion invariants survive; strict and normal-cone certificates require recomputation. The attainable gap to a broader fixed-marginal terminal-design relaxation is unresolved. Our near-optimality region is not a region of exact optimality. The separate strict-open example has $G_{13} = -6$, which cannot be a squared-inner-product Gram entry, so it is not a feature-space success certificate. Constructed and finite real-state examples do not establish success or failure prevalence over arbitrary data.

Lower expected Gram error need not yield an observable improvement in each PCA or classification experiment. Most such intervals remain unresolved, and the H2/H3 work counts concern offline policy evaluation, not online latency. The timing study uses 24 exposed development cases on one host with five rounds. IGC beats its cached reference but remains slower than public Cube; its directional-energy objective omits the direction-dependent step-length factor. No combined IGC-plus-lookahead controller was evaluated. Further scope qualifications are in Appendix L.

10 Conclusion

Cube rounding has both statistical and computational structure. Residual feasibility, terminal covariance and branch-sensitive landing explain when local variance minimization succeeds and fails. Rank-one matrix contributions connect this theory to Gram sketching and PCA, without guaranteeing uniform downstream gains. For a separate, specified Gram-aware direction rule, incremental reuse and two small algebraic reductions lower complete-draw cost while preserving observed outputs on paired tests. The resulting paper combines a risk theory with an implementation contribution, while keeping policy optimality, numerical equivalence and runtime evidence as distinct claims.

AI Use Statement

Generative AI assistants were used to develop and refine hypotheses and conceptual frameworks, formulate mathematical claims, assist with proof construction and critical review, design experiments, implement methods and analysis code, construct semisynthetic test instances, interpret results, search and summarize literature, and draft, edit and translate the manuscript. The current research and review workflow includes AI-agent assistance; AI-agent review is not an independent human certification. Exact arithmetic, protocol checks and numerical audits provide the documented computational checks. This disclosure does not assert an unrecorded human audit. The submitting authors remain responsible for the final text, claims, references and artifacts, and must make the corresponding disclosure in the submission form.

Reproducibility Statement

Appendix E specifies preprocessing, probability and auxiliary slots, seeds, policy evaluation, PCA ranks, downstream fitting and protocol deviations. Appendix F and Appendix G report mechanism and action-class audits. The proof appendices state the exact legal action and landing interfaces and rational certificates. Supplementary execution records distinguish original frozen results from later supplemental audits; these roles are not interchangeable. Appendix C documents the additional local timing study, algorithm and case-level results. This integration uses the repaired theory manuscript and completed development experiments; assembly did not launch new research experiments. The supplementary archive supplies algorithm sources, existing numerical records and reproduction instructions; its README states the coverage and dependencies of each component.

References

Davide Anguita, Alessandro Ghio, Luca Oneto, Xavier Parra, and Jorge Luis Reyes-Ortiz. A public domain dataset for human activity recognition using smartphones. In European Symposium on Artificial Neural Networks, 2013.

Abdelkader Bousabaa and Remi Sirolli. The layla-sphere method: Balanced sampling as a controlled martingale on the complex sphere. SSRN Electronic Journal, 2026. doi: 10.2139/ssrn.6955098.

F. Jay Breidt and Guillaume Chauvet. Improved variance estimation for balanced samples drawn via the cube method. Journal of Statistical Planning and Inference, 141(1):479–487, 2011.

David B. Brown, James E. Smith, and Peng Sun. Information relaxations and duality in stochastic dynamic programs. Operations Research, 58(4-part-1):785–801, 2010.

Guillaume Chauvet. On a characterization of ordered pivotal sampling. Bernoulli, 18(4): 1320–1340, 2012. doi: 10.3150/11-BEJ380.

Guillaume Chauvet and Yves Tillé. A fast algorithm for balanced sampling. Computational Statistics, 21:53–62, 2006. doi: 10.1007/s00180-006-0250-2.

Jean-Claude Deville and Yves Tillé. Efficient balanced sampling: The cube method. Biometrika, 91(4):893–912, 2004.

Petros Drineas and Michael W. Mahoney. Randnla: Randomized numerical linear algebra. Communications of the ACM, 59(6):80–90, 2016.

Alan Frieze, Colin McDiarmid, and Bruce Reed. Greedy matching on the line. SIAM Journal on Computing, 19(4):666–672, 1990. doi: 10.1137/0219045.

Rajiv Gandhi, Samir Khuller, Srinivasan Parthasarathy, and Aravind Srinivasan. Dependent rounding and its applications to approximation algorithms. Journal of the ACM, 53(3): 324–360, 2006.

Christopher Harshaw, Fredrik Savje, Daniel A. Spielman, and Peng Zhang. Design and analysis of the gram–schmidt walk: A general approach to balancing covariates in randomized experiments. Journal of the American Statistical Association, 2024. Advance online publication.

Caren Hasler, Esther Eustache, and Yves Tillé. In search of the most balanced sampling design. arXiv preprint arXiv:2607.26544, 2026.

H. Brendan McMahan, Maxim Likhachev, and Geoffrey J. Gordon. Bounded real-time dynamic programming: RTDP with monotone upper bounds and performance guarantees. In Proceedings of the 22nd International Conference on Machine Learning, pp. 569–576, 2005.

Edward M. Reingold and Robert E. Tarjan. On a greedy heuristic for complete matching. SIAM Journal on Computing, 10(4):676–681, 1981. doi: 10.1137/0210050.

Stéphane Ross, Joelle Pineau, Sébastien Paquet, and Brahim Chaib-draa. Online planning algorithms for POMDPs. Journal of Artificial Intelligence Research, 32:663–704, 2008. doi: 10.1613/jair.2567.

David Slate. Letter recognition. UCI Machine Learning Repository, 1991.

Arnaud Tripet and Yves Tillé. Balanced sampling with inequalities: Application to category bounding, matrix rounding, and spread sampling. Journal of the American Statistical Association, 121(553):796–806, 2026. doi: 10.1080/01621459.2025.2550667. Published online 27 October 2025.

Han Xiao, Kashif Rasul, and Roland Vollgraf. Fashion-MNIST: A novel image dataset for benchmarking machine learning algorithms. arXiv preprint arXiv:1708.07747, 2017.

Yi Yu, Tengyao Wang, and Richard J. Samworth. A useful variant of the davis–kahan theorem for statisticians. Biometrika, 102(2):315–323, 2015. doi: 10.1093/biomet/asv008.

A Further Success and Covariance Certificates

Theorem 9 (Explicit strict open Gram success). For

$$p = (1/2, 1/2, 1/4, 1/4, 1/4, 1/4), \quad x = (3, 1, -3, 2, -2, 0)$$

with all six-coordinate legal triples, there is an explicit rational $G_0 = ZZ^\top \succ 0$, where

$$Z = \begin{pmatrix} 3 & 0 & 0 & 0 & 0 & 0 \\ 0 & 5 & 0 & 0 & 0 & 0 \\ -2 & 1 & 3 & 0 & 0 & 0 \\ -2 & 2 & -1 & 3 & 0 & 0 \\ 2 & -2 & -2 & 0 & 2 & 0 \\ 0 & 0 & 0 & 2 & 0 & 3 \end{pmatrix},$$

for which myopic and complete-path optimal actions agree at every reachable state and $J(\mu) = J^* = 585029/34100$. The minimum immediate and Bellman gaps are $1/16$ and $5/48$. Every symmetric Gram perturbation E with

$$\|E\|_F < \frac{26244}{309512375}$$

preserves positive definiteness, the induced myopic tree, and the uniquely optimal terminal covariance in the registered class.

A.1 An exact criterion and certified three-way partition

For a fixed state, let $K(s) = \mathrm{conv}\{\Sigma_\pi(s) : \pi \in \Pi_L\}$ be the attainable terminal covariance set, where $\Sigma_\pi(s) = \mathrm{Cov}_\pi(I \mid S_0 = s)$. With contribution Gram matrix G, $J_G(\pi; s) = \langle G, \Sigma_\pi(s)\rangle$.

Theorem 10 (Normal-cone characterization). The myopic terminal covariance Σ_μ is complete-path optimal in the registered class if and only if

$$G \in N^{\min}_{K(s)}(\Sigma_\mu) := \{H : \langle H, \Sigma - \Sigma_\mu\rangle \geq 0,\ \forall \Sigma \in K(s)\}.$$

Given any valid outer relaxation $\overline{K}(s) \supseteq K(s)$ and any set $\mathcal{W}(s) \subseteq \{\Sigma_\pi(s) : \pi \in \Pi_L\}$ of covariances generated by feasible witness policies, suppose the outer optimization is computable and has finite value $L_{\mathrm{out}} = \inf_{\Sigma \in \overline{K}(s)} \langle G, \Sigma\rangle$, and the witness set is finite or has a computable linear optimization. The following partition is sound:

- certified-success if the outer lower bound equals the feasible myopic value;
- certified-failure if a feasible witness strictly beats myopic;
- ambiguous otherwise.

Writing $J^*_G(s) = \inf_{\pi \in \Pi_L} J_G(\pi; s)$ and $U_{\mathrm{wit}} = \inf_{\Sigma \in \mathcal{W}(s) \cup \{\Sigma_\mu\}} \langle G, \Sigma\rangle$ gives the computable bracket

$$J_G(\mu; s) - U_{\mathrm{wit}} \leq J_G(\mu; s) - J^*_G(s) \leq J_G(\mu; s) - L_{\mathrm{out}}.$$

If $L_{\mathrm{out}} > 0$, then $J_G(\mu; s)/J^*_G(s) \leq J_G(\mu; s)/L_{\mathrm{out}}$.

The normal cone is standard convex geometry. Here it is applied to the complete Cube path and paired with a sound three-way decision rule. This distinguishes sufficient certificates from a complete boundary.

We also obtain a quantitative failure threshold on a valid Gram path. From the strict base point, let $G(t) = G_0 + tzz^\top$ with $z = (20, -16, -40, -24, 24, 12)^\top$. For one fixed feasible witness w,

$$J_{G(t)}(w) - J_{G(t)}(\mu) = \frac{156159}{136400} - t\frac{949477}{8525}.$$

It crosses at $t_* = 156159/15191632$, while the first possible myopic action switch is only at $t_{\mathrm{act}} = 407/8128 > t_*$. Thus the unchanged myopic policy is certified to fail for $t \in (t_*, t_{\mathrm{act}})$. Along this rank-one path, $\|tzz^\top\|_F = t\|z\|_2^2$ with $\|z\|_2^2 = 3552$, so the conservative strict-success ball certifies only $t < t_{\mathrm{succ}} = 2187/91615663000$. The interval $[t_{\mathrm{succ}}, t_*]$ remains unresolved by these two certificates.

B Overview of Additional Mechanism Audits

Theory-targeted audits. The registered leakage family has 294 exact executions and 258 distinct geometries. All 42 zero-leakage cases are exact successes; every nonzero case has a strict witness gap, but all 42 smallest-leakage cases remain certified within 1.02. The ratio bound has no violations. Sixteen failures have a Bellman-optimal root and fail later, checking the full-path requirement in Eq. (1). These are semisynthetic mechanism audits, not natural prevalence estimates; full leakage, search and crossing results appear in Appendix F.

Effect of restricting the action class. A six-coordinate buffer exposes 20 of 35 root triples. Two of 12 exact systems have nonzero action-class gaps, at most 4.67% relative. A broader terminal-design relaxation is lower still, but its reachability by Cube trajectories is unknown. These audits justify the explicit $\mathcal{A}_L$ subscript; full endpoint and coupling details are in Appendix G.

Robustness and real-state scope. The all-triple, no-sentinel, no-retirement checks retain positive normalized myopic regret for $k = 2, 3, 4$ (6.52%, 9.60%, and 4.52%), but the non-monotone sequence is finite evidence rather than an all-size theorem. In 180 fixed-protocol Banknote, Wine-red, and Ionosphere windows, every state is a sound certified failure, mean

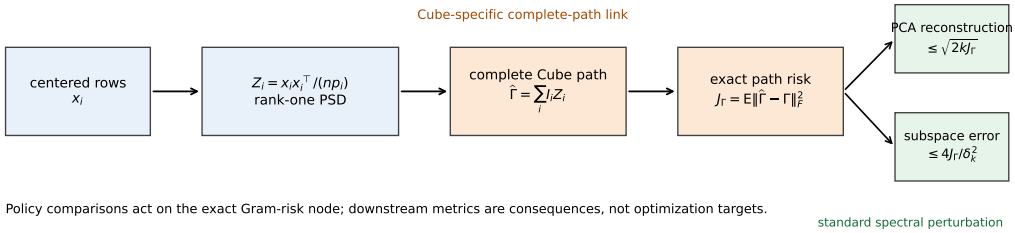

Figure 1: Connection to Gram estimation and PCA. The orange link is the Cube-specific complete-path result; the green implications follow from standard spectral perturbation arguments. The primary endpoint remains the prespecified Gram risk.

$J_\mu/J^* = 1.0747$, and no 1.02-safe skip is certified. The untouched WDBC replication likewise gives 60 sound failures (mean 1.0604); H2/H3 witnesses improve on myopic, but the fixed mapping forces $|F| \leq 1$ and therefore cannot test nontrivial success prevalence. Appendix G reports the full tables, tightness, and scope restrictions.

C Incremental Gram-Cube: Algorithm, Derivation and Evaluation

C.1 A complete sampling procedure

Assume $0 < \pi_i < 1$, $\sum_i \pi_i = k \in \mathbb{Z}$, finite feature rows, and strictly ordered auxiliary values. Integral initial coordinates may be handled outside this interface; the measured implementation requires all initial coordinates to be fractional.

1. Form $y_i = z_i/\sqrt{\pi_i}$ and $t_i = \|z_i\|^2/\pi_i$ once per draw. Order indices by increasing (t_i, i). Initialize $p = \pi$ and read the first four fractional units into the FIFO buffer.
2. For a full buffer, sort its local identifiers by auxiliary order. Set $G_{ii} = t_i^2$. Copy old Gram entries for surviving identifier pairs and compute all remaining off-diagonal entries as $(y_i^\top y_j)^2$. Store the current identifiers and Gram entries for the next update.
3. Center and scale the four auxiliary values locally. Form the endpoint basis B, $M = B^\top G_W B$ and $H = B^\top B$. Compute v by Eq. (10), with the stated zero/tie convention, then set $u = Bv$. Normalize and orient u with the reference rule.
4. Compute the maximal feasible positive and negative steps in the unit box. Choose $p + \alpha^+ u$ with probability $\alpha^-/(\alpha^+ + \alpha^-)$, and $p - \alpha^- u$ otherwise. Remove integral coordinates, retain surviving units in FIFO order, and refill from the unread queue. Repeat while a full buffer is available.
5. With three remaining fractional coordinates, use their unique mass–auxiliary circuit and the same endpoint coin. With two, drop the auxiliary constraint and apply mass-preserving pairwise landing. In exact arithmetic, a single fractional coordinate cannot remain when the total mass is integer; the existing numerical implementation also retains its single-coordinate completion routine.
6. Return selected indices and $\widehat{S} = \sum_{i:I_i=1} z_i z_i^\top/\pi_i$. Accumulate the upper triangle and mirror it, as in the comparators. The denominator is always the initial probability, not the terminal or current fractional state.

In exact arithmetic each randomized step is mean preserving, fixes at least one fractional coordinate, and preserves mass. Before pairwise landing it also preserves the scalar auxiliary balance. Termination takes at most N updates and yields k selected units with marginals π; linearity then gives $\mathbb{E}\widehat{S} = S$. These are standard Cube/martingale properties under this interface. Auxiliary balance may be relaxed at landing, and no zero final imbalance is asserted.

C.2 Proof of Proposition 8

Each column of B sums to zero and has zero inner product with s. Its middle two rows form the identity, so B has rank two. Because $(\mathbf{1}, s)$ has rank two, its nullspace is exactly the column space of B. In particular $H = B^\top B$ is positive definite. The endpoint pair has uniquely maximal separation when the four scores are strictly ordered; in exact arithmetic it is the pair found by the reference's maximum-span search. Thus the basis change removes search work without changing this reference basis.

For cache correctness, the initial probabilities and features are fixed for a complete draw. Therefore a pair's squared inner product does not change when p_t changes. Copying an entry with the same two identifiers is exact. Of the six unordered pairs in the next four-unit buffer, $\binom{r}{2}$ have both endpoints among the r survivors, giving the count $6 - \binom{r}{2}$. This accounts only for feature inner products, not all floating-point operations or total draw time.

For the direct solver let $K = \operatorname{adj}(H)M$. It is a positive multiple of $H^{-1}M$, which is similar to the real symmetric matrix $H^{-1/2}MH^{-1/2}$. Its eigenvalues are real. Its diagonal difference is Δ, and its off-diagonal entries are b, c, so their separation is $D = \sqrt{\Delta^2 + 4bc}$. Substitution in $(K - \lambda_- I)v = 0$, with $2\lambda_- = \operatorname{tr} K - D$, verifies both branches of Eq. (10). The sign choice avoids subtracting two nonnegative quantities in the selected diagonal component. If the selected vector vanishes, then $D = \Delta = 0$ and K has a repeated eigenvalue. Since K is similar to a symmetric matrix, it is diagonalizable; hence K is scalar and M is proportional to H. Every coefficient vector then minimizes the quotient and the declared default is admissible.

When minimizers and tie conventions agree, both implementations produce the same unoriented direction and maximal endpoint law at each common state. Coupling their oriented directions and endpoint coins inductively gives the same path and landing distribution. The argument does not equate the four-unit rule with the all-triple action generator in the success–failure examples. Nor does exact arithmetic settle numerical tie choices, normalization roundoff or endpoint snapping in binary64.

C.3 Cost and interpretation

Feature preparation costs $O(Nd)$, trace sorting costs $O(N \log N)$, and at most N constant-size buffer updates each require at most six d-dimensional inner products. Remaining local operations are constant size. Dense output costs $O(kd^2)$; adding these terms yields the claimed bound. The cache itself is constant size, whereas the implemented transformed feature array is $O(Nd)$, the state/order storage is $O(N)$, and the dense output is $O(d^2)$. The method is not a global $N \times N$ Gram precomputation. The two newly measured algebraic optimizations do not reduce the asymptotic order relative to the cached reference, and the 6-to-3 dot-product count is not an independently measured runtime factor in the V103 comparison.

C.4 Inputs, replication and complete-draw endpoint

The synthetic group has isotropic $(N, d) = (64, 16)$, anisotropic $(128, 16)$, and near-rank-one $(256, 32)$ features with noise levels 0.01 and 0.03. The real group has Iris $(148, 4)$, Wine $(176, 13)$, Breast Cancer $(128, 30)$ and Digits $(128, 64)$ subsets. The first three are standardized; Digits is divided by 16. Labels are unused. The larger synthetic group has isotropic $(512, 32)$, anisotropic $(1024, 64)$, isotropic $(512, 128)$ and near-rank-one $(1024, 64)$ cases, the latter with noise 0.03. Each family uses both equal $\pi_i = 1/4$ and a permutation of $(1, 3, 5, 7)/16$ repeated to fill the population, so $k = N/4$. The scalar auxiliary variable is the rank after ordering by the first raw feature with index ties. These are declared development inputs, not estimates of performance over all natural datasets.

The first two groups use 512 quality draws per case and method, 128 warmup draws, and five timing batches of 512 draws. The larger group uses 128 quality draws, 32 warmups and five batches of 128 draws. The eight methods therefore yield 73,728 quality records in V103; the direct IGC–cached-reference contrast contains 9,216 paired draws. Quality and

timing seed bases are respectively $(26000000001, 26100000001)$, $(27000000001, 27100000001)$ and $(28000000001, 28100000001)$ for the three groups. Each case/method timing batch uses the evaluator's fixed seed schedule. New seeds do not make the exposed inputs a holdout.

The time ratio in Table 2 is $\sum_{c=1}^{8}\sum_{r=1}^{5} T_{\text{candidate,c,r}} / \sum_{c=1}^{8}\sum_{r=1}^{5} T_{\text{reference,c,r}}$. Thus absolute-duration workloads receive more weight; a case-wise table is supplied below. Timing is descriptive: five rounds on one host do not support a claim about a hardware-independent speedup distribution. For comparisons with different sampling rules, exploratory quality intervals use 32 paired batches of 16 draws, or 16 batches of eight for the larger group. Their pointwise t intervals are unadjusted and conditional on these fixed inputs. An interval crossing zero is not evidence of equivalence or noninferiority.

The reported 9,216 equal outcomes are endpoint/risk comparisons. In addition, the evaluator observes four ordinary trajectories per case/method, checks their actual endpoint probabilities and constraints, and recomputes the initial-probability dense estimator independently. Those checks establish the scope tested, not an exhaustive distributional proof. The earlier and current papers' offline planning work proxies are not converted to online speedup factors.

Table 3: Case-level V103 timings. Microseconds are median batch time per draw; the ratio is the median of five paired round ratios, so it need not equal the ratio of the two displayed medians. MSE ratios are IGC/reference.

Case (N/d)	Design	Ref. μs	IGC μs	Time ratio	MSE ratio
Iso 64/16	Equal	10.84	9.35	0.8615	1.000
Iso 64/16	Unequal	10.54	9.09	0.8625	1.000
Aniso 128/16	Equal	21.62	18.80	0.8754	1.000
Aniso 128/16	Unequal	22.21	18.69	0.8483	1.000
Near-rank-one .01	Equal	57.58	50.13	0.8659	1.000
Near-rank-one .01	Unequal	56.53	50.16	0.8879	1.000
Near-rank-one .03	Equal	57.73	51.68	0.8953	1.000
Near-rank-one .03	Unequal	57.77	50.59	0.8773	1.000
Iris 148/4	Equal	21.61	17.71	0.8261	1.000
Iris 148/4	Unequal	20.96	17.40	0.8312	1.000
Wine 176/13	Equal	28.73	25.60	0.8873	1.000
Wine 176/13	Unequal	29.67	25.15	0.8485	1.000
Breast 128/30	Equal	29.72	25.91	0.9035	1.000
Breast 128/30	Unequal	28.08	24.92	0.8894	1.000
Digits 128/64	Equal	50.13	45.70	0.9116	1.000
Digits 128/64	Unequal	49.57	46.01	0.9259	1.000
Iso 512/32	Equal	123.30	111.64	0.8817	1.000
Iso 512/32	Unequal	122.16	107.76	0.8752	1.000
Aniso 1024/64	Equal	382.53	355.44	0.9323	1.000
Aniso 1024/64	Unequal	384.53	356.07	0.9286	1.000
Iso 512/128	Equal	419.40	409.13	0.9775	1.000
Iso 512/128	Unequal	424.34	406.29	0.9575	1.000
Near-rank-one 1024/64	Equal	382.62	355.19	0.9313	1.000
Near-rank-one 1024/64	Unequal	385.56	350.07	0.9080	1.000

C.5 Research scope

IGC retains the full local contribution Gram. Its computational improvement is complementary to complete-path risk analysis, but does not repair a suboptimal policy. The protected-block guarantee extends to another action class only when its required terminal invariants hold; the strict-open finite-policy certificate requires its own action-class analysis. The supplementary low-dimensional approximations are faster on some workloads but change the direction geometry. They are not used to claim the same quality guarantee as the reference computation.

This integration uses existing theory and previously executed experiments. No new experiment was run during manuscript assembly. Generalized eigenvector algebra, Gram caching and bounded-buffer Cube flight are established computational ideas; our additional contribution is their explicit specialization, operation accounting and measured implementation improvement for this Gram-aware rule. Independent-data and cross-hardware validation, an end-to-end combination with nonmyopic planning, and a broad equal-quality SOTA comparison remain unestablished.

D Frozen Gram/PCA Endpoint Plot

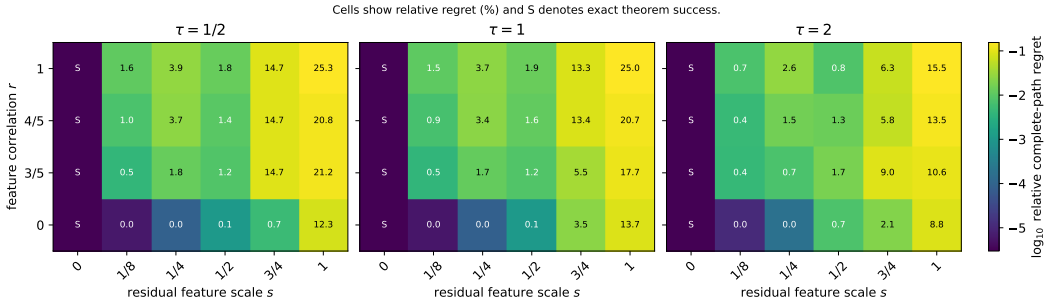

Figure 2: Frozen rank-one PSD phase diagram. Cells report complete-path relative reduction $(J_\mu - J^*)/J_\mu$; S marks theorem-success. The 72 executions contain 63 distinct geometries (12 success cells, three success geometries); the family had a development rehearsal.

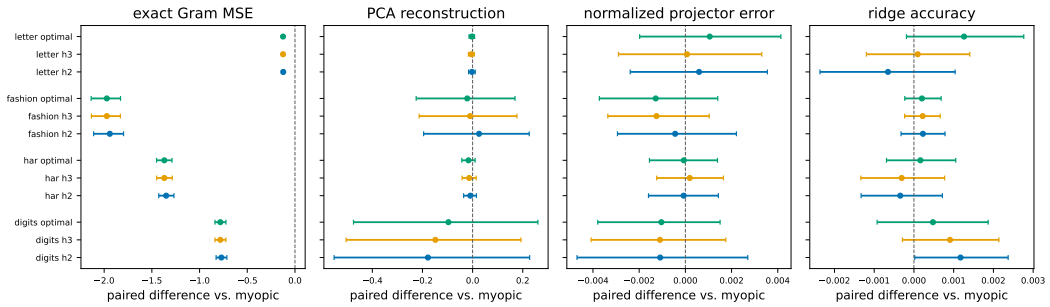

Figure 3: Four selected frozen endpoint contrasts relative to myopic (sampled Gram MSE is reported separately in Table 5). Points are partition-paired means and bars are dataset-wise 95% paired cluster-bootstrap intervals. Exact Gram risk improves consistently; the sampled PCA, subspace, and ridge endpoints are mostly unresolved.

E Executed Implementation Protocol and Deviations

Data, preprocessing, and fitting objects. Digits uses a label-stratified 70/30 split with master seed 20260809; HAR uses its official subject-disjoint files, Fashion its official 60,000/10,000 split, and Letter the conventional first 16,000/final 4,000 split. Fashion pixels are averaged over the eight contiguous row and column bins produced by array_split on each 28-pixel axis, yielding 64 features, then divided by 255. Dimensions above 64 (HAR: 561) are first projected with a fixed Gaussian PCG64 draw, thin QR and column signs chosen so the largest-absolute entry is positive. Its seed is $20260809 + \sum_j (j+1)\,\mathrm{ord}(name_j)$. The fixed matrix is saved before evaluation. The training prefix is trimmed to a multiple of six; only its mean and standard deviation standardize both training and test rows (scales $\le 10^{-12}$ become one).

To distinguish roles, denote the scalar Cube balance score by a_i here and feature vectors by $f_i \in \mathbb{R}^d$; the earlier scalar-auxiliary sections use x_i for a_i, whereas the Gram specialization uses x_i for f_i. Each six-row block is sorted by $f_i^T v$, $v_j = (-1)^j$ (zero-based j), then original retained row index. The ordered slots receive $p = (1/2, 1/2, 1/4, 1/4, 1/4, 1/4)$ and $a = (3, 1, -3, 2, -2, 0)$; the balance rows are $(1, \dots, 1)$ and a. Every legal support-minimal triple is offered; lexicographic ties were registered, but the original floating scorer implemented operational numerical ties (Appendix I). Landing is the registered exact-mass pairwise kernel.

Partition q uses PCG64 seed $20260809 + 1009q$ and its first 32 complete six-row blocks (192 rows). Block streams are independent and policy-coupled: SeedSequence([20260809,q,replicate,block]) supplies 16 uniforms, indexed by realized transition number. The primary rank is $k = \min(20, \lfloor d/2 \rfloor)$. The sketch determines only a PCA basis; one-vs-all ridge with fixed penalty 10^{-2} is fit on all retained training rows projected on that basis, with centered one-hot targets and their training class prior, and evaluated on the full test set. The 192-row window is the Gram estimand, not the downstream fitting sample.

Registration versus execution. The following deviations describe the original frozen run; supplementary completion must be recorded separately rather than rewriting its summaries. The retained/original training-row counts are Digits 1,248/1,252 (four omitted), HAR 7,350/7,352 (two), Fashion 60,000/60,000 (zero), and Letter 15,996/16,000 (four). Processed dimensions are 64, 64, 64, and 16, so their primary ranks are 20, 20, 20, and 8.

Table 4: Registered checks and actual frozen evidence. Pending items do not support completed-audit claims.

Registered	Executed	Deviation / reason	Claim impact
Primary and two sensitivity ranks	Primary rank only	Sensitivity loop omitted in original run	Main-rank findings retained; no rank-robustness claim
Rank-one samplewise PCA inequalities	Weighted endpoint expectations	Original executor accumulated expectations only	Zero violations refers to expectation checks only
Real-block sampled PCA inequalities	Per-sampled-sketch check	No identified execution omission	Zero sampled violations at primary rank
Runtime/action counts	Full-policy flight-cost cache misses and dataset total runtime	No isolated cold-cache trajectory timing	Offline work proxy only; no online latency claim

Table 5: Sampled Gram MSE: policy-minus-myopic differences and dataset-wise 95% paired partition-cluster intervals. Frozen propagation endpoints, not exact path risks. No simultaneous coverage is asserted.

Dataset	H2	H3	exact Bellman
Digits	$9.349\ [-25.069, 41.268]$	$3.908\ [-29.559, 35.777]$	$-3.156\ [-37.343, 29.948]$
UCI HAR	$-1.524\ [-3.653, 0.636]$	$-1.865\ [-4.134, 0.387]$	$-1.050\ [-3.377, 1.644]$
Fashion	$-0.473\ [-7.265, 8.975]$	$1.296\ [-4.134, 9.635]$	$1.565\ [-4.096, 9.993]$
Letter	$-0.029\ [-0.129, 0.080]$	$-0.076\ [-0.162, 0.019]$	$-0.087\ [-0.176, 0.015]$

Supplemental completion, 15 September 2026. The original omissions above remain historical deviations. Supplemental C2 now audits every endpoint in all 72 executions (63 distinct geometries), comprising 3,357 policy–terminal records. Reconstruction and subspace violations and exact identity/marginal violations are all zero; every quantity is finite and every target subspace bound applicable. The role is a deterministic supplemental replay using the same evaluator, followed by output-level parsing checks, not a second independent solver. See validation/endpoint_summary.json, validation/endpoint_output_audit.json and the full validation/endpoint_records/endpoints records. C1 is now complete under the separately versioned exact-mathematical-tie implementation: all four datasets, 20 partitions, 32 blocks and 32 coupled replicates, four policies and three ranks. The ranks are 20/10/40 for Digits, HAR and Fashion, and 8/4/15 for Letter. The 30,720 sampled records and 10,240 exact block records have zero sampled PCA-bound violations. Every dataset output and all 240 sidecars pass output-contract audits; no nonfinite quantity or nonapplicable target subspace bound occurs. These checks establish completed scope and stated inequalities, not stable downstream improvement or restored original path parity. See validation/rational_full_audits and Appendices I–J. Neither supplement rewrites the original frozen summaries.

F Detailed Theory-Targeted Audits

Testing the success conditions and crossing threshold. Twenty-nine focused exact unit tests cover the energy interval, ratio guarantee, and three-way label soundness. A rational search over 240 positive-definite instances found two exact myopic successes and one strict point; the explicit certificate above has 21 reachable states and positive margins throughout. A leakage grid has zero bound underestimation, while a separate search finds seven valid-Gram witness crossings among 240 fixed candidate directions.

A separate theory-targeted phase diagram uses six real-matrix Gram anchors, the fixed protected set $F = \{3, 5\}$, all 20 legal root triples, seven residual leakage levels, and seven exact rational correlation settings. Across all 294 exact Bellman grid executions (258 distinct Gram geometries, because all seven correlation settings coincide at zero leakage) there are zero additive-bound, ratio-bound, or three-way-label violations. All 42 zero-leakage points are certified exact successes. Every tested nonzero leakage point has a strict but sometimes tiny witness gap; nevertheless, at $\lambda = 2^{-10}$ all 42 points remain certified within 1.02. This coverage decreases to 38, 8, 4, 0, and 0 as λ grows through $2^{-8}, 2^{-6}, 2^{-4}, 2^{-2}, 1$. The experiment is real-anchored semisynthetic mechanism validation, not a prevalence claim. Of the 252 strict nonzero-leakage failures, 236 already violate root Bellman optimality, while 16 have a Bellman-optimal root action but fail later on the myopic trajectory. This directly checks the complete-path condition in Eq. 1.

Table 6: Theory-targeted leakage phase diagram (42 cases per level).

λ	mean J_μ/J^*	max J_μ/J^*	certified ≤ 1.02
0	1.000000	1.000000	42
2^{-10}	1.000168	1.000693	42
2^{-8}	1.000673	1.002768	38
2^{-6}	1.002679	1.011033	8
2^{-4}	1.010690	1.043386	4
2^{-2}	1.034640	1.154624	0
1	1.057490	1.220194	0

G Supporting Robustness and Real-State Audits

Action-class audit. With seven initially fractional coordinates, a six-coordinate buffer exposes 20 root triples rather than all 35. Across 12 exact systems, two have nonzero $V^*_{\mathcal{A}_6} - V^*_{\text{global}}$; the larger relative gap is 4.67%. The exact outer relaxation over all fixed-size-two terminal distributions with the same marginals attains 165/16, below the strict circuit value 585029/34100. Reachability of that optimizer by a larger Cube action class is not established, so this is an omission bound rather than a realized policy gap. Complete-circuit policies reach 14/15 binary endpoints and include every endpoint used by the outer optimizer individually; the residual gap is in trajectory-level coupling of endpoint probabilities.

All-triple robustness. Removing retirement, sentinels, and restricted triples leaves positive myopic regret for $k = 2, 3, 4$, with exact marginal and mass checks. Both the normalized and additive sequences are nonmonotone, so these are finite robustness results rather than an all-size theorem.

Fixed-protocol real states. The protocol uses 60 windows each from UCI Banknote, Wine-red, and Ionosphere. All 180 are sound certified failures; none is an exact success or a certified 1.02-safe skip. The median energy-bound/exact-regret ratio is 32.08 and no action evaluation is safely skipped. One smoke window per dataset was inspected before the run, so this evidence is validation rather than untouched evaluation. Exhausting all 15 probability-to-rank assignments shows that the fixed mapping always has $|F| \leq 1$; it therefore tests failure-side applicability, not nontrivial success prevalence.

Table 7: Theorem–experiment bridge. Each layer answers a different question.

Layer	Establishes	Does not establish
Theory	exact/open success; branch/landing failure	prevalence or all-size behavior
Leakage phase	finite-grid consistency and 1.02 coverage	natural prevalence
Robustness/real states	finite survival and applicability	theorem truth or deployment
Gram phase/blockwise	rank-one success/failure and exact-risk gains	downstream or efficient control

Table 8: No-retirement, no-sentinel, all-triple check.

k	N	gap/myopic	myopic states	comparator states
2	12	6.52%	455	391
3	18	9.60%	24,250	14,274
4	24	4.52%	534,207	460,674

Untouched WDBC replication. All 60 windows are sound certified failures with zero label or energy-bound violations. Mean, median, and maximum J_μ/J^* are 1.0604, 1.0555, and 1.1642. The preregistered H2/H3 witnesses improve on myopic in every window, and the best named witness matches the exact optimum. There are no certified 1.02 skips, and the same $|F| \le 1$ mapping restriction applies.

H Full Proofs and Exact Certificates

This appendix proves every theorem used in the main text for the registered finite action class. Generic martingale, Bellman, Fréchet, and convex- analytic steps are included so that the scope can be checked; they are not claimed as new results by themselves.

H.1 Complete flight-plus-landing identity and performance difference

Let $D < \infty$ be a deterministic upper bound on the flight stopping time τ, and write $M_t^\tau = M_{t\wedge\tau}$. For $\Delta_{t+1} = M_{t+1}^\tau - M_t^\tau$, mean preservation gives $\mathbb{E}[\Delta_{t+1} \mid \mathcal{F}_t] = 0$. The registered landing kernel is also mean preserving, so for $R_{\text{land}} = \widehat{M} - M_\tau$, $\mathbb{E}[R_{\text{land}} \mid \mathcal{F}_\tau] = 0$. If $s < t$, the tower property gives $\mathbb{E}\,\langle \Delta_{s+1}, \Delta_{t+1}\rangle = 0$, and conditioning on $\mathcal{F}_\tau$ gives $\mathbb{E}\,\langle \Delta_{s+1}, R_{\text{land}}\rangle = 0$. Expanding the squared norm of

$$\widehat{M} - M_0 = \sum_{t=0}^{D-1} \Delta_{t+1} + R_{\text{land}}$$

therefore proves the complete flight-plus-landing objective and $\mathbb{E}[\widehat{M}] = M_0$.

For completeness, fix the Bellman boundary value $V^*(S_\tau) = g(S_\tau)$ equal to the expected loss of the fixed landing kernel. Conditional on S_t,

$$c(S_t, \mu(S_t)) + \mathbb{E}[V^*(S_{t+1}) \mid S_t] - V^*(S_t) = A^*(S_t, \mu(S_t)).$$

Summing to the bounded stopping time and taking expectations telescopes the value terms and yields Eq. (1). Every advantage is nonnegative. Thus the expectation is zero if and only if the advantage is zero almost surely under the state–time occupancy measure. On a finite registered tree, this is equivalent to zero advantage at every nonterminal state–time pair reached with positive probability. This also proves why root optimality alone is insufficient.

H.2 Residual feasibility, exact success, and energy stability

Residual-feasibility theorem. Flight preserves both $\sum_i p_i = m$ and $\sum_i p_i x_i = b$. Suppose a protected coordinate $i \in F$ reached one at time t during flight. The remaining coordinates

Table 9: Real-state validation; label and energy-bound violations are zero.

Dataset	windows	failure labels	mean J_μ/J^*	max J_μ/J^*
Banknote	60	60	1.1307	1.4341
Wine-red	60	60	1.0511	1.1352
Ionosphere	60	60	1.0423	1.1728
Aggregate	180	180	1.0747	1.4341

$q_j = p_{t,j}$, $j \neq i$, would then satisfy

$$0 \le q_j \le 1, \qquad \sum_{j\neq i} q_j = m-1, \qquad \sum_{j\neq i} q_j x_j = b - x_i,$$

contradicting the definition of F. Hence no protected coordinate is fixed to one before landing. At the registered stopping time at most two coordinates are fractional. Integral mass rules out exactly one strictly fractional coordinate; if two remain, their fractional parts sum to one and the registered pairwise landing selects exactly one. Consequently at most one protected coordinate is selected and $\sum_{i\in F} I_i \le 1$ almost surely.

Mean preservation gives $\mathbb{E}[I_i] = p_i$. For distinct $i, j \in F$, mutual exclusion gives $\mathbb{E}[I_i I_j] = 0$ and hence $\mathrm{Cov}(I_i, I_j) = -p_i p_j$, while $\mathrm{Var}(I_i) = p_i(1-p_i)$. If the contributions vanish outside F, expanding the terminal squared error gives

$$\begin{aligned} J(\pi) &= \sum_{i\in F} p_i(1-p_i)\,\|Z_i\|_F^2 - 2\sum_{\substack{i<j\\ i,j\in F}} p_i p_j\,\langle Z_i, Z_j\rangle \\ &= \sum_{i\in F} p_i\,\|Z_i\|_F^2 - \left\|\sum_{i\in F} p_i Z_i\right\|_F^2, \end{aligned}$$

which is policy invariant and proves exact success.

Energy interval. Put $Y_F = \sum_{i\in F}(I_i - p_i)Z_i$ and $Y_R = \sum_{i\in R}(I_i - p_i)Z_i$. The previous covariance calculation fixes $\mathbb{E}\,\|Y_F\|_F^2 = J_F$. For Bernoulli variables with fixed marginals,

$$\max(0, p_i + p_j - 1) \le \mathbb{E}[I_i I_j] \le \min(p_i, p_j).$$

After subtracting $p_i p_j$, the absolute covariance is at most β_{ij}. Expanding the residual energy and upper bounding signed cross terms by their absolute values gives $0 \le \mathbb{E}\,\|Y_R\|_F^2 \le \overline{R}$.

Regard Y_F, Y_R as elements of the Hilbert space of square-integrable random matrices. Triangle and reverse-triangle inequalities yield

$$\sqrt{J(\pi)} \le \sqrt{J_F} + \sqrt{\overline{R}}, \qquad \sqrt{J(\pi)} \ge (\sqrt{J_F} - \sqrt{\overline{R}})_+,$$

and hence $L_E \le J(\pi) \le U_E$.

All policies have the same marginal variances and the same $F \times F$ covariance block. Every policy-dependent Gram term therefore corresponds to a pair $\{i, j\} \not\subseteq F$. The width of the admissible joint-inclusion interval for that pair is w_{ij}, so any two policy values differ by at most B_{Fr}. They also differ by at most $U_E - L_E$. Thus

$$0 \le J(\mu) - J^*_{\Pi_L} \le B = \min\{B_{\mathrm{Fr}}, U_E - L_E\}.$$

When $L_E > 0$, division by $J^*_{\Pi_L} \ge L_E$ proves the stated ratio.

To prove nonemptiness and full dimension, choose a nonzero contribution at one protected coordinate i with $0 < p_i < 1$ and initially set every other contribution to zero. Then $J_F > 0$ while $\overline{R} = B = 0$. The strict conditions $L_E > 0$ and $B/L_E < \varepsilon$ hold at this point. All displayed quantities are continuous functions of the finite contribution tuple, so the conditions persist on an open ambient neighborhood. Independent small perturbations

make every residual contribution nonzero. If $R \neq \varnothing$ this open set is not contained in the positive-codimension residual-zero stratum.

Finally, the argument used only fixed terminal marginals and mutual exclusion on F, not the mechanism that selected a flight direction. If these two invariants hold for an expanded class $\Pi' \supseteq \Pi_L$, the same interval, additive bound, and ratio apply to every policy in Π' and therefore to $J^*_{\Pi'}$. This proves the action-class robustness corollary. It does not extend certificates whose attainable covariance set changes with the class.

H.3 Strict open success and computable radius

We first prove the general stability statement behind the explicit rational certificate. Because the registered history tree and action sets are finite, a randomized history policy induces a convex combination of the terminal covariances induced by deterministic history policies. Their convex hull is therefore the finite covariance polytope K. Let Σ_0 be the unique covariance point minimizing $\langle G_0, \Sigma\rangle$. Define

$$\gamma = \inf_{\Sigma\in K, \Sigma\neq\Sigma_0} \frac{\langle G_0, \Sigma - \Sigma_0\rangle}{\|\Sigma - \Sigma_0\|_F}.$$

The normalized objective gap is positive on every other vertex. Its finite minimum is positive, and convexity plus the triangle inequality extends the same lower margin to all $\Sigma \in K\setminus\{\Sigma_0\}$.

For each state-action pair let $D(s,a)$ be the symmetric coefficient matrix defined by $c_G(s,a) = \langle G, D(s,a)\rangle$. At every nonterminal state reached under the base myopic policy, define the normalized immediate-action margin

$$\eta = \inf_{s,a\neq\mu_0(s)} \frac{\langle G_0, D(s,a) - D(s,\mu_0(s))\rangle}{\|D(s,a) - D(s,\mu_0(s))\|_F} > 0.$$

For a symmetric perturbation E with $\|E\|_F < \min(\gamma,\eta)$, Frobenius Cauchy–Schwarz gives

$$\langle G_0 + E, \Sigma - \Sigma_0\rangle \geq (\gamma - \|E\|_F)\,\|\Sigma - \Sigma_0\|_F > 0$$

and the analogous positive inequality for every on-path immediate-action comparison. Induction over the finite horizon shows that the perturbed myopic rule follows the same reachable tree, while Σ_0 remains the unique optimal covariance. Weyl's inequality adds positive definiteness whenever $\|E\|_2 < \lambda_{\min}(G_0)$.

Acyclic Markovization and the computable exposure margin. The registered state contains all sampling probabilities and determines the admissible actions, transition probabilities, stopping, costs and fixed landing law. A maximal flight fixes at least one fractional coordinate and never changes integral ones. Thus its fractional-coordinate rank $k(s)$ strictly decreases on every flight edge: states may merge but are never revisited. All state and action sets are finite.

For any randomized history policy, let $r_\pi(s)$ be its aggregate probability of visiting state s and $f_\pi(s,a)$ its state–action flow. Set $q_\pi(a \mid s) = f_\pi(s,a)/r_\pi(s)$ when $r_\pi(s) > 0$, with arbitrary legal choices elsewhere. Induction in decreasing rank shows that this randomized state policy reproduces all incoming and outgoing flows, hence the same landing-state and binary-terminal laws. Independently drawing its action at every state before sampling gives a mixture of deterministic state policies with the same law, since no state is revisited. Covariance is affine in this mixture because all policies have the same Bernoulli marginals. This is a standard Markovization argument; it does not equate one history's probability with an aggregate merged-state probability.

Let $r_\mu(s)$ denote the merged-state visit probability under μ_0 and put

$$\Delta_{\mathrm{DP}} = \min_{\substack{s\in\mathrm{Reach}_{\mathrm{nt}}(\mu_0)\\ a\neq\mu_0(s)}} r_\mu(s)\big(Q^*_{G_0}(s,a) - V^*_{G_0}(s)\big) > 0.$$

For a deterministic state policy ν with covariance different from Σ_0, choose a μ_0-reachable deviation state with maximal rank. At every higher-rank reachable state the policies agree,

so all incoming flow at that rank is identical under their laws. Deviations elsewhere at the same rank cannot affect this state. Consequently $r_\nu(s) = r_\mu(s)$. The performance-difference identity sums nonnegative Bellman advantages; this state's contribution alone is at least Δ_{DP}. Hence this deterministic-state-policy root gap is at least Δ_{DP}.

Fixed Bernoulli marginals make covariance diagonals identical, and every remaining difference is at most the joint-inclusion interval width, at most 1/2. Therefore $\|\Sigma_\nu - \Sigma_0\|_F < N/2$ and

$$\langle G_0, \Sigma_\nu - \Sigma_0 \rangle \geq \frac{2\Delta_{\mathrm{DP}}}{N} \|\Sigma_\nu - \Sigma_0\|_F .$$

The inequality includes identical covariance via zero equality. Writing any randomized history covariance as $\Sigma = \sum_\nu w_\nu \Sigma_\nu$ and using linearity and the triangle inequality gives

$$\langle G_0, \Sigma - \Sigma_0 \rangle \geq \frac{2\Delta_{\mathrm{DP}}}{N} \sum_\nu w_\nu \|\Sigma_\nu - \Sigma_0\|_F \geq \frac{2\Delta_{\mathrm{DP}}}{N} \|\Sigma - \Sigma_0\|_F .$$

Thus $2\Delta_{\mathrm{DP}}/N \leq \gamma$ for all policies. Their unnormalized randomized gaps may approach zero; no uniform positive unnormalized gap is claimed. The margin implementation processes decreasing rank and accumulates all incoming flows before a state is processed, exactly matching this merged-state occupancy definition.

Outward rational upper bounds on the immediate coefficient norms similarly give a rational lower bound on η. If $G_0 = ZZ^\top$, then

$$\frac{\det(G_0)}{\mathrm{tr}(G_0)^{N-1}} \leq \lambda_{\min}(G_0),$$

because every eigenvalue is at most the trace. For the matrix in the main text, exact enumeration gives 21 reachable states, minimum raw immediate and Bellman gaps 1/16 and 5/48, minimum reach-weighted Bellman gap 5/144, and hence

$$\underline{\gamma} = \frac{5}{432}, \quad \underline{\eta} = \frac{62500000000}{437165050353}, \quad \underline{\lambda} = \frac{26244}{309512375}.$$

Their minimum is the radius in the theorem. Since Z is lower triangular with nonzero diagonal, $G_0 = ZZ^\top \succ 0$. The verification script checks every legal action comparison at all 21 states reached by the certified policy, exact equality of myopic and optimal values, uniqueness of the immediate and Bellman minimizers, and the three rational lower margins displayed above. Its assertion is in tests/test_success_failure.py. From the artifact root, run PYTHONPATH=src python3 scripts/verify_paper_certificates.py; it checks this and the failure certificate below.

H.4 Normal cone, three-way labels, and crossing threshold

For any registered policy π, the terminal Gram–covariance identity is $J_G(\pi; s) = \langle G, \Sigma_\pi(s) \rangle$. A linear functional has the same infimum on a set and its convex hull, so Σ_μ is optimal exactly when

$$\langle G, \Sigma - \Sigma_\mu \rangle \geq 0 \quad \text{for every } \Sigma \in K(s),$$

which is the stated inward normal-cone condition.

Outer and witness feasibility give the chain

$$L_{\mathrm{out}} \leq J_G^*(s) \leq U_{\mathrm{wit}} \leq J_G(\mu; s).$$

If the left endpoint equals the myopic value, equality holds throughout and myopic is optimal. If a feasible witness is strictly smaller, myopic is strictly suboptimal. The labels cannot overlap, and leaving every other state ambiguous is therefore sound. Subtracting the chain from the myopic value proves the regret bracket; if $L_{\mathrm{out}} > 0$, division proves the ratio bound.

For the action-stable crossing, fixed transition and landing kernels make the terminal covariances of any fixed policies μ, w independent of t. With $\Delta\Sigma = \Sigma_w - \Sigma_\mu$,

$$J_{G_0+tH}(w) - J_{G_0+tH}(\mu) = d_0 + t d_1, \quad d_0 = \langle G_0, \Delta\Sigma \rangle , \quad d_1 = \langle H, \Delta\Sigma \rangle .$$

If $d_0 > 0, d_1 < 0$, its unique sign change is $t_* = d_0/(-d_1)$. Exact rational comparison of every on-path immediate action shows the myopic map is unchanged until t_{act}. Substituting the values in the main text proves $t_* < t_{\text{act}}$ and strict failure on (t_*, t_{act}). Since $\left\|tzz^\top\right\|_F = t\left\|z\right\|_2^2 = 3552t$, the strict-radius theorem covers exactly $t < t_{\text{succ}} = 2187/91615663000$; neither certificate decides $[t_{\text{succ}}, t_*]$.

H.5 Complete circuit actions for the scalar auxiliary geometry

For two constraint rows, the column belonging to coordinate i is $(1, x_i)^\top$. If active scores are pairwise distinct, every two columns are linearly independent and every three are dependent in $\mathbb{R}^2$. Thus every support-minimal nullspace circuit has size three. On a triple i, j, k, a direction is proportional to

$$(u_i, u_j, u_k) = (x_j - x_k, x_k - x_i, x_i - x_j).$$

Conversely, this nonzero vector annihilates both constraint rows, proving that all active triples are exactly the complete support-minimal circuit set. Repeated scores produce size-two circuits and require deduplication. Hence a buffer is complete only when it includes every active fractional coordinate; the six-coordinate theorems do so, while the seven-coordinate experiment measures the gap caused by omitting one active coordinate.

H.6 Exact branch- and landing-sensitive separation

We give a finite exact certificate rather than infer the theorem from floating-point simulation. At a state p, for every legal circuit a let p_a^+, p_a^- be the two maximal unit-box successors and let ρ_a^+, ρ_a^- be their mean-preserving probabilities. With exact landing value $L(p)$, the recursion is

$$V(p) = \begin{cases} L(p), & \text{if at most two coordinates are fractional,} \\ \min_a\{c(p,a) + \rho_a^+ V(p_a^+) + \rho_a^- V(p_a^-)\}, & \text{otherwise.} \end{cases}$$

Writing the two maximal step lengths as α_a^+, α_a^- and $D_a = \sum_i u_{a,i} Z_i$, the immediate conditional MSE is $c(p,a) = \alpha_a^+ \alpha_a^- \left\|D_a\right\|_F^2$. At an integral state $L(p) = 0$; if only r, s are fractional, exact-mass landing has $p_r + p_s = 1$ and $L(p) = p_r p_s \left\|Z_r - Z_s\right\|_F^2$. These formulas completely specify the recursion used by the certificate. Each maximal flight step fixes at least one fractional coordinate, so this is a finite acyclic recursion. All inputs, circuit directions, maximal steps, branch probabilities, flight costs, and landing values in the construction are rational; exact fraction arithmetic is therefore a proof certificate.

At the stated base point, the 20 legal root triples are exhaustive. Their exact immediate costs have the unique minimum $3/2$ at $(2,4,5)$; the next smallest is $32/5$. Exact Bellman backup instead selects $(1,2,6)$, and forcing the myopic root before continuing optimally costs an additional $369406/19173 > 0$. Under the myopic root the two successors are

$$\begin{aligned} \rho^+ = 3/5: &\quad (1/2, 5/6, 1/4, 0, 1/6, 1/4), \quad \text{optimal continuation } (2,3,6), \\ \rho^- = 2/5: &\quad (1/2, 0, 1/4, 5/8, 3/8, 1/4), \quad \text{optimal continuation } (3,4,6). \end{aligned}$$

The myopic continuation is $(1,3,5)$ on both branches, so the continuation failure is genuinely branch dependent. Repeating the exact recursion with terminal landing value set to zero selects root $(1,4,6)$, proving landing sensitivity. The complete exact values are

$$J_{\text{myopic}} = \frac{1113001}{4410}, \qquad J^* = \frac{298063}{1743}, \qquad J_{\text{myopic}} - J^* = \frac{29785853}{366030} > 0.$$

The following table records the strict comparisons needed by the computer-assisted part of the proof; "second" is the next smallest action value in the same exhaustive list.

comparison	best action/value	second action/value
root Bellman	$(1,2,6) : 298063/1743$	$(3,4,6) : 15385/88$
plus child	$(2,3,6) : 1182/7$	$(1,3,5) : 17271/98$
minus child	$(3,4,6) : 67343/308$	$(1,3,5) : 14477/60$
zero landing root	$(1,4,6) : 188627/1920$	$(1,2,6) : 399684419/4057704$

The exact verification visits 1,203 distinct flight states for the optimal recursion, 27 states under the myopic recursion, and 38 landing states; it checks all 20 root actions and the displayed rational equalities without floating-point comparisons. The verifier and its regression assertions are included in the artifact. The corresponding command uses the same test module and the short verifier command given above. These rational inequalities establish the theorem for the complete six- coordinate circuit class. The finite $k = 2, 3, 4$ replications in the main text are robustness evidence only and are not used to infer an all-size result.

I Numerical Tie Semantics and Supplemental Replay Provenance

Original numerical implementation and failed cross-host replay. The original frozen tables are retained as observations of their recorded binary64 implementation. Its raw floating-point minimum resolves numerical scores operationally; it does not implement lexicographic selection among all exactly equal mathematical action values. Cross-host primary-rank replay failed for sampled paths. Myopic primary records matched, and H3/exact Bellman expected block values matched within the registered tolerance, while their sampled endpoints differed. Across the four datasets, 15 H2 exact block values differed beyond tolerance. This limits cross-host reproducibility of the old operational policy; it is not evidence that the original recorded risk or terminal estimates were replaced by a replay.

Diagnostic comparison found zero changed instruction objects between the available original controller bytecode and current source. In the same new host environment, the original runner and supplemental rank loop matched on all 256 checked primary records. A single development block yielded 312 mathematically exact tie events: 158 floating-point selected actions differed from the declared lexicographic exact tie rule, with zero exact action-score excess at those comparisons (H2 Q is a truncated rollout score, not the actual complete-policy risk). The floating spreads ranged from 9.094947×10^{-13} to 7.275958×10^{-12}, most commonly 3.64×10^{-12}. Equal local or complete-risk values do not guarantee the same trajectory, sampled Gram error, or PCA basis. Original action histories and bases were not saved; old sampled-path parity cannot be restored from summaries and is not claimed.

Separate implementation correction and provenance. The supplemental scorer converts each entry of the already computed binary64 6×6 contribution Gram matrix to its exact binary rational value. Flight and landing costs, myopic tails, rollout Q-values and Bellman values are compared by exact rational value and then lexicographic action. This realizes mathematical ties for that actual numerical Gram objective. It does not recompute features as exact real numbers. States, admissible actions, landing transitions, seeds, policy roles, ranks and metrics remain registered; no platform core type is selected and no tolerance is enlarged. The scorer reuses the original state recursion and is not an independent second solver or a new scientific contribution.

The v1.1 output supplement records three sidecar types per partition: slot-sorted retained row IDs, the actual binary64 contribution Gram matrices, and the actual binary terminal selections on the fixed policy/replicate/block axes. These identify the realized scoring objective and sampled paths when combined with the retained source data and preprocessing protocol. Completed files are distinguished from interrupted partial outputs by checkpoint counts. These records do not guarantee bitwise PCA/eigensolver/ridge reproducibility across arbitrary BLAS platforms, nor restore missing original paths. The new full replay is separate evidence, not a replacement for the old frozen tables. The C1 full four-dataset replay is complete, with 30,720 sampled rank/policy records and 10,240 exact block records. Output-contract and sidecar audits pass for every dataset; both sampled PCA inequalities have zero violations and all reported quantities are finite. Each dataset has 60 completed sidecars (240 total), with no incomplete files. These are output-level audits, not a new independent solver or human certification. Appendix J retains every registered contrast.

See supplementary experiments/diagnostics/RATIONAL_TIE_CORRECTION_PROTOCOL_CN.md, reviews/NUMERIC_TIE_REVIEW_CHECKS.json, validation/parity_diagnostics/diagnostics/frozen_bytecode_comparison.json and the four validation/*_primary_parity.json reports.

The corrected scorer also changes H2 expected risk slightly; therefore not all old exact values are reproduced. Full old/new per-policy expected and sampled differences, stored intervals and offline work counts are retained in supplementary validation/rational_full_audits/NEW_LEGACY_COMPARISON.json.

Dataset	Original H2 mean	Corrected H2 mean
Digits	342.969522	342.970447
UCI HAR	54.8640823	54.86367
Fashion	74.2966501	74.298495
Letter	3.60848772	3.60848184

Table 10: H2 exact window-risk means under original operational and corrected mathematical tie implementations. These are distinct retained observations.

J All Registered Rank-Sensitivity Supplemental Contrasts

These are audit-corrected mathematical-tie supplemental results, separate from the original frozen floating-implementation tables. All four datasets, three ranks and H2/H3/exact Bellman comparisons are retained. Each table lists myopic absolute mean, policy-minus-myopic difference and its stored 95% paired partition-cluster interval; policy absolute means are retained in the complete supplemental Markdown and JSON exports. No bootstrap was recomputed for this presentation, no comparison was selected by its interval, and no multiplicity adjustment or simultaneous coverage is claimed. Inference is conditional on fixed data and preprocessing. Ridge accuracy is descriptive and expressed as a fraction, not percentage points. Exact Gram risk is rank independent and reported separately. Letter retains its external dataset role, but this implementation correction is not a newly untouched holdout.

Observed conditional effects. All 36 intervals for each of sampled Gram MSE, reconstruction excess, squared sine distance and rank-normalized squared sine distance cross zero. Ridge accuracy has 35 intervals crossing zero; the exception is HAR at $k = 10$ for H2, a descriptive gain of 0.00184669 with interval $[0.000432630, 0.00339231]$. With no multiplicity adjustment this isolated comparison does not establish stable downstream improvement. All 12 rank-independent exact-risk contrasts lie below zero. Thus lower conditional expected Gram risk persists under the corrected tie rule, while sampled spectral and task propagation remains largely unresolved. The original floating-implementation results and their development-only exception are retained separately, rather than overwritten by this replay.

Complete means, stored intervals and source identities are in supplementary validation/final_rank_tables/ALL_RANK_RESULTS_CN.md and validation/final_rank_tables/all_rank_results.json.

Rank	Policy	Myopic mean	Difference	95% paired CI
Digits				
20	H2	334.757	5.98574	$[-28.3999, 38.225]$
20	H3	334.757	8.51964	$[-20.5118, 36.2129]$
20	OPTIMAL	334.757	8.51964	$[-21.477, 36.1331]$
10	H2	334.757	5.98574	$[-28.3999, 38.225]$
10	H3	334.757	8.51964	$[-20.5118, 36.2129]$
10	OPTIMAL	334.757	8.51964	$[-21.477, 36.1331]$
40	H2	334.757	5.98574	$[-28.3999, 38.225]$
40	H3	334.757	8.51964	$[-20.5118, 36.2129]$
40	OPTIMAL	334.757	8.51964	$[-21.477, 36.1331]$
UCI HAR				
20	H2	55.8849	−1.42962	$[-3.34052, 0.559568]$
20	H3	55.8849	−1.47189	$[-3.29735, 0.476265]$
20	OPTIMAL	55.8849	−1.47189	$[-3.32848, 0.415282]$
10	H2	55.8849	−1.42962	$[-3.34052, 0.559568]$
10	H3	55.8849	−1.47189	$[-3.29735, 0.476265]$
10	OPTIMAL	55.8849	−1.47189	$[-3.32848, 0.415282]$
40	H2	55.8849	−1.42962	$[-3.34052, 0.559568]$
40	H3	55.8849	−1.47189	$[-3.29735, 0.476265]$
40	OPTIMAL	55.8849	−1.47189	$[-3.32848, 0.415282]$
Fashion				
20	H2	73.7764	0.418926	$[-5.93316, 9.69405]$
20	H3	73.7764	1.96616	$[-3.87143, 10.9037]$
20	OPTIMAL	73.7764	1.96616	$[-3.90282, 11.0161]$
10	H2	73.7764	0.418926	$[-5.93316, 9.69405]$
10	H3	73.7764	1.96616	$[-3.87143, 10.9037]$
10	OPTIMAL	73.7764	1.96616	$[-3.90282, 11.0161]$
40	H2	73.7764	0.418926	$[-5.93316, 9.69405]$
40	H3	73.7764	1.96616	$[-3.87143, 10.9037]$
40	OPTIMAL	73.7764	1.96616	$[-3.90282, 11.0161]$
Letter				
8	H2	3.72136	−0.0556654	$[-0.1582, 0.0500341]$
8	H3	3.72136	−0.0787191	$[-0.170213, 0.0200908]$
8	OPTIMAL	3.72136	−0.0787191	$[-0.169605, 0.0197675]$
4	H2	3.72136	−0.0556654	$[-0.1582, 0.0500341]$
4	H3	3.72136	−0.0787191	$[-0.170213, 0.0200908]$
4	OPTIMAL	3.72136	−0.0787191	$[-0.169605, 0.0197675]$
15	H2	3.72136	−0.0556654	$[-0.1582, 0.0500341]$
15	H3	3.72136	−0.0787191	$[-0.170213, 0.0200908]$
15	OPTIMAL	3.72136	−0.0787191	$[-0.169605, 0.0197675]$

Table 11: Sampled Gram MSE. All registered supplemental contrasts; policy-minus-myopic, stored unadjusted partition-cluster intervals.

Rank	Policy	Myopic mean	Difference	95% paired CI
Digits				
20	H2	9.69315	−0.140807	[−0.493335, 0.235705]
20	H3	9.69315	−0.16128	[−0.4777, 0.137435]
20	OPTIMAL	9.69315	−0.16128	[−0.483502, 0.136424]
10	H2	8.45653	−0.108539	[−0.398386, 0.213482]
10	H3	8.45653	−0.135532	[−0.404142, 0.126166]
10	OPTIMAL	8.45653	−0.135532	[−0.416028, 0.126704]
40	H2	8.87183	−0.159886	[−0.540164, 0.237575]
40	H3	8.87183	−0.18008	[−0.53274, 0.149678]
40	OPTIMAL	8.87183	−0.18008	[−0.537755, 0.147719]
UCI HAR				
20	H2	2.63342	−0.012163	[−0.0378879, 0.0137465]
20	H3	2.63342	−0.0194709	[−0.0394713, 0.0017545]
20	OPTIMAL	2.63342	−0.0194709	[−0.0394767, 0.00180094]
10	H2	2.49069	−0.016642	[−0.0552905, 0.0240439]
10	H3	2.49069	−0.0195248	[−0.0539476, 0.0165473]
10	OPTIMAL	2.49069	−0.0195248	[−0.0539896, 0.0163595]
40	H2	1.63368	−0.00135508	[−0.018072, 0.0159038]
40	H3	1.63368	−0.00404697	[−0.0235386, 0.0139528]
40	OPTIMAL	1.63368	−0.00404697	[−0.0235125, 0.0140413]
Fashion				
20	H2	2.5335	−0.00865306	[−0.232151, 0.195973]
20	H3	2.5335	−0.0120324	[−0.224528, 0.182726]
20	OPTIMAL	2.5335	−0.0120324	[−0.22532, 0.185045]
10	H2	2.58937	0.00696712	[−0.213072, 0.200746]
10	H3	2.58937	0.00840029	[−0.19315, 0.188747]
10	OPTIMAL	2.58937	0.00840029	[−0.19131, 0.190856]
40	H2	1.64869	−0.000986254	[−0.205243, 0.193826]
40	H3	1.64869	0.00982742	[−0.172626, 0.188296]
40	OPTIMAL	1.64869	0.00982742	[−0.172477, 0.190306]
Letter				
8	H2	0.432258	−0.00411495	[−0.0163151, 0.00804435]
8	H3	0.432258	0.000245906	[−0.0134039, 0.0143552]
8	OPTIMAL	0.432258	0.000245906	[−0.0135118, 0.0144443]
4	H2	0.495827	−0.0183962	[−0.0403331, 0.0040731]
4	H3	0.495827	−0.0100192	[−0.0336124, 0.0146676]
4	OPTIMAL	0.495827	−0.0100192	[−0.0344234, 0.0145246]
15	H2	0.0283482	−0.000685076	[−0.00232485, 0.00111229]
15	H3	0.0283482	0.000361135	[−0.00142516, 0.00223807]
15	OPTIMAL	0.0283482	0.000361135	[−0.00144955, 0.0022439]

Table 12: PCA reconstruction excess. All registered supplemental contrasts; policy-minus-myopic, stored unadjusted partition-cluster intervals.

Rank	Policy	Myopic mean	Difference	95% paired CI
Digits				
20	H2	4.21623	−0.0110528	[−0.0745473, 0.0544132]
20	H3	4.21623	−0.0182142	[−0.068945, 0.0355406]
20	OPTIMAL	4.21623	−0.0182142	[−0.0689739, 0.035162]
10	H2	2.53534	−0.0216104	[−0.0653601, 0.0291995]
10	H3	2.53534	−0.0289546	[−0.0794905, 0.0249819]
10	OPTIMAL	2.53534	−0.0289546	[−0.0797157, 0.0237713]
40	H2	5.98303	−0.0109359	[−0.0853706, 0.0643166]
40	H3	5.98303	−0.0225114	[−0.0883259, 0.0475384]
40	OPTIMAL	5.98303	−0.0225114	[−0.08941, 0.0454574]
UCI HAR				
20	H2	4.31022	−0.000144846	[−0.0301533, 0.0254676]
20	H3	4.31022	−0.00866784	[−0.037742, 0.0169002]
20	OPTIMAL	4.31022	−0.00866784	[−0.0386567, 0.0161395]
10	H2	2.362	−0.0108739	[−0.0491722, 0.0304261]
10	H3	2.362	−0.0124556	[−0.0488666, 0.0252832]
10	OPTIMAL	2.362	−0.0124556	[−0.0488453, 0.0260811]
40	H2	5.89648	0.00847245	[−0.0341714, 0.050197]
40	H3	5.89648	0.00208085	[−0.0434669, 0.0444231]
40	OPTIMAL	5.89648	0.00208085	[−0.0438113, 0.0447186]
Fashion				
20	H2	3.80741	−0.0148406	[−0.0618948, 0.0350753]
20	H3	3.80741	−0.0285202	[−0.0793258, 0.0234443]
20	OPTIMAL	3.80741	−0.0285202	[−0.0780405, 0.0241223]
10	H2	1.59261	0.0285106	[−0.0211832, 0.0733543]
10	H3	1.59261	0.022331	[−0.0212382, 0.0638928]
10	OPTIMAL	1.59261	0.022331	[−0.0203457, 0.0633088]
40	H2	6.04246	0.0212468	[−0.0344619, 0.07737]
40	H3	6.04246	0.0138256	[−0.0461318, 0.072148]
40	OPTIMAL	6.04246	0.0138256	[−0.0441984, 0.0746545]
Letter				
8	H2	0.857699	−0.00121994	[−0.0227638, 0.0208834]
8	H3	0.857699	0.0154285	[−0.0117347, 0.0432503]
8	OPTIMAL	0.857699	0.0154285	[−0.0109564, 0.0435059]
4	H2	0.607144	−0.0247156	[−0.0584908, 0.00918418]
4	H3	0.607144	−0.00996342	[−0.0477053, 0.0283081]
4	OPTIMAL	0.607144	−0.00996342	[−0.0470497, 0.0285751]
15	H2	0.2201	0.00197923	[−0.0165505, 0.0213385]
15	H3	0.2201	0.0103529	[−0.011261, 0.0320176]
15	OPTIMAL	0.2201	0.0103529	[−0.0106379, 0.0315019]

Table 13: Squared Frobenius sine distance. All registered supplemental contrasts; policy-minus-myopic, stored unadjusted partition-cluster intervals.

Rank	Policy	Myopic mean	Difference	95% paired CI
Digits				
20	H2	0.210812	−0.00055264	[−0.00380101, 0.00266352]
20	H3	0.210812	−0.000910711	[−0.00349513, 0.001728]
20	OPTIMAL	0.210812	−0.000910711	[−0.00347261, 0.00172442]
10	H2	0.253534	−0.00216104	[−0.00657054, 0.00300963]
10	H3	0.253534	−0.00289546	[−0.00805841, 0.0023655]
10	OPTIMAL	0.253534	−0.00289546	[−0.00803694, 0.00225605]
40	H2	0.149576	−0.000273398	[−0.00215285, 0.0016043]
40	H3	0.149576	−0.000562784	[−0.00225254, 0.0011394]
40	OPTIMAL	0.149576	−0.000562784	[−0.00222656, 0.00113807]
UCI HAR				
20	H2	0.215511	-7.24231×10^{-6}	[−0.00148558, 0.00126124]
20	H3	0.215511	−0.000433392	[−0.00187126, 0.000835193]
20	OPTIMAL	0.215511	−0.000433392	[−0.00191268, 0.000829217]
10	H2	0.2362	−0.00108739	[−0.00487441, 0.00301724]
10	H3	0.2362	−0.00124556	[−0.00488379, 0.00260335]
10	OPTIMAL	0.2362	−0.00124556	[−0.00495961, 0.0025491]
40	H2	0.147412	0.000211811	[−0.000857495, 0.00127641]
40	H3	0.147412	5.20213×10^{-5}	[−0.00105747, 0.00112503]
40	OPTIMAL	0.147412	5.20213×10^{-5}	[−0.00105014, 0.00109654]
Fashion				
20	H2	0.19037	−0.000742028	[−0.00303503, 0.00169569]
20	H3	0.19037	−0.00142601	[−0.00393512, 0.00118603]
20	OPTIMAL	0.19037	−0.00142601	[−0.00393437, 0.00123195]
10	H2	0.159261	0.00285106	[−0.00204302, 0.00734675]
10	H3	0.159261	0.0022331	[−0.00197221, 0.00633291]
10	OPTIMAL	0.159261	0.0022331	[−0.00207076, 0.00629011]
40	H2	0.151062	0.00053117	[−0.000852146, 0.00193794]
40	H3	0.151062	0.000345639	[−0.00110899, 0.0018367]
40	OPTIMAL	0.151062	0.000345639	[−0.00110981, 0.00183526]
Letter				
8	H2	0.107212	−0.000152493	[−0.00283779, 0.00268745]
8	H3	0.107212	0.00192856	[−0.00132479, 0.00544814]
8	OPTIMAL	0.107212	0.00192856	[−0.00135637, 0.00545245]
4	H2	0.151786	−0.00617889	[−0.0147332, 0.00242807]
4	H3	0.151786	−0.00249085	[−0.0117097, 0.00711604]
4	OPTIMAL	0.151786	−0.00249085	[−0.0120788, 0.00704749]
15	H2	0.0146734	0.000131949	[−0.00108937, 0.00140005]
15	H3	0.0146734	0.000690192	[−0.000733344, 0.00209202]
15	OPTIMAL	0.0146734	0.000690192	[−0.000686622, 0.00215188]

Table 14: Squared Frobenius sine distance / rank. All registered supplemental contrasts; policy-minus-myopic, stored unadjusted partition-cluster intervals.

Rank	Policy	Myopic mean	Difference	95% paired CI
Digits				
20	H2	0.903013	0.00118979	$[-0.000226563, 0.00256028]$
20	H3	0.903013	0.000240826	$[-0.00107511, 0.00159117]$
20	OPTIMAL	0.903013	0.000240826	$[-0.00108945, 0.00161697]$
10	H2	0.782468	0.00125287	$[-0.00201261, 0.00426899]$
10	H3	0.782468	0.00104931	$[-0.00191227, 0.00402243]$
10	OPTIMAL	0.782468	0.00104931	$[-0.00193535, 0.0041055]$
40	H2	0.933323	3.44037×10^{-5}	$[-0.000567661, 0.000587729]$
40	H3	0.933323	-8.02752×10^{-5}	$[-0.000742618, 0.000593463]$
40	OPTIMAL	0.933323	-8.02752×10^{-5}	$[-0.000725344, 0.000602064]$
UCI HAR				
20	H2	0.824701	4.18858×10^{-5}	$[-0.000753468, 0.000890789]$
20	H3	0.824701	-0.000326603	$[-0.00124653, 0.000633072]$
20	OPTIMAL	0.824701	-0.000326603	$[-0.00122688, 0.000616623]$
10	H2	0.783297	0.00184669	$[0.00043263, 0.00339231]$
10	H3	0.783297	0.000775153	$[-0.000835079, 0.00235571]$
10	OPTIMAL	0.783297	0.000775153	$[-0.000776783, 0.00235622]$
40	H2	0.875133	0.000391288	$[-0.000170194, 0.000931562]$
40	H3	0.875133	0.000415147	$[-0.000115067, 0.00105467]$
40	OPTIMAL	0.875133	0.000415147	$[-0.000102859, 0.00105564]$
Fashion				
20	H2	0.70662	0.000218281	$[-0.000239219, 0.000684281]$
20	H3	0.70662	0.000236562	$[-0.000175633, 0.000662816]$
20	OPTIMAL	0.70662	0.000236562	$[-0.000166262, 0.000672508]$
10	H2	0.646212	0.00114594	$[-0.0002155, 0.00258268]$
10	H3	0.646212	0.000928437	$[-0.000329086, 0.00216812]$
10	OPTIMAL	0.646212	0.000928437	$[-0.000305473, 0.00215895]$
40	H2	0.744662	2.26563×10^{-5}	$[-0.000232035, 0.000265004]$
40	H3	0.744662	-5×10^{-6}	$[-0.000336566, 0.000331254]$
40	OPTIMAL	0.744662	-5×10^{-6}	$[-0.000329535, 0.000332824]$
Letter				
8	H2	0.386514	-0.000869531	$[-0.00263322, 0.000860186]$
8	H3	0.386514	0.000755469	$[-0.000794922, 0.00229613]$
8	OPTIMAL	0.386514	0.000755469	$[-0.000789854, 0.00229962]$
4	H2	0.226615	0.000987891	$[-0.00078165, 0.0026848]$
4	H3	0.226615	0.000236328	$[-0.00133126, 0.00191682]$
4	OPTIMAL	0.226615	0.000236328	$[-0.00138289, 0.0018762]$
15	H2	0.537121	5.03906×10^{-5}	$[-0.000198437, 0.000305078]$
15	H3	0.537121	4.96094×10^{-5}	$[-0.000276562, 0.000374229]$
15	OPTIMAL	0.537121	4.96094×10^{-5}	$[-0.000278926, 0.000383203]$

Table 15: Ridge accuracy (fraction). All registered supplemental contrasts; policy-minus-myopic, stored unadjusted partition-cluster intervals.

Dataset	Policy	Myopic mean	Difference	95% paired CI
Digits	H2	343.741	−0.770454	[−0.825404, −0.713496]
Digits	H3	343.741	−0.782428	[−0.838548, −0.724233]
Digits	OPTIMAL	343.741	−0.782428	[−0.838724, −0.723593]
UCI HAR	H2	56.214	−1.35037	[−1.42845, −1.26919]
UCI HAR	H3	56.214	−1.36934	[−1.45042, −1.28623]
UCI HAR	OPTIMAL	56.214	−1.36934	[−1.45103, −1.28861]
Fashion	H2	76.2375	−1.93896	[−2.1084, −1.79487]
Fashion	H3	76.2375	−1.97232	[−2.13407, −1.82948]
Fashion	OPTIMAL	76.2375	−1.97232	[−2.13631, −1.82845]
Letter	H2	3.73169	−0.123209	[−0.132698, −0.113154]
Letter	H3	3.73169	−0.12517	[−0.135272, −0.115005]
Letter	OPTIMAL	3.73169	−0.12517	[−0.135075, −0.114767]

Table 16: Rank-independent exact complete-path Gram risk; all 12 registered supplemental contrasts.

K Supplementary Theory: Translation and an Explicit Failure Interval

Evidence status and scope. These proofs received a separate same-family AI-agent semantic review; independent implementation checks were run remotely with direct rational terminal-error recursion, without importing the original controller. All assertions passed. This is neither an independent human certification nor a cross-family semantic review. Common translation, covariance inequalities and uniform perturbation are standard tools; fixed-mass contribution differences are classical (Breidt & Chauvet, 2011). The added statements apply them to the protected correlated block and a declared Cube failure family. Neither proves a centered rank-one PSD feature-space strict-success neighborhood or real-data coverage. Supplementary validation/independent_results contains the translation record, three rational failure-point records and summary; point checks support implementation, while the interval follows from the uniform analytic proof.

K.1 A translation-centered protected-block certificate

Assume integral pathwise sample mass, preserved Bernoulli marginals, and $\sum_{i\in F} I_i \le 1$ for every admissible policy; put $R = [N] \setminus F$. Contributions lie in a finite-dimensional real Hilbert space. The registered actions, transitions, stopping, landing and deterministic tie rule are unchanged by common translation. For any A fixed before randomization let $Y_i = Z_i - A$, $q_i = p_i(1-p_i)$, β_{ij} be the maximum absolute Bernoulli covariance and w_{ij} its interval width. Define

$$
\begin{aligned}
P(A) &= \sum_{i\in F} q_i \left\|Y_i\right\|^2 - 2 \sum_{i<j\in F} p_i p_j \left\langle Y_i, Y_j \right\rangle, \\
T(A) &= \sum_{i\in R} q_i \left\|Y_i\right\|^2 + 2 \sum_{i<j\in R} \beta_{ij} \left|\left\langle Y_i, Y_j \right\rangle\right|, \\
D(A) &= 2 \sum_{i<j:\{i,j\}\not\subseteq F} w_{ij} \left|\left\langle Y_i, Y_j \right\rangle\right|, \\
L(A) &= (\sqrt{P(A)} - \sqrt{T(A)})_+^2, \quad U(A) = (\sqrt{P(A)} + \sqrt{T(A)})^2, \\
B(A) &= \min\{D(A), U(A) - L(A)\}.
\end{aligned}
$$

Proposition 11 (Translation-centered certificate). Terminal risk and every flight and landing cost are invariant under $Z_i \mapsto Z_i - A$; the myopic decisions are unchanged under the declared interface. For every policy, $L(A) \le J(\pi) \le U(A)$, and $J(\mu) - J^* \le B(A)$. If $L(A) > 0$, then $J(\mu)/J^* \le 1 + B(A)/L(A)$. If $Z_i = A$ for all $i \in R$, every policy has risk $P(A)$.

Proof. Set $e_i = I_i - p_i$. Exact mass gives $\sum_i e_i = 0$, so terminal error is unchanged. A flight increment is $\lambda \sum_i u_i Z_i$ with $\lambda \in \{\alpha, -\beta\}$ and $\sum_i u_i = 0$; landing changes likewise have zero total coordinate sum. Immediate minimizers and ties therefore coincide. With $X_F = \sum_{i\in F} e_i Y_i$ and $X_R = \sum_{i\in R} e_i Y_i$, protected exclusion gives $\mathbb{E}\left\|X_F\right\|^2 = P(A)$ and the Bernoulli absolute covariance bounds give $\mathbb{E}\left\|X_R\right\|^2 \le T(A)$. Triangle and reverse triangle inequalities in L^2 give the energy interval, including the necessary positive part. Two policies have identical diagonal and protected-block covariance entries; all other differences have magnitude at most w_{ij}, so their risk spread is at most $D(A)$ and the energy interval width. Taking the comparator infimum proves the additive bound. Its denominator is positive when $L(A) > 0$. If all residual Y_i vanish, risk equals $P(A)$; when this is zero no ratio is defined. □

Finite translation choice. For a nonempty predeclared finite list $\mathcal{C}$ containing zero, compute $L_{\mathcal{C}} = \max_A L(A)$, $U_{\mathcal{C}} = \min_A U(A)$ and $B_{\mathcal{C}} = \min\{\min_A D(A), U_{\mathcal{C}} - L_{\mathcal{C}}\}$. Intersecting valid intervals proves the corresponding bounds. Inclusion of zero ensures the additive bound is no worse; the ratio is no worse whenever the original lower denominator was positive. Each candidate costs $O(N^2 d)$ for d scalar contribution entries. A weighted residual mean may minimize the diagonal residual surrogate, but is not an optimizer of this full certificate. If no positive lower denominator is certified, use only the additive bound.

Explicit strict tightening (independent implementation checked). Use $p = (1/2, 1/2, 1/4, 1/4, 1/4, 1/4)$ and auxiliary scores $a = (3, 1, -3, 2, -2, 0)$ from the registered six-unit geometry. Mass is two and balance 5/4; residual feasibility gives $F = \{3, 5\}$. For $Z = (1, 1, 2, 1, 1, 1)$ and $A = 1$, only $Y_3 = 1$ is nonzero, yielding $P(1) = 3/16$, $T(1) = D(1) = 0$ and zero policy regret at positive risk. Without translation $P(0) = 11/16$, $T(0) \geq 7/8$, $L(0) = 0$, and $D(0) > 0$; the original additive certificate is strictly positive. The independent solver obtained $T(0) = 11/4$, $D(0) = 19/2$, $U(0) = 99/16$ and minimum and maximum policy risk both 3/16, and checked cost/myopic invariance at all 1,217 reachable states. Thus the bound is strictly improved. Original residual contributions are nonzero; this scalar PSD example does not assert centered-feature Gram realizability.

K.2 A conservative analytic positive-gap interval

Keep the original full-triple six-unit sampling geometry and pairwise landing; change only Z_6 in $Z(t) = (2, -4, 12, -16, 20, -24 + t)$. At the base point the imported certificate has unique myopic root $a = (2, 4, 5)$, immediate cost 3/2, next root cost 32/5, and optimal forced-root gap $g = Q_0^*(a) - V_0^* = 369406/19173 > 19$.

Proposition 12 (Uniform failure interval, conditional base imports). For every real $|t| \leq 1/40$, a remains the unique myopic root and

$$J_t(\mu_t) - V_t^* \geq Q_t^*(a) - V_t^* \geq g - 312|t| - 2t^2 > 8959/800 > 0.$$

Proof. For any binary endpoint $|e_i| \leq 1$ and the base absolute contributions sum to 78, so $|X_0| \leq 78$. As $X_t = X_0 + te_6$, $|X_t^2 - X_0^2| \leq 156|t| + t^2 =: E(t)$ uniformly over all paths and policies. The same bound holds for root immediate cost since every coordinate increment lies in $[-1, 1]$. Root a excludes six, so its cost is constant; every competitor has excess at least $49/10 - E(t) \geq 1599/1600 > 0$. The feasible policy family is independent of t. Uniform error survives infima both over all policies and those forced to take a, giving forced-root gap at least $g - 2E(t) \geq g - 6241/800 > 8959/800$. Every myopic continuation belongs to this forced-root family. □

The interval is a conservative robustness region, not the maximal phase boundary. It certifies neither unchanged globally optimal root nor unchanged branch-optimal actions throughout the interval. Root-local cost is insensitive to contribution six while terminal sixth-unit exposure persists; the imported base continuation penalty survives. The standard uniform perturbation proof is distinguished from any stronger analytic decomposition of covariance, branch and landing effects. Independent rational checks at $-1/40, 0, 1/40$ reproduced the root margins, forced-root positive gaps and terminal identities. These verify base imports and implementation, not an interval by interpolation.

K.3 Fixed-base-policy exposure and landing diagnostics

Freeze every myopic and optimal decision at $t = 0$, naming the resulting policies μ_0 and w_0. Their terminal laws are unchanged when only Z_6 changes. Direct exact arithmetic on the independently enumerated endpoint distributions gives

$$\begin{aligned} J_{\mu_0}(t) &= \frac{1113001}{4410} - \frac{195479}{17640}t + \frac{3}{16}t^2, \\ J_{w_0}(t) &= \frac{298063}{1743} - \frac{41}{6}t + \frac{3}{16}t^2, \\ J_{\mu_0}(t) - J_{w_0}(t) &= \frac{29785853}{366030} - \frac{74939}{17640}t. \end{aligned}$$

The equal quadratic coefficient is fixed by $p_6(1 - p_6) = 3/16$. For $\Delta\Sigma = \Sigma_{\mu_0} - \Sigma_{w_0}$, its sixth column is

$$(\Delta\Sigma_{16}, \ldots, \Delta\Sigma_{66}) = \left(-\frac{1717}{23520}, \frac{3629}{35280}, -\frac{7}{360}, \frac{1}{32}, -\frac{1}{24}, 0 \right).$$

Thus the gap slope is exactly $2\sum_{i<6} Z_i(0)\Delta\Sigma_{i6} = -74939/17640$: its parameter sensitivity comes from off-diagonal sixth-unit exposure, while the common marginal variance cancels.

This identifies a terminal-correlation mechanism for two fixed base policies. It does not assert that the actual μ_t or optimal policy follows either frozen continuation throughout the interval; the earlier forced-root bound, not this fixed witness identity, certifies the whole interval.

A separate independent audit retaining the original transitions but assigning zero objective weight to landing risk yields the unique optimal root $(1, 4, 6)$ and value $188627/1920$. This is an objective-term ablation, not a legal alternative landing-kernel experiment. Exact recomputation from the supplied terminal probabilities verified probability mass, marginals, coefficients and the covariance column. Sources are the supplementary validation/mechanism_results/mechanism_attempt01/fixed_policy_mechanism.json and repairs/FIXED_MECHANISM_ARITHMETIC_CHECK.json.

L Detailed Scope and Limitations

All results are relative to an explicitly registered action and landing class. The energy theorem extends to a larger action class only after fixed marginals and terminal mutual exclusion are proved there; the strict and normal-cone certificates must be recomputed. A fixed-marginal terminal-design outer relaxation strictly improves on the explicit circuit point, but it is unknown whether a larger Cube process can reach its optimizer.

The full-dimensional result guarantees near-optimality, not exact optimality throughout a full-dimensional region. The explicit strict radius is conservative and currently covers one local region. The six-unit failure instance is deliberately constructed, although it is non-half-integral, tie-free, branch-dependent, and landing-sensitive. Finite sentinel-free replications cannot replace an all-size proof. Moreover, under a fixed natural mapping, sound sufficient certificates can be too rare or too loose to save computation in practice. Priority language remains subject to the final literature review.

The strict-open certificate above belongs to general matrix-contribution geometry: its $G_{13} = -6$ cannot occur in the rank-one PSD specialization, where $G_{ij} = (x_i^T x_j)^2/(n^2 p_i p_j) \geq 0$. Even repairing this symmetric entry requires Frobenius distance at least $6\sqrt{2}$, far beyond the certified radius. Thus that ball is not a feature-space success certificate.

The Gram/PCA bridge has a separate limitation. Lower expected Gram error tightens the stated PCA guarantees, but it need not improve every randomized realization, every finite-eigengap problem, or every downstream classifier. The frozen study confirms this distinction: exact additive Gram risk improves on all four datasets, but most sampled PCA and classification intervals cross zero. Moreover, H2/H3 require roughly one to two orders of magnitude more offline complete-policy evaluation work than myopic; no cold-cache online latency or policy-specific peak-memory measurement was recorded. We therefore interpret them as mechanism probes, not as a scalable controller.

The added runtime study is confined to 24 exposed development cases on one host, with five timing rounds and no independent-data or cross-hardware confirmation. IGC is faster than the cached implementation of its own rule, but slower than public Cube. Its surrogate objective omits the direction-dependent maximal-step factor. We have not evaluated a combined IGC-plus-lookahead controller, and cannot attach the earlier planning risk gains to the new runtime savings. The additional computational contribution does not by itself demonstrate a major sampling-policy innovation.

M Proof of the PCA Consequences

For (4), optimality of $\widehat{P}_k$ for $\widehat{\Gamma}$ and Cauchy–Schwarz give

$$\mathcal{L}_\Gamma(\widehat{P}_k) - \mathcal{L}_\Gamma(P_k) \leq \|P_k - \widehat{P}_k\|_F \, \|\widehat{\Gamma} - \Gamma\|_F \leq \sqrt{2k}\, \|\widehat{\Gamma} - \Gamma\|_F.$$

Jensen's inequality completes the expectation bound. Equation (5) specializes the population-gap Davis–Kahan variant of Yu et al. (2015), whose constant is two, and then uses the Frobenius branch of its numerator.