Diffusion budgets for contact-freezing lattice ellipsoids

Scott Schubr
Luminary Oddities

October 2026

Abstract

We study bounded-covariance matrix martingales whose ellipsoids acquire and retain lattice contacts. On a common set of Haar lattices of probability tending to one, we bound the total nonnegative gain in a heat-occupation frame by the time-averaged square-root diffusion cost. A cost of $o(\log d)$ forces this gain to vanish at the horizon $(16 \log d + 8 \log\log d + O(1))/d^2$. The bound is uniform over admissible policies on each selected lattice and also gives an averaged-budget theorem. Its proof combines deterministic pair envelopes, radial and angular Haar estimates, same-noise coupling and a potential-weighted restart. Applications include every fixed-strength normalized curvature-preserving tilt and rotating low-rank feedback. We also solve the instantaneous curvature-preserving optimization problem with the same diffusion cost, including noncommuting curvature and occupation operators. For fixed instantaneous operators, matching two-sided bounds give the sharp square-root small-budget law with an explicit linear error.

Keywords. stochastic control; matrix martingales; random lattices; diffusion cost; occupation measures.
2020 Mathematics Subject Classification. 93E20, 60H10, 11H06.

1 Introduction

Contact-freezing matrix martingales provide a way to grow lattice-free ellipsoids while accumulating boundary constraints. Klartag constructed this process and used one-vector Gaussian contact estimates to obtain lattice packing bounds [1]. Here we study a quantitative question about perturbing its diffusion: how much reduction of a specified occupation-frame trace can be purchased with a small square-root covariance budget?

The answer concerns the full-information policy class: a controller may use the complete initial lattice and its evolving contacts. The initial ensemble is Haar measure on lattices of covolume equal to the volume of the unit ball. The spectral window, horizon and gain functional are fixed explicitly below. Deterministic lattice envelopes produce one high-probability set on which the principal estimate holds simultaneously for every admissible policy.

Reading the main result. Fix the covariance ceiling, spectral-window constant and horizon parameter before taking the dimension to infinity. Theorem 5.1 selects one set of lattices of Haar probability $1 - o(1)$. On each selected lattice, every admissible full-information policy with nonnegative instantaneous gain obeys the same gain–budget inequality. Its budget is the expected time average of $\|\sqrt{C_t} - I_{F_t}\|_{\mathrm{HS}}^2$; its gain is the expected time integral of the occupation trace defined in Section 2. A uniformly $o(\log d)$ budget gives uniformly vanishing gain. Admissibility includes existence of the stopped process. Corollaries 5.2 and 6.2 treat averaged budgets and moving low-rank perturbations.

There are two complementary ingredients. The dynamical theorem bounds accumulated gain using a heat potential whose frozen boundary values record contact time. An instantaneous matrix optimization then constructs curvature-preserving feedback penalized by exactly the same square-root cost. This pairing makes the budget operational: repeatedly moving a low-rank active subspace still has bounded cost, whereas a general optimized controller must be evaluated by its actual accumulated cost.

The background mean-value inputs are the classical Siegel and Rogers formulas; we state precisely the forms used, as recorded in [2, 3]. The scalar martingale comparison extends the contact-tail method of [1, Proposition 4.1] to two fixed vectors before Haar integration. The cost is the squared Bures-Wasserstein distance to the identity [4]. Recent packing constructions also use invariant deformation spaces and orbit-weighted contacts [5]; the ensemble and occupation trace in the present theorems are specified independently of those constructions.

2 Finite-shell processes and the occupation potential

Equip $V = \mathrm{Sym}_d(\mathbb{R})$ with $\langle H, J\rangle = \mathrm{Tr}(HJ)$ and its Frobenius norm. Operators on V have operator norm and Hilbert–Schmidt norm, denoted $\|\cdot\|_{\mathrm{op}}$ and $\|\cdot\|_{\mathrm{HS}}$. These operator norms are distinct from norms of matrices in V.

Fix a lattice L, $a_0 > 0$, $T > 0$ and $0 < w < a_0$. Suppose every nonzero lattice vector has length greater than $r_0 = a_0^{-1/2}$. Put

$$r_1 = (a_0 - w)^{-1/2}, \qquad S = (L \cap \{r_0 < |x| \le r_1\})/\{\pm 1\}.$$

Choose one representative of each class. This is a finite set. For $x \in S$ write $u_x = x/|x|$, $e_x = u_x u_x^T$, and

$$h_x = a_0 - |x|^{-2}, \qquad z_x(t) = \langle A_t, e_x\rangle - |x|^{-2}.$$

Start at $A_0 = a_0 I_d$. At the first time $z_x = 0$, declare the class hit. At all subsequent times require increments to lie in

$$F_t = \{H \in V : \langle H, e_x\rangle = 0 \text{ for every class hit by time } t\}.$$

Let P_t be its orthogonal projection. Simultaneous hits are all frozen. Projection values at isolated hitting times can be assigned predictable versions without changing stochastic integrals.

Definition 2.1. An admissible policy is a predictable positive semidefinite covariance C_t on F_t, extended by zero on $F_t^{\perp}$, together with a continuous solution

$$dA_t = C_t^{1/2}\, dW_t, \qquad t < \tau = T \wedge \sigma, \quad \sigma = \inf\{t : \|A_t - a_0 I_d\|_{\mathrm{op}} \ge w\}, \tag{1}$$

where W is standard Brownian motion in V. The process is stopped at τ. We require $0 \le C_t \le c I_{F_t}$ for a fixed $c \ge 1$ and preservation of all unhit inequalities up to their first hits.

Before σ, vectors longer than r_1 have $x^T A_t x > 1$, so S includes every possible new contact. Freezing ensures $x^T A_t x = 1$ after a hit. Thus $\{y : y^T A_t y < 1\}$ has no nonzero lattice points before the stop. Admissibility includes existence of the stopped solution for the chosen feedback.

For $R > 0$, $z \ge 0$, define

$$\begin{aligned} g_R(z) &= \int_0^R 2\overline{\Phi}(z/\sqrt{s})\, ds \\ &= (R + z^2)2\overline{\Phi}(z/\sqrt{R}) - 2z\sqrt{R}\,\phi(z/\sqrt{R}), \\ w_R(z) &= 4\overline{\Phi}(z/\sqrt{R}), \end{aligned} \tag{2}$$

where ϕ and $\overline{\Phi}$ are the standard normal density and upper tail. For unhit classes at $t < T$ set $w_x(t) = w_{T-t}(z_x(t))$; use zero weights after contact or stopping. Define operators on V by

$$\begin{aligned} K_t^{\mathrm{raw}} &= \sum_{x \text{ unhit}} w_x(t) e_x \otimes e_x, \quad K_t = P_t K_t^{\mathrm{raw}} P_t, \\ L_t &= \operatorname{Tr} K_t^{\mathrm{raw}} - \operatorname{Tr} K_t \geq 0, \quad a_t = \operatorname{Tr}((I_{F_t} - C_t) K_t). \end{aligned} \tag{3}$$

Here $(e \otimes e)H = \langle e, H \rangle e$. All subsequent dynamical gain conclusions assume $a_t \geq 0$ almost everywhere. Set

$$\mathcal{A}(L) = \mathbb{E}_W \int_0^\tau a_t \, dt, \qquad B(L) = T^{-1} \mathbb{E}_W \int_0^\tau \|C_t^{1/2} - I_{F_t}\|_{\mathrm{HS}}^2 \, dt. \tag{4}$$

The subscript W holds the lattice fixed.

Lemma 2.2 (Heat derivatives and potential identity)**.** *Let $\Gamma_0 = \sum_{x \in S} g_T(h_x)$. Freeze each contribution $g_{T-t}(z_x(t))$ at its value $T - \tau_x$ when class x first hits at τ_x. Stop all contributions at τ. Their sum Ψ_t satisfies*

$$\Psi_t = \Gamma_0 + M_t - \frac{1}{2} \int_0^{t \wedge \tau} (L_s + a_s) \, ds, \tag{5}$$

where M is a square-integrable martingale for each fixed lattice. In particular $\Psi \geq 0$. If N_t counts new classes hit by $t \leq \tau$, then

$$\mathbb{E}_W \int_0^\tau N_s \, ds + \tfrac{1}{2} \mathcal{A}(L) \leq \Gamma_0. \tag{6}$$

Moreover, for $R > 0$, $z \geq 0$,

$$\begin{aligned} g_{zz} = w_R, \quad \partial_R g_R = w_R/2, \\ g_R'(z) = 4[z\overline{\Phi}(z/\sqrt{R}) - \sqrt{R}\phi(z/\sqrt{R})], \\ |g_R'(z)| \leq \sqrt{2R/\pi}\, w_R(z). \end{aligned} \tag{7}$$

Proof. The explicit expression in (2) extends smoothly across $z = 0$ for $R > 0$ and gives $g_R(0) = R$, $g_R'(0) = -\sqrt{8R/\pi}$, $g_R''(0) = 2$. Differentiation gives the first three identities. To prove the last one, let $m(r) = \phi(r)/\overline{\Phi}(r)$. The conditional variance of a standard normal given it is at least r equals $1 + rm(r) - m(r)^2 > 0$. Hence $(m(r) - r)' = m(r)(m(r) - r) - 1 < 0$. Also $m(r) > r$ for $r \geq 0$ by integrating $s\phi(s)$ over $[r, \infty)$, and $m(0) = \sqrt{2/\pi}$. Substituting $r = z/\sqrt{R}$ proves the bound.

Itô's formula on each unhit contribution gives drift $-w_x/2 + (w_x/2)\langle C_t P_t e_x, P_t e_x \rangle$. Sum and use (3) to obtain (5). At contact $g_R(0) = R$, so freezing introduces no jump. For $q \leq \tau$ the frozen part is

$$\sum_{\tau_x \leq q} (T - \tau_x) = \int_0^q N_s \, ds + (T - q) N_q.$$

The unfrozen part is nonnegative. Take expectations and let $q \uparrow \tau$ to get (6). All passages can first be made with $q \leq T - \eta$. The bounds $0 \leq g_R \leq R$, $|g_R'| \leq \sqrt{8R/\pi}$, $0 \leq g_R'' \leq 2$ and $|\partial_R g_R| \leq 1$, together with finite S and bounded C, imply dominated convergence and square integrability on $[0, T]$. Terminal contacts contribute zero frozen cost. □

3 Haar normalization and initial occupation

Let $\kappa_d = \mathrm{vol}(B_1^d)$, and let $\mathbb{H}_d$ denote invariant probability on lattices of covolume κ_d. In the remainder, d tends to infinity, $\gamma \in \mathbb{R}$ and $b_0 > 0$ are fixed, and

$$\rho = \log\log d, \quad a_0 = (1-\rho/d)^{-2}, \quad T = \frac{16\log d + 8\rho + \gamma}{d^2}, \quad w = b_0\sqrt{dT}. \tag{8}$$

All assertions are for sufficiently large d, so $T > 0$ and $w < a_0$. Let $\mathcal{H}_d$ be the initial hole event $L \cap \overline{B}_{r_0} = \{0\}$, ignoring its null boundary.

We use three precise classical mean-value consequences. For nonnegative integrable f and $d \geq 3$,

$$\mathbb{E}_{\mathbb{H}_d} \sum_{0 \neq x \in L} f(x) = \kappa_d^{-1} \int f(x)\,dx, \tag{9}$$

$$\mathbb{E}_{\mathbb{H}_d} \sum_{x,y \in L \text{ independent}} F(x,y) = \kappa_d^{-2} \iint F(x,y)\,dx\,dy. \tag{10}$$

The second sum is ordered. Formula (10) is Rogers' independent-vector formula, stated in [2, equation (7)]. For even compactly supported f, putting $Y_f = \frac{1}{2}\sum_{0\neq x\in L} f(x)$ gives the variance bound

$$\mathrm{Var}\, Y_f \leq \frac{\zeta(d/2)^2}{2\kappa_d\zeta(d)} \int |f|^2. \tag{11}$$

Indeed write the full sum as a primitive sum of $F(x) = \sum_{k\geq 1} f(kx)$. The primitive second moment in [3, Introduction, Rogers formula] gives variance $(2\kappa_d\zeta(d))^{-1}\int F^2$ for the half-sum. Minkowski's inequality gives $\|F\|_2 \leq \zeta(d/2)\|f\|_2$. Approximation handles all the nonnegative kernels used below.

The normalized radial intensity of all vectors in the shell is, for $h = a_0 - r^{-2}$,

$$\mu_d(dh) = \frac{d}{2}(a_0-h)^{-(d+2)/2}dh \leq Cde^{-\rho+\nu h}dh, \qquad \nu = \frac{d+2}{2(a_0-w)}, \quad 0 \leq h \leq w. \tag{12}$$

This follows from the radial integral $d\int r^{d-1}dr$ and the mean-value theorem applied to $-\log(a_0-h)$. Constants in the paper may depend on b_0, γ, c, but not on d, the lattice or the policy. Direct expansion gives

$$\nu^2 T = 4\log d + 2\rho + \gamma/4 + o(1). \tag{13}$$

The error uses $w\log d = o(1)$.

Lemma 3.1. *There is a fixed $G_0 < \infty$ such that*

$$\mathbb{H}_d(\mathcal{H}_d \cap \{\Gamma_0 \leq G_0\}) = 1 - o(1).$$

Here Γ_0 is defined by the deterministic shell formula even on lattices outside $\mathcal{H}_d$.

Proof. By (9), the expected number of nonzero vectors in B_{r_0} is $r_0^d = (1-\rho/d)^d = O(e^{-\rho}) = o(1)$. For the mean of Γ_0, (12), Tonelli and integration by parts give

$$\int_0^\infty e^{\nu h}\overline{\Phi}(h/\sqrt{s})\,dh = \frac{e^{\nu^2 s/2}\Phi(\nu\sqrt{s}) - 1/2}{\nu} \leq \nu^{-1}e^{\nu^2 s/2}.$$

Thus $\mathbb{E}\Gamma_0 \leq Cde^{-\rho}\nu^{-3}e^{\nu^2 T/2} = O(1)$ by (13). Also $g_T(h) \leq Tw_T(h)/2$ and $w_T(h)^2 \leq 4e^{-h^2/(2T)}$. Completing the square shows

$$\int_0^w w_T(h)^2\mu_d(dh) \leq Cde^{-\rho}\sqrt{T}\,e^{\nu^2 T/2} \leq Cd^2(\log d)^{1/2}.$$

The even kernel $f(x) = \mathbf{1}_{\{r_0<|x|\leq r_1\}}g_T(a_0-|x|^{-2})$ in (11) yields $\mathrm{Var}\,\Gamma_0 \leq CT^2d^2(\log d)^{1/2} = o(1)$. Choose G_0 greater than a uniform bound for $\mathbb{E}\Gamma_0$ by one and apply Chebyshev's inequality. □

4 Same-noise coupling and deterministic pair envelopes

For any admissible policy define on the same probability space

$$Z_t = a_0 I_d + \int_0^{t\wedge\tau} P_s\, dW_s, \qquad E_t = A_t - Z_t.$$

The tangent spaces here are those of the actual process A. Fix $\kappa = 1/2$, let

$$b = \kappa\frac{\log d}{d}, \quad \eta = \inf\{t : \|E_t\|_{\mathrm{F}} \geq b\}, \quad q = \eta \wedge \tau, \quad p(L) = \mathbb{P}_W(\eta < \tau). \tag{14}$$

The Frobenius norm controls every quadratic form of E. Doob's inequality and the stochastic isometry imply

$$\mathbb{E}_W \sup_{t\leq T} \|E_t\|_{\mathrm{F}}^2 \leq 4B(L)T, \qquad p(L) \leq \min\{1, 4B(L)T/b^2\} \leq \min\{1, C_* B(L)/\log d\}. \tag{15}$$

This does not require independence of the feedback and the lattice.

Lemma 4.1 (Fixed-vector and fixed-pair bounds)**.** *For two fixed shell vectors, write $h = h_x$, $k = h_y$, $\theta = u_x \cdot u_y$. For every deterministic $t < T$,*

$$\mathbb{E}_W[\mathbf{1}_{\{t<q\}} w_x(t) w_y(t)] \leq 4\exp\left(-\frac{[(h+k-2b)_+]^2}{4T(1+\theta^2)}\right), \tag{16}$$

$$\mathbb{E}_W[\mathbf{1}_{\{t<q\}} w_x(t)^2] \leq 4\exp\left(-\frac{[(h-b)_+]^2}{2T}\right). \tag{17}$$

Proof. For $z \geq 0$ and $\alpha \geq 0$, the Gaussian tail bound gives

$$4\overline{\Phi}(z/\sqrt{R}) \leq 2e^{-\alpha z + \alpha^2 R/2}.$$

The sum of the reference gaps $\langle Z, e_x + e_y\rangle - |x|^{-2} - |y|^{-2}$ is a martingale with initial value $h + k$ and quadratic variation at most

$$\|P_s(e_x + e_y)\|_{\mathrm{F}}^2 ds \leq 2(1+\theta^2)ds.$$

On $\{t < q\}$ both actual gaps differ from their reference gaps by at most b. Apply the exponential bound to the actual nonnegative gaps on the both-unhit event, then replace them by the reference gaps minus b. Only at this point drop the event. The exponential martingale estimate for the reference stopped at q gives a bound

$$4\exp\{-\alpha(h+k-2b) + \alpha^2[(T-t) + (1+\theta^2)t]\}.$$

Its quadratic coefficient is at most $T(1+\theta^2)$; optimization over $\alpha \geq 0$ proves (16). For (17), use $w_x^2 \leq 2w_x$, reference quadratic variation at most t by time t, and optimize $4\exp\{-\alpha(h-b) + \alpha^2 T/2\}$. The reference gaps after the event is dropped may be negative; the exponential argument remains valid. □

Define deterministic envelopes on each lattice, using full signed vectors in the shell and ordered independent pairs:

$$\begin{aligned} Q_j(L) = {} & 4\sum_x \exp\left(-\frac{[(h_x - b)_+]^2}{2T}\right) \\ & + 4 \sum_{x,y \text{ independent}} (u_x \cdot u_y)^{2j} \exp\left(-\frac{[(h_x + h_y - 2b)_+]^2}{4T(1+(u_x\cdot u_y)^2)}\right), \quad j = 1, 2. \end{aligned} \tag{18}$$

The constants deliberately absorb the antipodal factors. On $\mathcal{H}_d$, $r_1 < 2r_0$ for large d. Every shell vector is primitive: otherwise its primitive divisor has length at most $r_1/2 < r_0$. Distinct antipodal classes therefore are linearly independent. Expanding the squared raw-frame norm and using Lemma 4.1 gives

$$\mathbb{E}_W[\mathbf{1}_{\{t<q\}}\|K_t\|_{\mathrm{HS}}^2] \le Q_2(L). \tag{19}$$

Compression is used on the whole operator; no pairwise domination after projection is needed.

Lemma 4.2 (Radial and angular estimates)**.** *There is a fixed integer m such that*

$$\int_{\mathcal{H}_d} Q_2 \, d\mathbb{H}_d \le Cd^{2+\kappa}(\log d)^m, \qquad \int_{\mathcal{H}_d} Q_1 \, d\mathbb{H}_d \le Cd^{3+\kappa}(\log d)^m. \tag{20}$$

Proof. Use (9) for the diagonal and (10) for the independent pairs, dropping the nonnegative hole indicator after applying the deterministic bounds. The diagonal is at most

$$Cde^{-\rho}\int_0^\infty e^{\nu h-[(h-b)_+]^2/(2T)}dh \le Cde^{-\rho+\nu b}(b+\sqrt{T})e^{\nu^2T/2}.$$

Indeed split at b, translate the rest by b, and complete the square. Since $\nu b = (\kappa/2)\log d + o(1)$, this costs at most $d^{2+\kappa/2}$ times a fixed power of $\log d$.

For the pair integral, put $V = T(1+\theta^2)$ and combine the two radial variables as $s = h+k$. Formula (12) bounds it by

$$\begin{aligned} &Cd^2e^{-2\rho}\int_0^\infty se^{\nu s-[(s-2b)_+]^2/(4V)}ds \\ &\qquad \le Cd^2e^{-2\rho+2\nu b}(V+\nu V^{3/2}+b\sqrt{V}+b^2)e^{\nu^2V}. \end{aligned} \tag{21}$$

On $s \le 2b$ bound the integral by $2b^2e^{2\nu b}$. On the rest set $s = v+2b$ and complete the square

$$\nu v - v^2/(4V) = \nu^2V - (v-2\nu V)^2/(4V).$$

The integrals of $v+2b$ are bounded by a constant times $V+\nu V^{3/2}+b\sqrt{V}$. This proves (21) including the cutoff region.

For independent uniform sphere directions the density of θ is $c_d(1-\theta^2)^{(d-3)/2}$ on $[-1,1]$, with $c_d \le C\sqrt{d}$. Since $\nu^2T = O(\log d)$, for $j = 1, 2$ Gaussian integration gives

$$\mathbb{E}_\theta[\theta^{2j}e^{\nu^2T\theta^2}] \le C\sqrt{d}\int_{\mathbb{R}} t^{2j}e^{-dt^2/4}dt \le Cd^{-j}. \tag{22}$$

For large d the exponent inequality follows from $\log(1-t^2) \le -t^2$. Also $2\nu b = \kappa\log d + o(1)$ and $e^{\nu^2T} = O(d^4e^{2\rho})$. Uniformly for $T \le V \le 2T$, the product of d^2 and the parenthesis in (21) is bounded by a fixed power of $\log d$. Thus the pair contribution is at most $Cd^{4+\kappa-j}(\log d)^m$, proving (20). Every angular range and both radial tails have been included. □

5 The diffusion-budget theorem

Theorem 5.1. *Fix $b_0 > 0$, $\gamma \in \mathbb{R}$ and $c \ge 1$. With (8), there are measurable sets $\mathcal{G}_d \subseteq \mathcal{H}_d$ of Haar probability $1-o(1)$ and constants $G_0, C_* < \infty$ such that the following holds simultaneously for every admissible policy on each $L \in \mathcal{G}_d$ with $a_t \ge 0$:*

$$\begin{aligned} 0 \le \mathcal{A}(L) \le (1+\sqrt{c})\sqrt{B(L)}\,d^{-1/4} + 2G_0p_B(L) + 2d^{-1/8}\sqrt{p_B(L)}, \\ p_B(L) = \min\{1, C_*B(L)/\log d\}. \end{aligned} \tag{23}$$

Consequently every uniformly $o(\log d)$ per-lattice budget has uniformly vanishing gain on this common set.

Proof. Use Lemmas 3.1 and 4.2, with $\kappa = 1/2$, to choose the common set by the deterministic conditions

$$\Gamma_0 \le G_0, \qquad T\sqrt{Q_2(L)} \le d^{-1/4}, \qquad C_c T^2 Q_1(L) \le d^{-1/4}, \tag{24}$$

where C_c will cover the bracket constant below. Markov's inequality and $T = O(\log d/d^2)$ show that the second condition excludes at most $Cd^{-1}(\log d)^{m+2}$ Haar measure, and the third excludes at most $Cd^{-1/4}(\log d)^{m+2}$. Both tend to zero. There is no policy in the definition of $\mathcal{G}_d$.

By functional calculus,

$$\|I_F - C\|_{\mathrm{HS}} \le (1+\sqrt{c})\|I_F - \sqrt{C}\|_{\mathrm{HS}}.$$

Cauchy–Schwarz on time and Brownian probability and (19) imply

$$\mathbb{E}_W \int_0^q a_s\, ds \le (1+\sqrt{c})\sqrt{B(L)T}\sqrt{TQ_2(L)} \le (1+\sqrt{c})\sqrt{B(L)}d^{-1/4}. \tag{25}$$

The martingale in Lemma 2.2 has integrand $C_t^{1/2} P_t U_t$, where

$$U_t = \sum_{x \text{ unhit}} g'_{T-t}(z_x(t))e_x.$$

All scalar coefficients have the same sign. Since $\langle e_x, e_y\rangle = \theta^2 \ge 0$, (7), compression contraction and Lemma 4.1 give

$$V(L) := \mathbb{E}_W\langle M\rangle_q \le C_c T^2 Q_1(L) \le d^{-1/4}. \tag{26}$$

In detail, $\|U_t\|_{\mathrm{F}}^2$ expands into pairs bounded by $(2T/\pi)w_x w_y \theta^2$; multiply by c, integrate over $t < q$, and use the full-vector factors already absorbed in Q_1.

Let $F_{\mathrm{exit}} = \{\eta < \tau\}$, $p = \mathbb{P}_W(F_{\mathrm{exit}})$, and let Γ_η be the sum of unhit heat contributions at η on this event. The potential identity gives $\Psi_q \le \Gamma_0 + M_q$ and $\Gamma_\eta \le \Psi_\eta$. Since $\mathbb{E}_W M_q = 0$,

$$\mathbb{E}_W[\Gamma_\eta \mathbf{1}_{F_{\mathrm{exit}}}] \le \Gamma_0 p + \mathbb{E}_W[M_q(\mathbf{1}_{F_{\mathrm{exit}}} - p)] \le G_0 p + \sqrt{V(L)p(1-p)}. \tag{27}$$

Restart (5) conditionally at η, subtracting old frozen contact constants. Its terminal potential is nonnegative, and $L_s, a_s \ge 0$. Optional sampling therefore gives

$$\mathbb{E}_W\left[\mathbf{1}_{F_{\mathrm{exit}}} \int_\eta^\tau a_s ds\right] \le 2\mathbb{E}_W[\Gamma_\eta \mathbf{1}_{F_{\mathrm{exit}}}].$$

This restart uses the same remaining horizon $T - s$ and the same controlled path. Combine this inequality with (25), (27), (26) and (15); monotonicity in p after replacing $\sqrt{p(1-p)}$ by $\sqrt{p}$ proves (23). □

Corollary 5.2 (Averaged budgets)**.** *For a measurable choice of admissible policy on $\mathcal{G}_d$, suppose $\mathbb{E}_{\mathbb{H}_d}[B \mid \mathcal{G}_d] \le b_d$. Then*

$$\begin{aligned}\mathbb{E}_{\mathbb{H}_d}[\mathcal{A} \mid \mathcal{G}_d] \le{}& (1+\sqrt{c})\sqrt{b_d}d^{-1/4} + 2G_0 \min\{1, C_* b_d/\log d\} \\ &+ 2d^{-1/8}\sqrt{\min\{1, C_* b_d/\log d\}}.\end{aligned} \tag{28}$$

If $b_d = o(\log d)$, both this mean and the conditional probability of $\mathcal{A} > g$ tend to zero for each fixed $g > 0$.

Proof. Apply Jensen's inequality to the square roots and to the concave minimum function in (28). The last assertion follows by Markov's inequality because $\mathcal{A} \ge 0$. □

6 Curvature-preserving feedback and rotating modes

For $A > 0$ define the positive definite curvature operator on F by

$$G_{A,F}H = P_F(A^{-1}HA^{-1}), \qquad H \in F. \tag{29}$$

Its quadratic form is $\operatorname{Tr}(A^{-1}HA^{-1}H) > 0$ for nonzero H. The ellipsoid $\mathcal{E}_A = \{y : y^T Ay \le 1\}$ has volume $\kappa_d(\det A)^{-1/2}$. Matrix Itô differentiation gives

$$d\log \operatorname{vol}(\mathcal{E}_{A_t}) = -\tfrac{1}{2}\operatorname{Tr}(A_t^{-1}dA_t) + \tfrac{1}{4}\operatorname{Tr}(C_t G_{A_t,F_t})dt. \tag{30}$$

Thus the affine condition $\operatorname{Tr}(C_tG_t) = \operatorname{Tr} G_t$ preserves this local drift.

Proposition 6.1 (Normalized feedback)**.** *Fix $0 < \epsilon < 1$ and $\delta > 0$. On nonempty F_t, put*

$$D_t = K_t - \frac{\operatorname{Tr}(K_tG_t)}{\operatorname{Tr}(G_t^2)}G_t, \qquad f_t = \frac{\|D_t\|_{\mathrm{HS}}^2}{\sqrt{\|D_t\|_{\mathrm{HS}}^2 + \delta^2}}, \qquad C_t = I_{F_t} - \epsilon\frac{D_t}{\sqrt{\|D_t\|_{\mathrm{HS}}^2 + \delta^2}}.$$

On $F_t = \{0\}$ set all quantities to zero. This feedback admits the stopped finite-shell construction, satisfies $\operatorname{Tr}(C_tG_t) = \operatorname{Tr} G_t$, and has $a_t = \epsilon f_t \ge 0$. On the common good set of Theorem 5.1 with $c = 2$,

$$J(L) := \mathbb{E}_W \int_0^\tau f_t dt = O_\epsilon(1/\log d). \tag{31}$$

Proof. Orthogonal projection in the operator Hilbert space gives $\operatorname{Tr}(DG) = 0$ and $\operatorname{Tr}(DK) = \|D\|_{\mathrm{HS}}^2$. The spectrum of C is in $[1-\epsilon, 1+\epsilon]$, and $\|C - I_F\|_{\mathrm{HS}} \le \epsilon$. Hence

$$\|\sqrt{C} - I_F\|_{\mathrm{HS}} \le \frac{\epsilon}{1 + \sqrt{1-\epsilon}},$$

so $B(L)$ is uniformly bounded. Theorem 5.1 gives $\mathcal{A} = O_\epsilon(1/\log d)$, and $\mathcal{A} = \epsilon J$ proves (31).

For existence fix a contact phase, where F is constant. On every compact interval $t \le T - \eta$ in the spectral band, the heat formula, positive regularizer and positive covariance lower bound make the coefficients smooth in A. The standard local strong SDE construction continues until a new hit or spectral exit. Freeze all simultaneous hits and repeat in the smaller F. There are finitely many classes, hence finitely many phases. Bounded covariance implies a continuous martingale limit at T; the contact inequalities extend by continuity. This constructs the single stopped process in the statement. □

The potential also links gain to retained dimension. Put $D_0 = \dim V = d(d+1)/2$ and $M_0 = a_0 + w$. Since every contact removes at most one dimension, $\dim F_t \ge D_0 - N_t$. Equation (29) has eigenvalues at least M_0^{-2}. Combining (30) with (6), any curvature-preserving admissible policy satisfies

$$\mathbb{E}_W \log \operatorname{vol}(\mathcal{E}_{A_\tau}) \ge \log \operatorname{vol}(\mathcal{E}_{A_0}) + \frac{D_0\mathbb{E}_W\tau - \Gamma_0 + \mathcal{A}(L)/2}{4M_0^2}. \tag{32}$$

This formula identifies the precise occupation gain controlled by the theorem.

Corollary 6.2 (Rotating perturbations)**.** *Let $0 \le C \le cI_F$ and $a_t \ge 0$. If*

$$T^{-1}\mathbb{E}_W \int_0^\tau \operatorname{rank}(C_t - I_{F_t})dt = o(\log d),$$

uniformly on a set of lattices, then the gain tends uniformly to zero on its intersection with $\mathcal{G}_d$. The perturbed subspace may vary with time.

Proof. Functional calculus gives $\|\sqrt{C}-I_F\|_{\mathrm{HS}}^2 \le M_c \operatorname{rank}(C-I_F)$, where $M_c = \max_{0\le s\le c}(\sqrt{s}-1)^2$. Apply Theorem 5.1. □

For a concrete rotating two-mode policy choose a predictable two-dimensional subspace $E_t \subset F_t$, with compressed operators K_E, G_E. Define

$$B_0 = K_E - \frac{\operatorname{Tr}(K_E G_E)}{\operatorname{Tr}(G_E^2)} G_E, \quad B_E = \frac{B_0}{\sqrt{\|B_0\|_{\mathrm{HS}}^2 + \delta^2}}, \quad C = I_F - \epsilon B_E,$$

extending B_E by zero off E and using $C = I_F$ when such a subspace is unavailable. Then $\operatorname{Tr}(B_E G) = 0$, $\operatorname{rank}(C - I_F) \le 2$, and

$$a = \frac{\epsilon \|B_0\|_{\mathrm{HS}}^2}{\sqrt{\|B_0\|_{\mathrm{HS}}^2 + \delta^2}} \ge 0.$$

Its budget is bounded by $2\max_{s\in[1-\epsilon,1+\epsilon]}(\sqrt{s}-1)^2$. Every admissible realization therefore has vanishing gain on $\mathcal{G}_d$, regardless of how many different modes are visited. Smooth or locally Lipschitz frame choices on contact phases give the local SDE construction above; other predictable choices are covered when the stipulated solution exists.

7 Optimization with the actual square-root cost

The budget has a direct finite-dimensional variational meaning. Let F now be any real Hilbert space of finite dimension r, and let $K \ge 0$, $G > 0$ be self-adjoint operators on it. Define

$$a(C) = \operatorname{Tr}((I - C)K), \qquad b(C) = \|\sqrt{C} - I\|_{\mathrm{HS}}^2 = \operatorname{Tr}(C + I - 2\sqrt{C}).$$

The latter equals the squared Bures-Wasserstein distance from C to I [4, equation (1)].

Theorem 7.1 (Exact penalized controller)**.** *Suppose $r \ge 1$, $c > 1$, $\beta > 0$. The unique maximizer of $a(C) - \beta b(C)$ subject to*

$$0 \le C \le cI, \qquad \operatorname{Tr}(CG) = \operatorname{Tr} G$$

is $C_\lambda = f(K + \beta I + \lambda G)$, where functional calculus uses

$$f(h) = \begin{cases} c, & h \le \beta/\sqrt{c}, \\ \beta^2/h^2, & h > \beta/\sqrt{c}. \end{cases} \tag{33}$$

The real multiplier λ is the unique solution of $\operatorname{Tr}(GC_\lambda) = \operatorname{Tr} G$. The optimizer satisfies $a(C_\lambda) \ge \beta b(C_\lambda) \ge 0$. For $r = 0$, the unique empty controller has zero gain and cost.

Proof. Up to constants, the negative penalized objective with multiplier λ is

$$\mathcal{L}_\lambda(C) = \operatorname{Tr}(H_\lambda C) - 2\beta \operatorname{Tr}\sqrt{C}, \qquad H_\lambda = K + \beta I + \lambda G.$$

Diagonalize C with orthonormal eigenvectors v_j and eigenvalues $x_j \in [0, c]$. Let $s(h) = \min_{0\le x\le c}(hx - 2\beta\sqrt{x})$. Scalar differentiation shows that its unique minimizer is (33), and it is positive. Equivalently, diagonalizing H_λ, pinching C into its eigenspaces preserves $\operatorname{Tr}(H_\lambda C)$ and increases $\operatorname{Tr}\sqrt{C}$ by concavity. Within a scalar eigenspace the scalar minimization gives $f(h)I$. Uniqueness also follows from strict convexity of $-\operatorname{Tr}\sqrt{C}$: on positive definite matrices its second derivative in every nonzero symmetric direction is strictly positive, by the divided differences of $x \mapsto 1/(2\sqrt{x})$. A minimizer cannot have a zero eigenvalue, since increasing that

eigenvalue by $s > 0$ decreases the square-root term by $2\beta\sqrt{s}$ against an $O(s)$ linear term. This proves that the unique minimizer of $\mathcal{L}_\lambda$ is C_λ.

The weighted trace $\psi(\lambda) = \operatorname{Tr}(GC_\lambda)$ is continuous. For $\lambda_1 < \lambda_2$, adding the two minimizing inequalities gives $(\lambda_2 - \lambda_1)(\psi(\lambda_2) - \psi(\lambda_1)) \le 0$. Since $G > 0$, for sufficiently negative λ the operator H_λ is negative definite and $C_\lambda = cI$; as $\lambda \to +\infty$, its least eigenvalue tends to infinity and $C_\lambda \to 0$. Thus ψ takes the value $\operatorname{Tr} G$, strictly between its two limits.

For uniqueness at that value, suppose two multipliers have the same weighted trace. Equality holds in both minimizing inequalities; uniqueness of the matrix minimizer forces the same C. Its eigenvalues are all positive and at least one eigenvalue x is strictly below c, since its weighted trace is less than $c \operatorname{Tr} G$. For a corresponding unit eigenvector v, the feasible variation $C + svv^T$ is two-sided for small s. Stationarity of both Lagrangians implies, by subtraction, $(\lambda_2 - \lambda_1)\langle Gv, v\rangle = 0$, a contradiction. This argument includes clipping thresholds. At the root, the Lagrangian minimum proves optimality under the trace constraint. Comparing to the feasible identity gives $a - \beta b \ge 0$. □

This formula retains G inside H_λ; it requires no simultaneous diagonalization of G and K. A congruence that whitens G changes the square-root cost centered at the identity and would solve a different problem. Used as feedback, the optimizer is subject to the admissibility conditions of Section 2; its nonnegative gain permits direct application of Theorem 5.1 with its actual $B(L)$.

Proposition 7.2 (The small-budget value with an explicit error)**.** *Fix $K \ge 0$, $G > 0$ on a nonzero finite-dimensional F and $c > 1$, and set*

$$D = K - \frac{\operatorname{Tr}(KG)}{\operatorname{Tr}(G^2)} G, \qquad S = \|D\|_{\mathrm{HS}}, \qquad M = \|D\|_{\mathrm{op}}.$$

For $\delta \ge 0$ define the attained maximum

$$V(\delta) = \max\{a(C) : 0 \le C \le cI,\ \operatorname{Tr}(CG) = \operatorname{Tr} G,\ b(C) \le \delta\}.$$

If $D = 0$, then $V(\delta) = 0$ for every δ. Otherwise, whenever

$$v := \frac{M\sqrt{\delta}}{S} \le \min\{1, (c-1)/2\},$$

one has

$$2S\sqrt{\delta} - M\delta \le V(\delta) \le 2S\sqrt{\delta} + M\delta. \tag{34}$$

The lower bound is realized as a feasible gain by $C_t = I - tD$, where $t = 2\sqrt{\delta}/S - M\delta/S^2$. In particular $V(\delta) = 2\|D\|_{\mathrm{HS}}\sqrt{\delta} + O(\delta)$ as $\delta \downarrow 0$, for the fixed operators and ceiling. For $F = \{0\}$, define $V(\delta) = 0$.

Proof. The feasible set is compact and contains I, so the maximum exists. The trace constraint gives $a(C) = -\operatorname{Tr}[D(C - I)]$. Put $E = \sqrt{C} - I$, so $b(C) = \operatorname{Tr} E^2$ and $C - I = 2E + E^2$. Cauchy–Schwarz and $-\operatorname{Tr}(DE^2) \le M \operatorname{Tr} E^2$ yield

$$a(C) \le 2S\sqrt{b(C)} + Mb(C) \le 2S\sqrt{\delta} + M\delta.$$

If $D = 0$ every feasible gain is zero. Otherwise put $u = \sqrt{\delta}/S$, so $v = Mu$ and $t = 2u - Mu^2 \ge 0$. As $\operatorname{Tr}(DG) = 0$, the proposed C_t preserves curvature exactly. Moreover $tM = 2v - v^2 \le 1$ and $tM \le 2v \le c - 1$, proving $0 \le C_t \le cI$. For each eigenvalue d_j of D,

$$(\sqrt{1 - td_j} - 1)^2 = \frac{t^2 d_j^2}{(1 + \sqrt{1 - td_j})^2}.$$

Since $\sqrt{1-tM} = 1 - v$, summation gives

$$b(C_t) \le \frac{t^2 S^2}{(2-v)^2} = \delta, \qquad a(C_t) = tS^2 = 2S\sqrt{\delta} - M\delta.$$

These identities also cover $\delta = 0$ and the allowed zero-eigenvalue endpoint. □

The proof uses elementary Hilbert–Schmidt duality and scalar square-root identities. It identifies the local value and a feasible direction for the same cost as Theorem 7.1; the long-time, dimension-dependent policy statement remains Theorem 5.1.

Two exact demonstrations

Take $c = 2$, $\beta = 3/(2\sqrt{5})$, and

$$G = I_2, \quad K = \begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix}, \quad \lambda = \tfrac{1}{2} - \beta.$$

Then $H_\lambda = \mathrm{diag}(1/2, 3/2)$, both eigenvalues exceed $\beta/\sqrt{2}$, and

$$C_\lambda = \mathrm{diag}(9/5, 1/5), \quad a(C_\lambda) = 4/5, \quad b(C_\lambda) = 4 - 8/\sqrt{5} > 0. \tag{35}$$

For a noncommuting example retain c, β, C_λ and set

$$G = \begin{pmatrix} 1 & 1/2 \\ 1/2 & 1 \end{pmatrix}, \quad \lambda = -1, \quad K = \mathrm{diag}(1/2, 3/2) - \beta I + G.$$

Here $G > 0$, $K > 0$ (its first leading minor and determinant are positive), and $KG \neq GK$. Nevertheless H_λ is the same diagonal matrix, $\mathrm{Tr}(GC_\lambda) = \mathrm{Tr}\, G = 2$, and both values in (35) hold. Repeating either block gives additive gain and cost. The accompanying portable checker verifies these identities, the clipping branch and numerical noncommuting instances.

8 The budget as a policy diagnostic

The parameters give a quantitative diagnostic for feedback design. On the selected Haar ensemble, persistent positive occupation gain entails leaving the uniformly sublogarithmic-budget class. The finite-shell potential, common envelopes and explicit optimizer express that diagnostic using one diffusion cost from instantaneous design through the full stopped trajectory.

Declarations

AI assistance. OpenAI ChatGPT/Codex assisted with literature discovery, mathematical exploration and proof development, manuscript drafting, verification-code development and internal critical checking. The author assumes responsibility for all content.

Competing interests. Scott Schubr owns Luminary Oddities, which operates the Research Archive hosting the preprint of this article.

Funding. This work received no specific grant from any funding agency, commercial or not-for-profit sector.

Data and source availability. Complete proofs are included in the article. The accompanying source package contains editable LaTeX, reproducible build instructions and the portable demonstrations.

References

[1] B. Klartag, *Lattice packing of spheres in high dimensions using a stochastically evolving ellipsoid*, arXiv:2504.05042v2 (2026). https://arxiv.org/abs/2504.05042v2.

[2] B. Klartag, *Poisson approximation of random lattices*, arXiv:2605.10479v1 (2026), equation (7). https://arxiv.org/abs/2605.10479v1.

[3] D. Kelmer and S. Yu, *The second moment of the Siegel transform in the space of symplectic lattices*, arXiv:1802.09645v1 (2018), Introduction (classical Rogers formula). https://arxiv.org/abs/1802.09645v1.

[4] R. Bhatia, T. Jain and Y. Lim, *On the Bures-Wasserstein distance between positive definite matrices*, arXiv:1712.01504v1 (2017). https://arxiv.org/abs/1712.01504v1.

[5] N. Gargava, M. Gröbner and M. Simkin, *A cyclotomic gluing construction for lattice sphere packings*, arXiv:2610.12619v1 (2026). https://arxiv.org/abs/2610.12619v1.