The non-reciprocal entropy–production bound: proven cases, the $Q = I$ premise, an exact reduction, and an open residual

Yutian Shen [14*] Xiuting Sun [2*] Weiting Liu [3*]

[1]ZJU–Xinchang Joint Innovation Centre (TianMu Laboratory), Gaochuang Hi-Tech Park, Shaoxing 312500, China

[2]School of Aerospace Engineering and Applied Mechanics, Tongji University, Shanghai 200092, China

[3]School of Mechanical Engineering, Zhejiang University, Hangzhou 310027, China

[4]Zhejiang Key Laboratory of Additive Manufacturing Technology and Equipment, Gaochuang Hi-Tech Park, Shaoxing 312500, China

*Lead and corresponding author (Yutian Shen).

Email:mz120210804@stu.yzu.edu.cn

Co-corresponding authors (Xiuting Sun; Weiting Liu).

Abstract

Non-reciprocal couplings can make added dissipation *raise* the barrier that protects an attractor against rare escape — a protective reversal impossible at equilibrium. We ask, and bound, how strong this effect can be per unit of entropy production. For a linear stochastic system $\dot{x} = Lx + \sqrt{\varepsilon}\,\xi$, $\langle \xi\xi^{\top} \rangle = Q$, with Hurwitz $L =- A + N$ ($A = A^{\top} \succ 0$, $N = - N^{\top}$), the barrier's sensitivity to a dissipation increment $M \succcurlyeq 0$ is the linear form $\langle M, K \rangle$ of a response operator K built from the stationary Gramian, and a reversal occurs iff $\lambda_{\max}(K) > 0$. Under isotropic noise $Q = I$ we conjecture the universal bound $\sup_N [\lambda_{\max}(K)]_+ / \dot{S}_i \leq 1/\left(128\, a_1^3\right)$, $a_1 = \lambda_{\min}(A)$, and prove it — sharp constant included — in the following cases: the value $1/\left(128 a_1^3\right)$ and the saturating $7:1$ geometry in the small-coupling limit in every dimension (an explicit copositivity certificate; unconditional for $d \leq 4$, otherwise conditional on a rank-two reduction we verify numerically), and the full two-mode case for all couplings. An exact Schur complement lowers the governing matrix inequality by one dimension, and we show the isotropy hypothesis is *load-bearing* — for general Q the ratio is unbounded, so $Q = I$ is not a normalisation. A free-coordinate reduction recasts the general statement as a matrix inequality on a compact box and reduces $d = 3$ to a matrix sum-of-squares certificate (whose existence the Hol–Scherer theorem guarantees where the matrix is strictly positive). The residual — $d \geq 3$ at all couplings — remains open; it is numerically unbroken over $> 10^7$ configurations up to $d = 12$. Ensemble simulations reproduce every prediction from trajectory data on circuit, optical, mechanical, and quantum realisations.

Keywords: non-reciprocity; stochastic thermodynamics; entropy production; rare-event escape; Freidlin–Wentzell quasipotential; Lyapunov equation; copositivity; sum-of-squares.

Introduction

At equilibrium, dissipation is purely corrosive of metastability: adding friction along the direction in which a system escapes its attractor can only lower the Freidlin–Wentzell barrier that renders escape rare. Non-reciprocal couplings — non-conservative interactions in which the action of one mode on another is not matched by the reaction, as arise generically in driven, active, and reservoir-engineered matter [10–13] — overturn this intuition. A suitably oriented *increment* of dissipation can then *raise* the barrier, suppressing rare escape rather than promoting it. This protective reversal is a strictly non-equilibrium effect — it is absent in equilibrium — and it is sustained only at the cost of a nonzero entropy-production rate. The operative question is therefore quantitative: how much barrier protection can non-reciprocity purchase per unit of dissipation, and is that exchange rate bounded?

We answer this for linear (Ornstein–Uhlenbeck) dynamics, the regime in which the stationary distribution is Gaussian, the quasipotential is quadratic, and every relevant object is available in closed form. The sensitivity of the barrier to a dissipation increment $M \succcurlyeq 0$ is the linear form $\langle M, K \rangle$ of a single *response operator* K assembled from the stationary covariance; a protective reversal exists precisely when $\lambda_{\max}(K) > 0$, which in turn demands $N \neq 0$. Our central result is a conjectured universal ceiling on the protection-per-entropy ratio,

$$\sup_{N} \frac{[\lambda_{\max}(K)]_+}{\dot{S}_i} \;\le\; \frac{1}{128\, a_1^3}, \qquad a_1 = \lambda_{\min}(A),$$

for isotropic noise $Q = I$; we refer to it as the *SL bound*. We *prove* it in the small-coupling limit in every dimension and, fully, for $d = 2$; in general it remains a conjecture, supported by an exact reduction and by extensive numerics (made precise below). The constant is sharp: it is approached only by a *single* slow channel whose partner mode dissipates at exactly $7a_1$ — a clean $7{:}\,1$ selection rule, with the reversal switching on at the $3{:}\,1$ ratio. The ceiling is fixed entirely by the slowest mode and is otherwise independent of dimension.

We prove [eq:main] in well-separated cases, and we are deliberate about the boundary between what is proven and what remains open. The value $1/\left(128a_1^3\right)$ and the $7{:}\,1$ equality are theorems in every dimension in the small-coupling limit, through an explicit copositivity certificate on a degree-three moment cone (Theorem 2; unconditional for $d \leq 4$, and for $d \geq 5$ conditional on a rank-two structure we verify numerically); the two-mode case is proven for all couplings (Proposition 2); and an exact Schur complement lowers the governing matrix inequality by one dimension (Proposition 3). We then show that the isotropy hypothesis cannot be relaxed: for general Q the ratio is unbounded (Proposition 1), so $Q = I$ is not a normalisation — and, when the dissipation is anisotropic, it is a non-equilibrium condition. A free-coordinate parametrisation eliminates the Lyapunov constraint and recasts the general statement as a matrix inequality on a compact box; in $d = 3$ this reduces to a matrix sum-of-squares certificate that the Hol–Scherer theorem guarantees wherever the matrix is strictly positive. The matrix is verified positive only pointwise (numerically) and is rank-deficient at the saturating point, so the $d = 3$ construction is not yet closed (Proposition 5). The general case ($d \geq 3$, all couplings) thus remains open — reduced to that conditional certificate for $d = 3$, and supported only numerically for $d \geq 4$ (unbroken over $> 10^7$ configurations up to $d = 12$).

Because [eq:main] concerns the linear covariance dynamics alone, it is realised identically across physically disparate systems. We reproduce every prediction by direct simulation, from trajectory data only, on a non-reciprocal RLC circuit, coupled optical microcavities, Brownian oscillators, and quantum non-reciprocal oscillators, and we confirm its robustness to non-Gaussian and non-Markovian baths. The remainder of the paper develops these results in the order just outlined.

Setup and the standing hypothesis

Let $A = A^\top \succ 0$, $N = -N^\top$ on $\mathbb{R}^d$, $L = -A + N$ (automatically Hurwitz: $\mathrm{Re}\lambda(L) \leq -\lambda_{\min}(A)$), and noise covariance $Q \succ 0$. The stationary covariance $\Sigma \succ 0$ solves the Lyapunov equation [8,9]

$$L\Sigma + \Sigma L^\top = -Q, \qquad \Sigma = \int_0^\infty e^{Ls}\, Q\, e^{L^\top s}\, ds.$$

Let v_1 be a unit top eigenvector of Σ, and Y solve the adjoint Sylvester equation

$$L^\top Y + YL = v_1 v_1^\top, \qquad Y = -\int_0^\infty e^{L^\top s}\, v_1 v_1^\top e^{Ls}\, ds \preccurlyeq 0.$$

Define the *response operator* and the *entropy–production rate*

$$K = \Sigma Y + Y\Sigma, \qquad \dot{S}_i = 2\,\mathrm{tr}\quad(\Phi^\top Q^{-1} \Phi\, \Sigma), \quad \Phi = L + 1/2\, Q \Sigma^{-1}.$$

The barrier sensitivity to a unit-trace increment $M \succcurlyeq 0$ is $\langle M, K\rangle$, maximised over M by $\lambda_{\max}(K)$; reversal (a protective barrier rise) occurs iff $\lambda_{\max}(K) > 0$. For $Q = I$, $\dot{S}_i = -2\mathrm{tr}(\Sigma^{-1}\Omega^2)$ with $\Omega = 1/2\,(L\Sigma - \Sigma L^\top)$. The escape barrier is the Freidlin–Wentzell quasipotential of the rare-transition problem [1,2], and the ratio $\lambda_{\max}(K)_+/\dot{S}_i$ — protection gained per unit entropy production — is in the spirit of thermodynamic uncertainty relations [3–7]. Non-reciprocity ($N \neq N^\top$) is essential: for $N = 0$ no increment reverses the barrier [10–12].

Theorem 1 (the SL bound — proven cases; $Q = I$). *Let $A \succ 0$, $N = -N^\top$, $Q = I$, $a_1 = \lambda_{\min}(A)$, and $C(A) = \sup_N [\lambda_{\max}(K)]_+/\dot{S}_i$. The bound $C(A) \leq 1/\left(128\, a_1^3\right)$ holds, with the constant sharp, in each of the following cases:*

1. *(Small-coupling limit, every* d*.) As* $t \to 0$ *the supremum equals* $1/\left(128\, a_1^3\right)$*, attained by a single slow channel with partner mode at* $7a_1$ *(the* $7{:}\,1$ *law), via an explicit copositivity certificate on the degree-3 Stieltjes moment cone (Theorem 2). This is unconditional for* $d \le 4$ *and, for* $d \ge 5$*, conditional on the rank-two structure of Lemma 2, which we have verified numerically.*
2. *(Full two modes.) For* $d = 2$ *and all couplings,* $\dot{S}_i = 4t^2/(a_1 + a_2)$ *is exact and the efficiency is monotone in the coupling, so* $C(A) \le 1/\left(128\, a_1^3\right)$ *(Proposition 2).*

Conjecture 1 (the SL bound). For every $A \succ 0$ and $N = -N^\top$ at $Q = I$, and in every dimension, $C(A) \le 1/\left(128\, a_1^3\right)$.

The open part is $d \ge 3$ at all couplings. It is genuinely non-perturbative — the efficiency is not ray-monotone for $d \ge 3$, so it does not reduce to the small-coupling case. For $d = 3$ the inequality is exactly equivalent to two polynomial nonnegativities on the box $[0,1]^2$; the leading coefficient is proven nonnegative and the matrix form $H(r_2, r_3) \succcurlyeq 0$ is verified *pointwise* (numerically). Wherever $H \succ 0$ strictly, the Hol–Scherer theorem guarantees a matrix-SOS certificate (Proposition 5); but H is rank-deficient at the saturating $7{:}\,1$ point — the boundary of the PSD cone, where strict positivity fails — so this guarantee does not by itself close $d = 3$ (the $P1$ block is strictly SOS-feasible in a floating-point solve; the $q16$ block is not yet certified). For $d \ge 4$ only numerical evidence is available. The conjecture is unbroken over $> 10^7$ configurations up to $d = 12$ with no counterexample (Remark 3).

Standing hypothesis ($Q = I$*) — and why it cannot be relaxed.*

The constant $1/\left(128 a_1^3\right)$ is specific to isotropic white noise and is *not* a normalisation.

Proposition 1 (no general-Q bound). *Fix L (hence* A, a_1*) and scale* $Q \mapsto cQ$*,* $c > 0$*. Then* $\Sigma \mapsto c\Sigma$*, while* v_1 *and* Y *are unchanged, so* $K \mapsto cK$ *and* $\lambda_{\max}(K) \mapsto c\,\lambda_{\max}(K)$*; and* $\Phi = L + 1/2\, Q\Sigma^{-1}$ *and* $\dot{S}_i = 2\mathrm{tr}(\Phi^\top Q^{-1}\Phi\Sigma)$ *are both invariant. Hence the metric* $128\, a_1^3\, \lambda_{\max}(K)/\dot{S}_i$ *scales linearly in* c *and is unbounded over* $Q \succ 0$*.*

This already fails for Q commuting with A. (Numerical witness: $A = \mathrm{diag}(1, 3.9118)$, $N_{12} = 0.12$, $Q = c\,\mathrm{diag}(5.957, 2.059)$ gives metric $6.89, 13.79, 34.46, 137.85, 689.27$ for $c = 1, 2, 5, 20, 100$ — exactly linear; a non-commuting Q gives the metric scaling by c as well.) Thus $Q = I$ must be assumed throughout; there is no Q-independent analogue of the constant (Fig. 1).

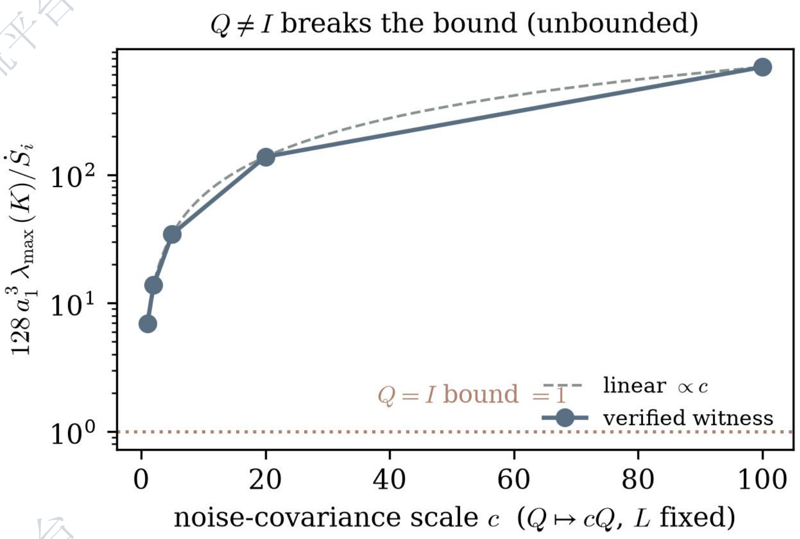

The isotropy hypothesis is load-bearing (Proposition 1). For a fixed operator L and a noise covariance scaled as $Q \mapsto cQ$, the metric $128\, a_1^3\, \lambda_{\max}(K)/\dot{S}_i$ grows linearly in c (verified witness, log scale) and exceeds the $Q = I$ value 1 without bound; there is no general-Q version of the constant.

Remark 1 (homogeneity). With $Q = I$, under $A \mapsto sA$, $N \mapsto sN$ one has $\Sigma \mapsto \Sigma/s$, v_1 fixed, $Y \mapsto Y/s$, $K \mapsto K/s^2$, $\dot{S}_i \mapsto s\dot{S}_i$; the efficiency scales as s^{-3} and $a_1 \mapsto sa_1$, so $C(A) = a_1^{-3} C(A/a_1)$ and it suffices to treat $a_1 = 1$, i.e. $A \succcurlyeq I$.

Throughout, $\{X, Z\} = XZ + ZX$, $[X, Z] = XZ - ZX$, and $Q = I$ unless stated otherwise.

The small-coupling bound: a theorem in all dimensions

Fix $A = \mathrm{diag}(a_1, ..., a_d)$, $1 = a_1 < a_2 \leq \cdots$ (degeneracy by continuity, Remark 2). Put $N = tN_0$, $t \to 0$.

Lemma 1. *$\dot{S}_i = 4t^2 \sum_{i<j} N_{0,ij}^2/\left(a_i + a_j\right) + O(t^3)$. In $d = 2$ this is exact: $\dot{S}_i = 4t^2/(a_1 + a_2)$.*

Proof. At $t = 0$, $\Sigma_0 = \mathrm{diag}(1/2a_i)$, so $\Phi_0 =- A + 1/2\, \Sigma_0^{-1} = 0$ and $\dot{S}_i = O(t^2)$. Differentiating [eq:lyap]: $(\Phi_1)_{ij} = 2a_j N_{0,ij}/\left(a_i + a_j\right)$, whence $\dot{S}_i = 2t^2 \mathrm{tr}(\Phi_1^\top \Phi_1 \Sigma_0) + O(t^3) = 4t^2 \sum_{i<j} N_{0,ij}^2/\left(a_i + a_j\right) + O(t^3)$. The $d = 2$ exactness is a direct symbolic computation. □

At $t = 0$, $K_0 =- 1\big/2a_1^2\, e_1 e_1^\top$ with kernel $\mathcal{K} = \mathrm{span}(e_2, ..., e_d)$ and $K_1|_{\mathcal{K}} = 0$ (quadratic onset). Degenerate second-order perturbation theory on $\mathcal{K}$ gives $P_{ab} = (K_2)_{ab} + 2a_1^2 (K_1)_{a1} (K_1)_{1b}$ $(a, b \geq 2)$, with $\lambda_{\max}(K) = t^2 \lambda_{\max}(P) + o(t^2)$.

Lemma 2 (rank-two structure). *P depends only on the slow-mode couplings $p_j := N_{0,1j}$ (the $N_{0,jk}$, $j,k \geq 2$, drop out), and $P = 1/2\,(\xi\xi^\top - 4\eta\eta^\top)$ with $\xi_j = p_j(a_j - a_1)/[a_1(a_1 + a_j)^2]$, $\eta_j = p_j/(a_1 + a_j)^2$.*

(Symbolic for general spectrum, $d = 3,4$; confirmed against high-precision Richardson extraction of $\lambda_{\max}(K)/t^2$. Since fast–fast couplings only increase $\dot{S}_i$, the optimum has them zero.)

Set $a_1 = 1$, fast–fast couplings zero. With $w_j = p_j^2/(1 + a_j)^4 \geq 0$, $b_j = a_j - 1 \geq 0$, and moments $\mu_k = \sum_j w_j b_j^k$ of $\nu = \sum_j w_j \delta_{b_j}$ on $[0,\infty)$, the efficiency $\mathcal{E} = \lambda_{\max}(P)/\left(4\sum_j p_j^2/(1 + a_j)\right)$ obeys, with $\lambda_{\max}(P) = 1/4\,\left(T + \sqrt{T^2 + 8V}\right)$, $T = \mu_2 - 4\mu_0$, $V = 2(\mu_0\mu_2 - \mu_1^2)$,

$$\mathcal{E} \leq 1/128 \Leftrightarrow F := R^2 - 2RT - 8V \geq 0, \qquad R = \mu_0 + 3/2\,\mu_1 + 3/4\,\mu_2 + 1/8\,\mu_3.$$

The equivalence squares $\sqrt{T^2 + 8V} \leq R - T$; this is reversible because $R - T > 0$ on the moment cone (each atom contributes a positive amount).

Theorem 2 (small-coupling bound). *$F \geq 0$ for every nonnegative measure on $[0,\infty)$. Hence $C_{\text{pert}}(A) := \lim_{t\to 0}\sup_{N_0}\lambda_{\max}(K)_+/\dot{S}_i = 1/(128\,a_1^3)$ in every dimension — the reduction to F through Lemma 2 being unconditional for $d \leq 4$ and verified numerically for $d \geq 5$ — with equality iff ν is a single atom at $b = 6$, a single slow-mode channel with $a_j = 7a_1$ (the $7{:}1$ law).*

Proof. F depends only on $(\mu_0, \mu_1, \mu_2, \mu_3)$. By the two-node Gauss–Radau / Tchakaloff theorem [17], every point of the degree-3 Stieltjes moment cone of $[0,\infty)$ is represented by a measure with at most two atoms; so $F \geq 0$ on the cone $\Leftrightarrow$ $\left[Q(b_i, b_j)\right]_{2\times 2}$ is copositive for all $b_1, b_2 \geq 0$, where

$$Q(x,y) = 1/64\,(2+x)^3(2+y)^3 - 1/8\left[(2+x)^3(y^2-4) + (2+y)^3(x^2-4)\right] - 8(x-y)^2.$$

Since $Q(x,x) = 1/64\,(2+x)^4(x-6)^2 \geq 0$, the 2×2 copositivity criterion reduces this to $\Psi(x,y) := Q(x,y) + 1/64\,(2+x)^2(2+y)^2|x-6||y-6| \geq 0$. *Same side* ($(x-6)(y-6) \geq 0$): with $f_1(x,y) = x(y-2)^2 + 2y^2 - 24y - 24$, $G_+ := Q + 1/64\,(2+x)^2(2+y)^2(x-6)(y-6) = 1/32\,f_1(x,y)f_1(y,x)$; as f_1 is increasing in x and $f_1(6,y) = 8y(y-6)$, $f_1 \leq 0$ on $[0,6]^2$ and ≥ 0 on $[6,\infty)^2$, so both factors share sign and $G_+ \geq 0$. *Opposite side*: $G_- := Q - 1/64\,(2+x)^2(2+y)^2(x-6)(y-6) = 1/2\,(x-y)^2(xy + 2x + 2y - 12)$, and $xy + 2x + 2y - 12$ is increasing with value 0 at $(0,6)$, hence ≥ 0 there. Thus $\Psi \geq 0$, $F \geq 0$, and equality is the single atom $b = 6$. □

Remark 2. If $\lambda_{\min}(A)$ is repeated, $C(A) \leq 1/(128a_1^3)$ follows by continuity; degeneracy strictly lowers the efficiency (e.g. $a = (1,1,7)$ gives $\mathcal{E}\cdot 128 \leq 0.0315$).

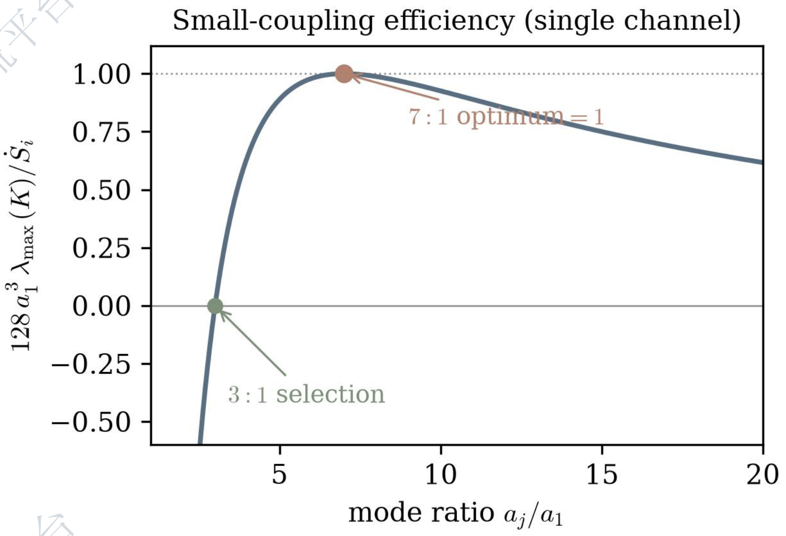

The small-coupling efficiency for a single slow-mode channel, $128\, a_1^3\, \lambda_{\max}(K)/\dot{S}_i = 16(r-3)/(1+r)^2$ *with* $r = a_j/a_1$ *(Theorem 2): reversal switches on at the* $3{:}\,1$ *ratio, and the bound is saturated exactly at the* $7{:}\,1$ *ratio (value 1, i.e.* $\mathcal{E} = 1/128\, a_1^3$*).*

The two-mode case (full coupling)

Proposition 2. *For* $A = \mathrm{diag}(a, b)$, $0 < a \le b$, $N = \begin{pmatrix} 0 & t \\ -t & 0 \end{pmatrix}$, $\lambda_{\max}(K)/\dot{S}_i \le 1/(128a^3)$ *for all* t.

Proof sketch. $\dot{S}_i = 4t^2/(a+b)$ exactly; the claim is $\lambda_{\max}(K) \le M := t^2/[32a^3(a+b)]$. As K is 2×2 this is $2M - \mathrm{tr}K \ge 0$ and $M^2 - M\mathrm{tr}K + \det K \ge 0$; the first holds with a strict margin, the second is nonnegative and tight as $t \to 0, b \to 7a$. Equivalently $\mathcal{E}(t)$ decreases in t, so $\sup_t \mathcal{E} = \mathcal{E}(0^+) \le 1/(128a^3)$ by Theorem 2. □

An exact dimension reduction (Schur)

The bound $128a_1^3\lambda_{\max}(K) \le \dot{S}_i$ is exactly $M \succcurlyeq 0$ for $M := \dot{S}_i I - 128a_1^3 K$.

Proposition 3 (Schur reduction). *In the eigenbasis with* $v_1 = e_1$, $K_{11} < 0$ *whenever* $\lambda_{\max}(K) > 0$ *(one has* $K_{11} = -2Q_{\mathrm{c}}/D_\ell$ *with* Q_{c} *copositive and* $D_\ell > 0$*). Hence* $M_{11} = \dot{S}_i - 128a_1^3 K_{11} > 0$*, and* $M \succcurlyeq 0 \Leftrightarrow M/M_{11} := M_{rr} - M_{1r}M_{1r}^{\top}/M_{11} \succcurlyeq 0$*, an exact equivalence reducing the PSD test to dimension* $d-1$*. For* $d = 3$ *this is the pair* $\mathrm{tr}(M/M_{11}) \ge 0$ *and* $\det(M/M_{11}) \ge 0$*.*

(Verified $14201/14201$ over $d = 3$–5; $M_{11} > 0$ in all cases.)

Free coordinates and the residual inequality

Work in the eigenbasis of Σ, $\Sigma = \operatorname{diag}(\sigma_1 \geq \cdots > 0)$, $v_1 = e_1$.

Lemma 3 (free coordinates). *With* $\Omega = 1/2\,(L\Sigma - \Sigma L^{\top})$, $L = (\Omega - 1/2\,I)\Sigma^{-1}$ *(*$L_{pp} =- 1/2\sigma_p$, $L_{pq} = \Omega_{pq}/\sigma_q$*); conversely every descending* $\sigma > 0$ *and antisymmetric* Ω *give an* L *solving [eq:lyap] with* $\Sigma = \operatorname{diag}(\sigma)$ *— the Lyapunov constraint is eliminated, and* $A = 1/2\,\Sigma^{-1} + 1/2\,[\Sigma^{-1}, \Omega]$.

Lemma 4. $A_{pp} = 1/(2\sigma_p)$ *for every* p *(the commutator in* $A = 1/2\,\Sigma^{-1} + 1/2\,[\Sigma^{-1}, \Omega]$ *has zero diagonal); hence* $\sigma_1 = 1/(2e_1^{\top}Ae_1)$, $\lambda_{\max}(\Sigma) \leq 1/(2a_1)$, *and* $A \succcurlyeq I \Rightarrow \sigma_p \leq 1/2$ *for all* p*. Also* $\dot{S}_i = 2\sum_{p<q} \Omega_{pq}^2 \left(1/\sigma_p + 1/\sigma_q\right)$.

Proposition 4 (box reduction). *The SL bound (Conjecture 1) is implied by: for all diagonal* $0 \prec \Sigma \preccurlyeq 1/2\,I$ *and antisymmetric* Ω, $\{\Sigma, Y\} \preccurlyeq 1/128\,\dot{S}_i\,I, L^{\top}Y + YL = e_1e_1^{\top}$, $L = (\Omega - 1/2\,I)\Sigma^{-1}$, $\dot{S}_i =- 2\operatorname{tr}(\Sigma^{-1}\Omega^2)$.

Proof. By Remark 1 it suffices to prove $\lambda_{\max}(K) \leq \dot{S}_i/128$ on $\{A \succcurlyeq I, \lambda_{\min}(A) = 1\}$; by Lemma 4, $A \succcurlyeq I \Rightarrow \sigma_p \leq 1/2$, so this surface lies in the box. □

Inequality [eq:box] is verified to machine precision over $> 5 \times 10^5$ adversarial configurations ($d = 2$–8, ill-conditioned and near-degenerate included); multistart minimisation over the box gives $\min \lambda_{\min}\left(1/128\,\dot{S}_i I - \{\Sigma, Y\}\right) = 4.3 \times 10^{-32}$, attained at $\sigma = (1/2\,, 1/2\,, 1/14)$ ($\sigma_3/\sigma_1 = 1/7$), $\Omega \to 0$ — the single-channel optimum, where it is tight.

The dimension three reduction

In free coordinates ($\sigma = (1, r_2, r_3)$, $\Omega = (w, x, y)$) with the relaxation $a_1 \to g$, $A - gI \succcurlyeq 0$ (equivalent since the bound is monotone in g), the $d = 3$ bound is *exactly equivalent* to two polynomial nonnegativities on the compact box $[0,1]^2$:

$$(d = 3\text{bound}) \Leftrightarrow \left[P1 \geq 0 \text{ and } q16 \geq 0\right] \quad \text{on} \quad \{0 \leq r_2, r_3 \leq 1,\ (w, x, y) \in \mathbb{R}^3\}.$$

The coefficient $c_0 = 2Q_{\mathrm{c}}D_{\ell}^2 SN^2 \geq 0$ is proven (copositive Q_{c} times squares). The boundary obstruction at $\Omega = 0$ is removed by the matrix reformulation: writing $P1 = m_{\Omega}^{\top}H(r_2, r_3)\,m_{\Omega}$,

$$P1 \geq 0\ \forall\Omega \Leftrightarrow H(r_2, r_3) \succcurlyeq 0 \text{ on } [0,1]^2,$$

a matrix positivity in two compact variables.

Proposition 5 ($d = 3$: reduction to a conditional certificate). $H(r_2, r_3) \succcurlyeq 0$ *at every tested point of* $[0,1]^2$ *(a pointwise, numerical semidefinite-feasibility check). On any compact region where* $H \succ 0$ *strictly, the Hol–Scherer matrix-Putinar theorem [21] (the matrix analogue of Putinar's Positivstellensatz [18,19]) guarantees a matrix-SOS* *certificate* $H = \Sigma_0 + r_2(1 - r_2)\Sigma_1 + r_3(1 - r_3)\Sigma_2 + \cdots$ *(*Σ_i *PSD matrix-SOS); likewise for q16. At the saturating 7:1 point, however,* H *is rank-deficient — a boundary of the PSD cone, where strict positivity fails — so the guarantee does not by itself close* $\mathrm{d} = 3$*: a global certificate must treat that point separately. The full* $\mathrm{d} = 3$ *bound is thus reduced to a finite SOS computation conditional on the (numerically observed) pointwise positivity and on handling the boundary.*

Status of the explicit construction.

Building these certificates is a finite computation wherever H is strictly positive (and modulo the boundary point), but heavy: the relevant Gram cones are large ($\sim 2 \times 10^3$ total PSD dimension across blocks, polynomial degrees 18 for $P1$ and 29 for $q16$). Direct attempts (degree-8 box modules; the full Schmüdgen module; Clarabel/SCS; symbolic Schur determinants) are respectively infeasible at low degree, out-of-memory (> 16 GB), or non-convergent. A direct SCS formulation in standard form (bypassing the cvxpy canonicalisation blow-up) instead solves the $P1$ block locally (four PSD cones, largest 504, in 28 GB) and certifies it *strictly* SOS-feasible (margin $t \approx + 0.016 > 0$); the surviving obstacle is the exact rationalisation of a *tight*, rank-deficient certificate, which needs low-rank extraction rather than naive rounding. Closing $d = 3$ in print therefore requires a *reduced-basis* (Newton polytope + block symmetry) certificate and exact rational rounding [20,22], not brute force — a self-contained computational-algebra task.

The analytic obstruction and the open residual

Remark 3 (the obstruction). Write $S = 1/2\,\{\Sigma, Y\}$ (so [eq:box] is $\lambda_{\max}(S) \le \dot{S}_i/256$) and $P = 1/2\,Y \preccurlyeq 0$. Exactly, $[\Sigma, P] = 1/2\,[\Sigma, Y]$ and

$$S + \sigma_1^2 e_1 e_1^\top = 2(\Sigma P \Omega - \Omega P \Sigma), \qquad u^\top(\Sigma P \Omega - \Omega P \Sigma)u = 2(\Sigma u)^\top P(\Omega u).$$

For the top eigenvector u of S, $\lambda_{\max}(S) = 4(\Sigma u)^\top P(\Omega u) - \sigma_1^2 u_1^2$, and Cauchy–Schwarz in the $(-P)$-inner product gives $4(\Sigma u)^\top P(\Omega u) \le 4\sqrt{(u^\top \Sigma(-P)\Sigma u)(\Omega u)^\top(-P)(\Omega u)}$. The numerator lives in the $(-P)$-metric (the observability Gramian of e_1, $\mathrm{tr}(-P) = \sigma_1/2$), while $\dot{S}_i$ lives in the Σ^{-1}-metric of Ω; the two are linked only through $L^\top P + PL = 1/2\, e_1 e_1^\top$, and no metric comparison closes [eq:box] tightly. The residual is genuinely non-perturbative: ray-monotonicity fails for $d \ge 3$, K obeys no closed Lyapunov equation in $\{\Sigma, Y, \Omega, e_1\}$, and the reversal is not confined to any 2×2 principal block.

Status of the theory

Statement	**Status**
Value $1/(128a_1^3)$ = small-coupling sup, all d	**Theorem** 2 (copositivity; unconditional $d \le 4$, numerical $d \ge 5$)
$7:1$ equality case ($a_j = 7a_1$)	**Theorem** 2
Full $d = 2$ (all N)	**Proposition** 2 (proof sketch)

Statement	**Status**
$Q \neq I$: ratio unbounded ($Q = I$ premise)	**Proposition** 1
Exact Schur reduction $M \succcurlyeq 0 \Leftrightarrow M/M_{11} \succcurlyeq 0$	**Proposition** 3
Robust Hurwitz Re$\lambda \leq -\, a_1$; homogeneity degree -3	**Theorem**
SL bound $C(A) \leq 1/\left(128a_1^3\right)$, general A, N	**Conjecture** 1
• $d = 3$, all couplings	**Open**; reduced to a conditional SOS certificate (5); $P1$ block strictly SOS-feasible (floating-point)
• $d \geq 4$, all couplings	**Open**; numerical only (> 10^7 evals, $d \leq 12$, no counterexample)

The physically meaningful content — maximal protection-per-entropy = $1/\left(128a_1^3\right)$, set by the slowest mode and attained at the 7: 1 ratio — is proven in every dimension *in the small-coupling limit* (Theorem 2; conditional on Lemma 2 for $d \geq 5$) and fully in $d = 2$. The residual is the closed-form proof of [eq:box] for $d \geq 3$ at all couplings, equivalently the construction of the (conditional) $d = 3$ certificate and a general-d copositivity argument; it is stated here as an open problem. All non-symbolic claims were checked numerically to machine precision; the homogeneity, copositivity, free-coordinate, Schur, and Q-scaling identities were verified symbolically.

Numerical verification

We integrate the non-reciprocal two-mode stochastic dynamics $\dot{x} = Lx + \xi$, $\langle \xi\xi^{\top} \rangle = Q$, by ensemble Euler–Maruyama and recover every prediction *from trajectory data alone* (sample covariance $\widehat{\Sigma}$, regression-estimated drift $\hat{L}$, and the entropy-production rate). The same dynamics is realised, with identical mathematics, by three classical platforms (a quantum realisation follows below): a two-node non-reciprocal RLC analog circuit (A = self-damping, N = the inverter-induced sign-asymmetric cross-paths, Q = injected noise); two non-reciprocal coupled optical microcavities (A = cavity loss $\kappa_i/2$, N = a directional/synthetic-gauge coupling, Q = balanced vacuum/thermal input); and two non-reciprocal overdamped Brownian oscillators ($A = \gamma^{-1}\mathrm{diag}(k_i)$ stiffness, N = a non-conservative coupling, $Q = 2\gamma^{-1}k_B T\, I$). The thermal/vacuum noise of the physical platforms is isotropic ($Q \propto I$) when the mode dampings are comparable, meeting the standing hypothesis (Proposition 1); strong damping asymmetry instead tracks the dampings ($Q \propto$ diag) and requires the input to be balanced by design — the quantum analogue of which is the vacuum floor discussed below — while $N \neq N^{\top}$ always requires a non-equilibrium drive.

Across 24 M pooled samples ($a = (1,7)$, $t = 1$): $\widehat{\Sigma}$ matches theory to 0.16%, $\hat{L}$ to 0.23%, $\dot{S}_i$ to 1.6%. Sweeping the mode ratio (Fig. [fig:simlaw]), the measured metric vanishes below 3: 1, peaks at 7: 1, and approaches 1 as $t \to 0$ — the laws of Theorem 2. Setting $N = 0$ gives $\lambda_{\max}(K) = 0$ (no reversal); scaling $Q \mapsto cQ$ makes the metric grow linearly and exceed 1 without bound (Proposition 1). Finally, the rare-event barrier $V^* = b^2/(2v_1^{\top}\Sigma v_1)$ is confirmed by first-passage statistics: a dissipation increment along the barrier-raising direction of K lifts V^* and lengthens the escape time 2.6 ×, while an orthogonal increment does nothing (Fig. 4) — the programmable, directional protection that non-reciprocity enables. The same dimensionless results hold for the optical (loss ratio 7: 1; bound preserved including detuning, 4D) and Brownian (stiffness ratio 7: 1) platforms.

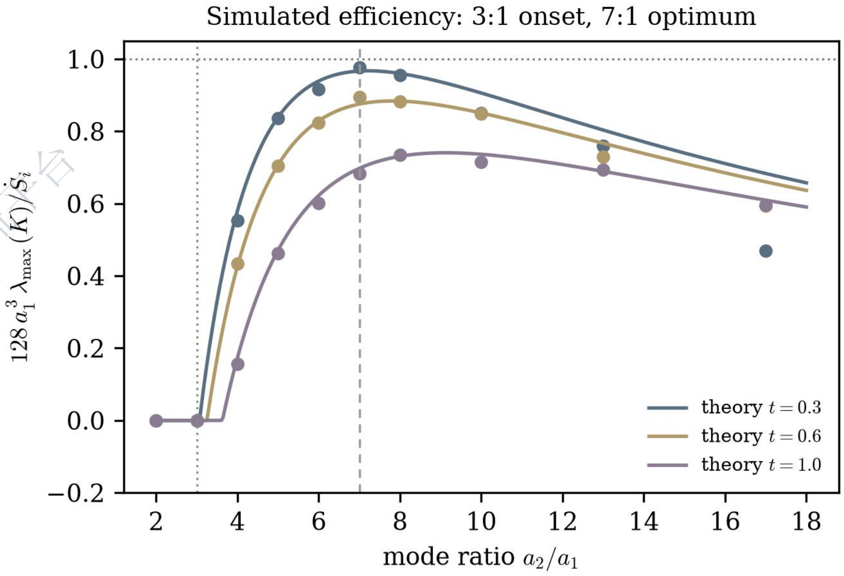

Rare-event escape (first-passage time): the protective dissipation increment raises the barrier and the escape time 2.6 ×; an orthogonal one does not.

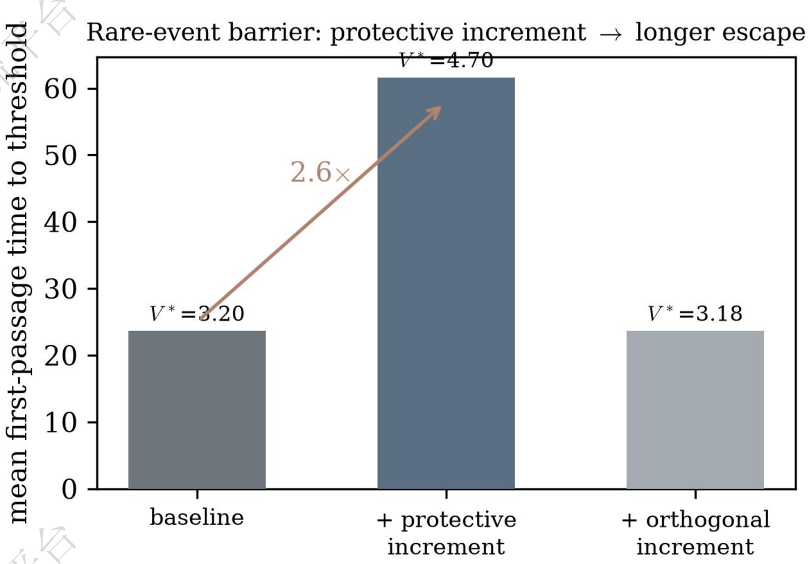

Rare-event escape (first-passage time): the protective dissipation increment raises the barrier and the escape time 2.6 ×; an orthogonal one does not.

Experimental parameters and conditions.

Table 1 gives the operating parameters of the platforms. In every case the control variables are the mode-dissipation *ratio* a_2/a_1 (loss ratio κ_2/κ_1, stiffness ratio k_2/k_1) and the non-reciprocal coupling strength; the noise is isotropic ($Q \propto I$) and supplied by the platform's native fluctuations. The simulations use an ensemble-vectorised Euler–Maruyama integrator (time step $\Delta t = 2 \times 10^{-3}$ in units of $1/a_1$, here $\Delta t = 10^{-3}$ for the Q-sweep), $R = 3\text{–}4 \times 10^3$ replicas, 2×10^4 for first-passage runs, and $6\text{–}8 \times 10^3$ steps after a burn-in of 2×10^3 ($\sim 2.4 \times 10^7$ pooled samples per point); the covariance and drift are estimated by sample statistics and linear regression, the entropy-production rate from the reconstructed irreversible drift.

Operating parameters of the four platforms (three classical, one quantum). The thermal/vacuum noise makes $Q \propto I$ automatic at moderate ratios; $N \neq N^\top$ requires a non-equilibrium element (inverter / directional coupler / non-conservative force).

Platform	**Dissipation** A	**Non-recip.** N	**Noise** Q	**Operating point**
RLC circuit	self-damping R_0/a_i, $R_0 = 100\,\mathrm{k}\Omega$, $C = 10\,\mathrm{nF}$, $\tau = 1\,\mathrm{ms}$	± cross-path (one via inverter)	digital i.i.d. ($Q = I$)	$a_2/a_1 = 7$, $t \lesssim a_1$
Optical microcavity	loss $\kappa_i/2$; $\kappa_1 = 2\pi \cdot 2$, $\kappa_2 = 2\pi \cdot 14$ MHz	directional / gauge $\Gamma = 2\pi \cdot 1\,\mathrm{MHz}$	vacuum/thermal input ($\propto I$)	$\kappa_2/\kappa_1 = 7$; collapse at $\Gamma_{\max}$

Platform	**Dissipation** A	**Non-recip.** N	**Noise** Q	**Operating point**
Brownian oscillator	$\gamma^{-1}k_i$; $\gamma = 10^{-3}$ pN s/μm, $k_1 = 1, k_2 = 7$ pN/μm	non-conservative $c = 0.8$ pN/μm	thermal $2\gamma^{-1}k_B T I$, $T = 300$ K	$k_2/k_1 = 7$; $1/a_i = 1.0, 0.14$ ms
Quantum oscillator	mode loss $\kappa_i/2$; quadratures $[x, p] = i$	reservoir-engineered Γ, J (Metelmann–Clerk)	engineered vacuum-balanced $\propto I$ (FDT $\propto \mathrm{diag}\kappa$ gives no reversal)	physical states only; metric $\leq 0.29 \ll 1$

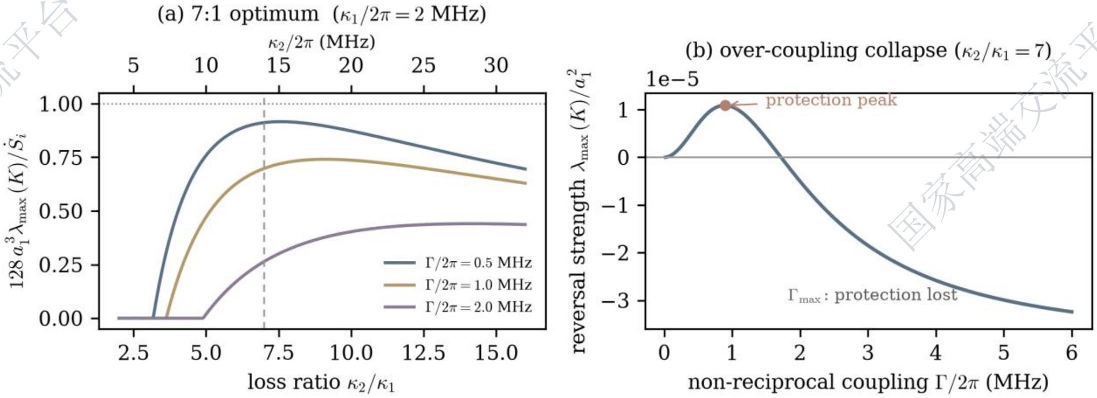

Optical microcavities (physical units). (a) Efficiency vs. loss ratio κ_2/κ_1 ($\kappa_1/2\pi = 2$ MHz; upper axis κ_2 in MHz) for three non-reciprocal couplings: 7: 1 optimum, value → 1 as $\Gamma \to 0$. (b) Over-coupling collapse at $\kappa_2/\kappa_1 = 7$: the reversal strength $\lambda_{\max}(K)$ rises, peaks, then falls through zero at $\Gamma_{\max}$ — too much non-reciprocity destroys protection.

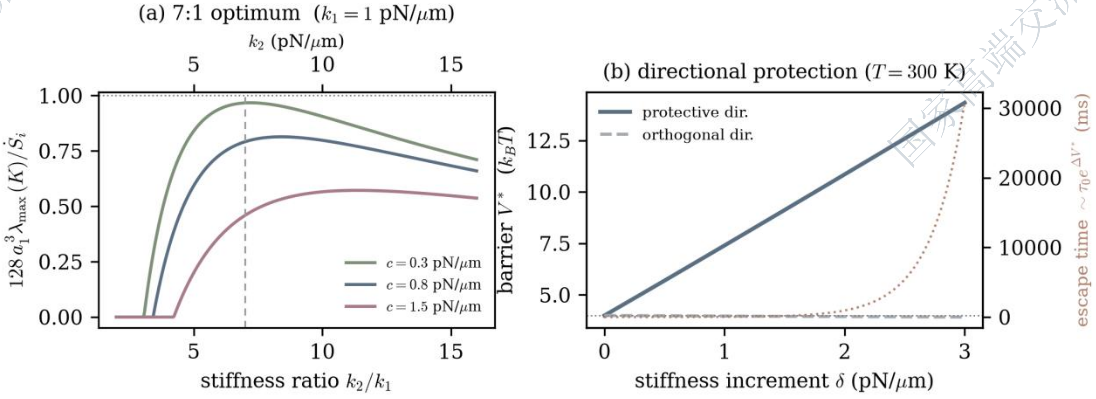

Brownian oscillators (physical units, $T = 300$ K). (a) Efficiency vs. stiffness ratio k_2/k_1 ($k_1 = 1$ pN/μm; upper axis k_2) for three non-conservative couplings: 7: 1 optimum. (b) A stiffness increment along the protective direction of K raises the escape barrier V^* (from

$4\,k_B T$) and the mean escape time $\sim \tau_0 e^{\Delta V^*}$ (from $\tau_0 = 1$ ms to $\sim 10^4$ ms); an orthogonal increment leaves it flat.

Quantum oscillators.

Because the bound governs the linear covariance dynamics, it extends to quantum Gaussian systems. We take two non-reciprocal bosonic modes (reservoir-engineered [13], Metelmann–Clerk: $H = \sum_i \omega_i n_i + J\left(a_1^\dagger a_2 + \mathrm{h.c.}\right)$, jumps $\sqrt{\kappa_i}\, a_i$ and $\sqrt{\Gamma}(a_1 + i a_2)$), build the genuine master-equation drift and the vacuum/zero-point diffusion via the standard Gaussian formulas [14–16] (quadratures r, $[x, p] = i$), and keep only *physical* states ($\Sigma + i/2\,\mathbb{J} \succcurlyeq 0$, $\mathbb{J}$ the symplectic form — distinct from the circulation Ω used above). The bound is never violated: over $> 10^4$ physical states the maximum metric is 0.29 (Fig. [fig:quantum]). Two quantum facts stand out. First, an *equilibrium* bath obeys the fluctuation–dissipation relation $D \propto \mathrm{diag}(\kappa_i) \neq I$, so there is *no* reversal (metric = 0): the $Q = I$ premise is a *non-equilibrium* condition, requiring engineered isotropic (balanced) input noise. Second, the Heisenberg uncertainty *forbids* the strongly-asymmetric $7{:}1$ optimum — it would demand sub-vacuum (squeezed) fluctuations in the high-loss mode — so the physical quantum reversal sits far below the classical $1/128$ ceiling. A direct quadrature simulation (3×10^7 samples) recovers the metric to 0.05%.

Robustness to non-ideal baths.

The white-Gaussian assumption can be relaxed. *Non-Gaussian* white baths (telegraph, Laplace, uniform; noise kurtosis $1.0, 6.1, 1.8$ vs. 3) leave Σ and the metric *identical* to $< 0.8\%$ (Fig. 8a): a linear system's covariance depends only on the second moment Q, and the state itself Gaussianises by the central-limit theorem — the bound is *exactly* robust to non-Gaussianity. *Non-Markovian* (Ornstein–Uhlenbeck colored) baths with memory time τ_c only *suppress* the metric: it decreases monotonically from the white value at $\tau_c \to 0$ (Fig. 8b, augmented-Lyapunov exact), and the bound is never exceeded — bath memory weakens, never defeats, the ceiling.

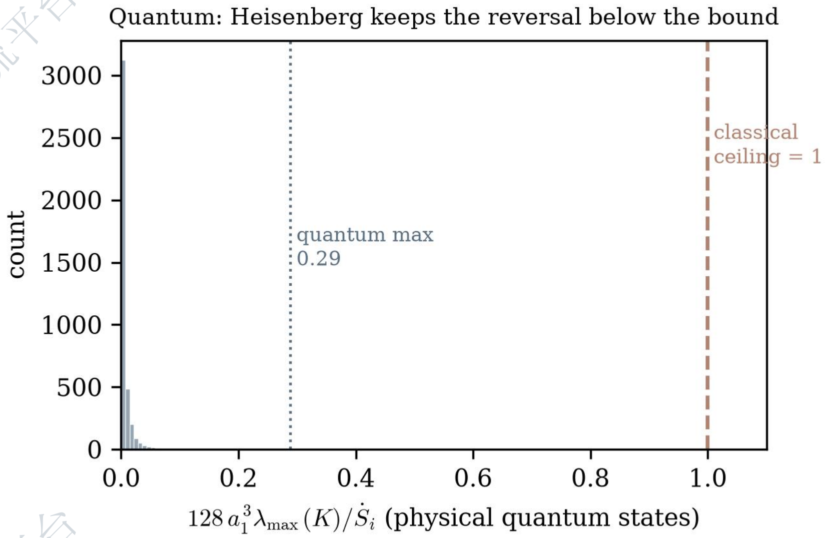

Bath robustness. (a) Non-Gaussian white baths (noise kurtosis on bars): Σ and the metric are invariant. (b) Non-Markovian colored baths: memory time τ_c only suppresses the metric; the bound 1 is never approached from above.

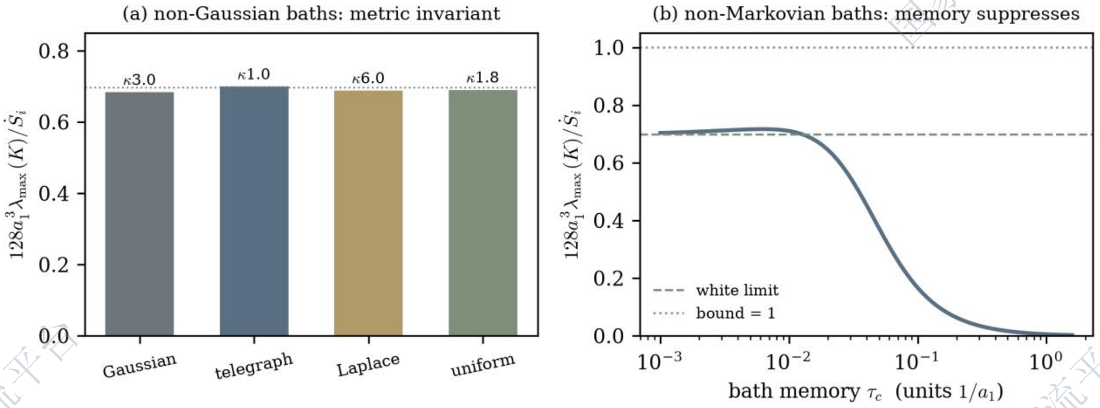

Bath robustness. (a) Non-Gaussian white baths (noise kurtosis on bars): Σ and the metric are invariant. (b) Non-Markovian colored baths: memory time τ_c only suppresses the metric; the bound 1 is never approached from above.

Global optimum, joint baths, and transient relaxation.

Three further checks close the numerical picture (Fig. 9). *(i)* $d = 3$ *global optimum.* A full-parameter global optimisation over $A = \mathrm{diag}(1, a_2, a_3)$ and *all* of N (with $Q = I$) returns a maximum metric of 1.0000, attained at $a = (1, 7.00, \text{decoupled})$ with $N \to 0$ — the $7{:}1$ single channel — while 3×10^5 random configurations never exceed 0.997. This is numerical evidence that in $d = 3$ the bound is *saturated* precisely at the perturbative $7{:}1$ optimum (consistent with Theorem 1(i) and the $d = 3$ reduction), not a proof of the $d = 3$ case. *(ii) Joint non-ideal bath.* Driving each channel with an Ornstein–Uhlenbeck process whose innovations are non-Gaussian (telegraph, Laplace) — colored *and* non-Gaussian at once — the measured metric is identical across innovation laws (kurtosis $3/1/$

6) and tracks the exact colored curve, staying ≤ 1 at every memory time: the bound survives both non-idealities simultaneously. *(iii) Transient relaxation.* Evaluating the instantaneous metric along the exact relaxation $\Sigma(t) = \Sigma_\infty + e^{Lt}(\Sigma_0 - \Sigma_\infty)e^{L^\top t}$, the bound is never breached during relaxation from the attractor ($\Sigma_0 \to 0$, the rare-event-relevant start; $\max_t = 0.9991$ over $> 4 \times 10^3$ operators) or from above (0.9999): the steady-state ceiling caps the entire physical relaxation. The instantaneous ratio is not pointwise monotone, and arbitrary non-physical initial covariances can pass 1 by ~ 10% — consistent with the bound being, strictly, a steady-state (quasipotential) statement.

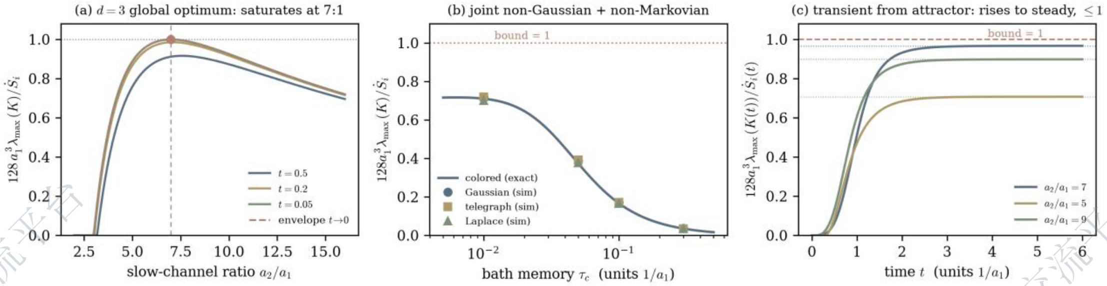

Further verifications. (a) $d = 3$ *full-parameter global optimum: the metric saturates the bound (= 1) exactly at the 7: 1 single channel as the coupling* $t \to 0$*; random configurations stay strictly below. (b) Joint non-Gaussian + non-Markovian bath: three innovation laws give identical metrics, tracking the exact colored curve, ≤ 1.*
(c) Transient relaxation from the attractor: the instantaneous metric rises to the steady-state value and stays ≤ 1.

Discussion and outlook

The picture is sharp. Non-reciprocity turns dissipation into a programmable, directional resource for stabilising rare states, but the conversion obeys a universal speed limit — at most $1/(128a_1^3)$ units of barrier protection per unit of entropy production, fixed by the slowest dissipative mode and attained only by the 7: 1 slow-channel geometry. The result is dual in spirit to the thermodynamic uncertainty relations [3–7]: there, dissipation lower-bounds current fluctuations; here, it upper-bounds a protective response. It also makes the role of the noise explicit. The ceiling is a property of *isotropic* forcing; an equilibrium bath, whose fluctuation–dissipation noise is anisotropic, supports no reversal at all, so the effect lives strictly out of equilibrium. The same mechanism survives quantisation, where the Heisenberg uncertainty principle imposes an additional vacuum floor that holds the physical quantum reversal well below the classical ceiling.

What remains is a sharply posed problem. For $d = 3$ we have reduced the bound to the nonnegativity of two explicit polynomials on a compact box, shown the leading obstruction to vanish, and found the *P1* block strictly sum-of-squares feasible in a floating-point solve; what is missing is an *exact*, boundary-aware certificate — the saturating point is rank-deficient — together with the *q16* block. A closed $d = 3$ proof would still not settle $d \geq 4$, which so far rests on numerics alone and would require a genuine general-d argument. The $7: 1$ selection rule and the $1/\left(128a_1^3\right)$ ceiling should be directly observable in the circuit, photonic, and colloidal platforms analysed here.

Acknowledgments

This work was supported by the General Program of the National Natural Science Foundation of China (Grant No. 5237050287).

Author contributions

Y. Shen conceived the problem, developed the theory and its proofs, carried out all computations and simulations, and wrote the manuscript. X. Sun provided theoretical input. W. Liu provided the laboratory and technical platform. All authors discussed the results and approved the final manuscript.

References

[1] M. I. Freidlin and A. D. Wentzell, *Random Perturbations of Dynamical Systems*, 3rd ed., Grundlehren der mathematischen Wissenschaften Vol. 260 (Springer, Berlin, 2012).

[2] H. Touchette, The large deviation approach to statistical mechanics, *Phys. Rep.* **478**, 1–69 (2009). doi:10.1016/j.physrep.2009.05.002

[3] U. Seifert, Stochastic thermodynamics, fluctuation theorems and molecular machines, *Rep. Prog. Phys.* **75**, 126001 (2012). doi:10.1088/0034-4885/75/12/126001

[4] A. C. Barato and U. Seifert, Thermodynamic uncertainty relation for biomolecular processes, *Phys. Rev. Lett.* **114**, 158101 (2015). doi:10.1103/PhysRevLett.114.158101

[5] T. R. Gingrich, J. M. Horowitz, N. Perunov, and J. L. England, Dissipation bounds all steady-state current fluctuations, *Phys. Rev. Lett.* **116**, 120601 (2016). doi:10.1103/PhysRevLett.116.120601

[6] J. M. Horowitz and T. R. Gingrich, Thermodynamic uncertainty relations constrain non-equilibrium fluctuations, *Nat. Phys.* **16**, 15–20 (2020). doi:10.1038/s41567-019-0702-6

[7] G. T. Landi and M. Paternostro, Irreversible entropy production: from classical to quantum, *Rev. Mod. Phys.* **93**, 035008 (2021). doi:10.1103/RevModPhys.93.035008

[8] M. Lax, Fluctuations from the nonequilibrium steady state, *Rev. Mod. Phys.* **32**, 25–64 (1960). doi:10.1103/RevModPhys.32.25

[9] H. Risken, *The Fokker–Planck Equation: Methods of Solution and Applications*, 2nd ed., Springer Series in Synergetics Vol. 18 (Springer, Berlin, 1989).

[10] M. Fruchart, R. Hanai, P. B. Littlewood, and V. Vitelli, Non-reciprocal phase transitions, *Nature* **592**, 363–369 (2021). doi:10.1038/s41586-021-03375-9

[11] C. Scheibner, A. Souslov, D. Banerjee, P. Surówka, W. T. M. Irvine, and V. Vitelli, Odd elasticity, *Nat. Phys.* **16**, 475–480 (2020). doi:10.1038/s41567-020-0795-y

[12] C. Coulais, D. Sounas, and A. Alù, Static non-reciprocity in mechanical metamaterials, *Nature* **542**, 461–464 (2017). doi:10.1038/nature21044

[13] A. Metelmann and A. A. Clerk, Nonreciprocal photon transmission and amplification via reservoir engineering, *Phys. Rev. X* **5**, 021025 (2015). doi:10.1103/PhysRevX.5.021025

[14] C. Weedbrook, S. Pirandola, R. García-Patrón, N. J. Cerf, T. C. Ralph, J. H. Shapiro, and S. Lloyd, Gaussian quantum information, *Rev. Mod. Phys.* **84**, 621–669 (2012). doi:10.1103/RevModPhys.84.621

[15] A. Serafini, *Quantum Continuous Variables: A Primer of Theoretical Methods* (CRC Press, Boca Raton, 2017).

[16] C. W. Gardiner and P. Zoller, *Quantum Noise*, 3rd ed., Springer Series in Synergetics (Springer, Berlin, 2004).

[17] K. Schmüdgen, *The Moment Problem*, Graduate Texts in Mathematics Vol. 277 (Springer, Cham, 2017). doi:10.1007/978-3-319-64546-9

[18] M. Putinar, Positive polynomials on compact semi-algebraic sets, *Indiana Univ. Math. J.* **42**, 969–984 (1993). doi:10.1512/iumj.1993.42.42045

[19] J. B. Lasserre, Global optimization with polynomials and the problem of moments, *SIAM J. Optim.* **11**, 796–817 (2001). doi:10.1137/S1052623400366802

[20] P. A. Parrilo, Semidefinite programming relaxations for semialgebraic problems, *Math. Program.* **96**, 293–320 (2003). doi:10.1007/s10107-003-0387-5

[21] C. W. Scherer and C. W. J. Hol, Matrix sum-of-squares relaxations for robust semi-definite programs, *Math. Program.* **107**, 189–211 (2006). doi:10.1007/s10107-005-0684-2

[22] H. Peyrl and P. A. Parrilo, Computing sum of squares decompositions with rational coefficients, *Theoret. Comput. Sci.* **409**, 269–281 (2008). doi:10.1016/j.tcs.2008.09.025