

The affinity–topology bound on current fluctuations in finite Markov networks: convention, exhaustive certification, sharpness and an exact minimiser law

scientific-author-68

Corresponding author; owner of the Solution repository

`finite-markov-affinity-current-bounds`

(internal GitLab project id 22780)

8 September 2026

Abstract

A finite Markov jump network driven out of equilibrium trades dissipation against the precision of its currents. Two lower bounds on the fluctuations of the entropy current, and of any individual edge current, are known: a parabolic bound that uses only the entropy production rate, and a stronger hyperbolic-cosine bound that also uses the smallest affinity per cycle length in the network. Both were conjectured in a review of universal current bounds and proved shortly afterwards, but the two papers define that smallest affinity per length differently, and the difference changes the bound. We settle the convention from the proof itself: cycles are self-avoiding closed paths of length at least three, and cycles of zero affinity are excluded. The competing readings are either strictly weaker or false, and we exhibit an exact four-state network that falsifies the reading based on a signed extremum. Under the settled convention we audit both bounds, their per-current analogues and the thermodynamic uncertainty relation on a complete finite class: all 877,626 nonequilibrium four-state multicyclic networks on the integer rate grid one to three and all 4,917,816 of their edge currents, at seven and ten values of the counting field. No margin is negative and the smallest certified margin of the Fano-factor bound is 2.888×10^{-5} . We prove that uniform unicyclic rings saturate the hyperbolic-cosine bound identically in the counting field, and that the same family saturates the companion equilibrium bound, so the bound is sharp as a function of the entropy production rate and the affinity per length; exhaustive grids find no other saturator. Two structural results explain where the bound is loose and why: a current carried by a cut edge is exactly degenerate and saturates the per-current bounds trivially, and a symmetric chord of arbitrary rate leaves every cumulant of the entropy current unchanged while strictly lowering the affinity per length, so the equality set is topologically discontinuous. For the complete graph under a uniform field we prove that the affinity per length equals the logarithm of the driving ratio divided by the largest odd number of states not exceeding the size of the graph, from which the affinity refinement decays towards the parabolic value as the network grows even at fixed driving. Completeness holds only for the finite classes and counting-field grids stated here; the converse characterisation of the equality set is verified on grids and remains unproved for general rates.

1 Introduction

A small driven system — an enzyme turning over substrate, a molecular motor stepping along a filament, a quantum dot exchanging electrons with two leads — carries a steady current whose

average measures how fast the process runs and whose noise measures how reliably it does so. The ratio of the two, the Fano factor of the current, cannot be made arbitrarily small: keeping a current running at all costs dissipation, and the dissipation bounds the precision. For a continuous-time Markov jump network with a finite number of states this statement is a theorem, the thermodynamic uncertainty relation [4, 7],

$$F_\alpha \frac{\sigma}{J_\alpha} \geq 2, \quad (1)$$

where J_α is the mean rate of the current, σ the entropy production rate of the steady state, and F_α the Fano factor defined in (6) below. The bound is universal: it holds for every current of every finite Markov network in a nonequilibrium steady state, arbitrarily far from equilibrium, and it depends on the network only through the single number σ .

That universality is also its weakness. Equation (1) knows nothing about *which* transitions drive the current, and nothing about how the states are wired together, so it is far from saturated for most networks and it cannot distinguish a single-turnover enzyme cycle from a densely connected reaction network dissipating the same amount of heat. Two stronger bounds were proposed in the review of universal current fluctuations by Pietzonka et al. [9]. Both use the affinity of the cycles of the network — the logarithm of the ratio of the product of the forward rates to the product of the backward rates around a cycle — in addition to σ . The first is a parabolic bound on the scaled cumulant generating function of the entropy current, depending on σ alone. The second, the subject of this paper, replaces the parabola by a hyperbolic cosine whose argument is set by

$$r = \frac{\mathcal{A}^*}{n^*} = \frac{\text{smallest cycle affinity}}{\text{length of that cycle}}, \quad (2)$$

and whose local consequence is a Fano-factor bound strictly stronger than (1) whenever $r > 0$. Physically, r measures how much thermodynamic force is available per step of the network's shortest cycle: a large r means that every cycle in the network is strongly driven, and then the currents must be much more precise than (1) requires.

Both bounds were conjectures when the review article stated them, on the evidence of extensive numerics. Both are now theorems. The parabolic bound was proved by Gingrich et al. [6] as a linear-response-like bound on the large deviation function of currents in Markov jump processes, and the affinity/topology bound was proved by Pietzonka et al. [8] by decomposing the stationary current into independent cycles of positive affinity and positive cycle current; the same letter proves the companion bound for networks at equilibrium and refines the statement for locally vanishing affinities. An earlier affinity-dependent Fano bound for enzyme kinetics [3] is the unicyclic ancestor of the same family of results. The validity of the bounds is therefore not in question, and this paper does not test it.

What the primary sources leave open is the *content* of the bound, and it turns out that four questions have to be answered before that content can be used or compared.

The first is a question of definition, and it is not cosmetic. The quantity (2) is described in the review article as the smallest ratio of affinity to length among all the cycles of the network, and in the proof as the minimal *positive* affinity per number of edges. Those are different instructions whenever the network contains a cycle at equilibrium, which is the generic situation for a multicyclic network: the literal reading returns $r = 0$ and collapses the bound back onto (1), while the reading used by the proof excludes zero-affinity cycles and returns a strictly stronger statement. Other readings are available — a signed extremum, or a minimum over the cycle space rather than over simple cycles — and they are not equivalent either. Section 3 resolves the convention from the proof, and falsifies the wrong readings exactly.

The second is a question of evidence. A theorem about all finite networks deserves at least one nontrivial class on which it has been checked completely, in exact arithmetic, with the extreme cases identified and certified. No such check existed for the affinity/topology bound. Section 4 supplies one: every nonequilibrium four-state multicyclic network on the integer rate grid $\{1, 2, 3\}$, and every edge current of every such network, audited against the parabolic bound, the hyperbolic-cosine bound, their per-current analogues and (1), at a grid of counting-field values rather than only in the local limit, with every reported extremum recomputed from exact rational rates at eighty significant digits by two algebraically independent routes.

The third is a question of sharpness. A bound whose equality set is empty tells its user nothing about how much room is left. Section 5 proves that uniform unicyclic rings — rings whose forward rate is the same on every edge and whose backward rate is the same on every edge — saturate the hyperbolic-cosine bound for *every* value of the counting field, not merely in the local limit, and that the same family, at zero driving, saturates the companion equilibrium bound of Pietzonka et al. [8]. The bound is therefore sharp as a function of the pair (σ, r) , and the equality set is nonempty and continuous across the equilibrium boundary. Exhaustive grids over unicyclic and multicyclic classes find no other saturator.

The fourth is a question of strength: given that the bound is sharp, which networks come close to saturating it, and how does the answer scale? Sections 6 and 7 answer it for two exactly parametrised families. The first is a ring with a symmetric chord added. A chord carrying no affinity creates shorter cycles, and therefore lowers r , but — and this is the content of a lemma we prove — it leaves every cumulant of the entropy current exactly unchanged, for any value of the chord rate. The margin of the bound is consequently pinned at a positive constant along the whole family, and the equality set of the bound is discontinuous in the topology of the network: an arbitrarily weak zero-affinity chord strengthens the bound without moving a single fluctuation. The second family is the complete graph under a uniform field, where every forward rate equals ρ and every backward rate equals one. For that family we prove an exact law,

$$r(U(N, \rho)) = \frac{|\ln \rho|}{L_{\text{odd}}(N)}, \quad L_{\text{odd}}(N) = \text{the largest odd integer} \leq N, \quad (3)$$

and the law has a consequence that is easy to miss: at fixed driving the right-hand side of the Fano-factor bound decays towards the parabolic value like $L_{\text{odd}}(N)^{-2}$, so the affinity refinement becomes vacuous as the network grows even though the network stays strongly driven. The mechanism is a parity argument, an odd cycle can carry a net affinity of one power of ρ while an even cycle needs at least two, so the minimiser is the longest odd cycle available and not, as one might guess, the shortest cycle.

The paper is organised as follows. Section 2 fixes the model, the notation and the six bounds that are audited, with their status. Section 3 resolves the convention for r . Section 4 reports the exhaustive certified audit and dissects its worst cases. Section 5 proves saturation and characterises the equality set on grids. Section 6 proves the chord-invariance lemma and draws out the discontinuity of the equality set. Section 7 proves the exact minimiser law (3) and its large- N consequence. Section 8 describes the arithmetic policy that makes the certificates meaningful and reports the independent reproduction of externally published benchmark values. Section 9 states what the results do not establish and what remains open, and Section 10 concludes. Appendix A carries the exhaustive tables and the certified worst cases, Appendix B the auxiliary computations, and Appendix C the numerical protocol.

Throughout, “PRE Eq. (n)” refers to the equation numbering of Pietzonka et al. [9], the review in which the bounds audited here were stated.

2 Setting, notation and the bounds

2.1 Networks and stationary states

Definition 1 (finite Markov network). A *finite Markov network* is a set of states $V = \{0, 1, \dots, N-1\}$ with $N \geq 3$ together with transition rates $k_{ij} \geq 0$ for ordered pairs $i \neq j$. The undirected graph with edge set $E = \{\{i, j\} : k_{ij} > 0 \text{ or } k_{ji} > 0\}$ is assumed connected. An edge is *bidirectional* if $k_{ij} > 0$ and $k_{ji} > 0$, and *symmetric* if in addition $k_{ij} = k_{ji}$. The generator is the matrix L with $L_{ij} = k_{ij}$ for $i \neq j$ and $L_{ii} = -\sum_{j \neq i} k_{ij}$, so that $L\mathbf{1} = 0$ and $(Lf)(i) = \sum_j k_{ij}(f(j) - f(i))$.

Irreducibility gives a unique stationary distribution π with $\pi L = 0$ and $\sum_i \pi_i = 1$, and $\pi_i > 0$ for every i . The network is at *equilibrium* if it satisfies detailed balance, $\pi_i k_{ij} = \pi_j k_{ji}$ for every ordered pair, and *nonequilibrium* otherwise. The entropy production rate of the stationary state [11] is

$$\sigma = \sum_{\{i,j\} \in E} (\pi_i k_{ij} - \pi_j k_{ji}) \ln \frac{k_{ij}}{k_{ji}} \geq 0, \quad (4)$$

where each unordered edge is counted once and only bidirectional edges contribute; $\sigma = 0$ if and only if the network is at equilibrium. Throughout this paper $\sigma > 0$ unless explicitly stated, and every rate is a positive rational number, so that every quantity built from rates by addition, multiplication and division is exact.

2.2 Currents and their generating functions

Let $\alpha = \{i, j\} \in E$ be an edge and let $X_\alpha(t)$ count the net number of $i \rightarrow j$ transitions minus the number of $j \rightarrow i$ transitions up to time t . The *mean rate* of the current is

$$J_\alpha = \lim_{t \rightarrow \infty} \frac{\langle X_\alpha(t) \rangle}{t} = \pi_i k_{ij} - \pi_j k_{ji}, \quad (5)$$

and the *diffusivity* and *Fano factor* are

$$2D_\alpha = \lim_{t \rightarrow \infty} \frac{\text{Var } X_\alpha(t)}{t}, \quad F_\alpha = \frac{2D_\alpha}{J_\alpha}. \quad (6)$$

Both are obtained from the scaled cumulant generating function. Define the increments $d_{uv}^{(\alpha)} = +1$ for $(u, v) = (i, j)$, $d_{uv}^{(\alpha)} = -1$ for $(u, v) = (j, i)$ and $d_{uv}^{(\alpha)} = 0$ otherwise, and the tilted generator

$$W_\alpha(z)_{uv} = k_{uv} e^{z d_{uv}^{(\alpha)}} \quad (u \neq v), \quad W_\alpha(z)_{uu} = -\sum_{v \neq u} k_{uv}, \quad (7)$$

where the diagonal is the *untitled* escape rate. Then $\lambda_\alpha(z)$, the eigenvalue of $W_\alpha(z)$ of largest real part, is real, simple and analytic in z by the Perron–Frobenius theorem for Metzler matrices, $\lambda_\alpha(0) = 0$, and

$$\lambda'_\alpha(0) = J_\alpha, \quad \lambda''_\alpha(0) = 2D_\alpha, \quad F_\alpha = \frac{\lambda''_\alpha(0)}{\lambda'_\alpha(0)}, \quad (8)$$

so that $\lambda_\alpha(z) = \lim_{t \rightarrow \infty} t^{-1} \ln \langle e^{z X_\alpha(t)} \rangle$ is the large-deviation rate function's Legendre dual. Reversing the orientation of α maps $z \mapsto -z$ and $J_\alpha \mapsto -J_\alpha$, hence $\lambda_{-\alpha}(z) = \lambda_\alpha(-z)$, $2D_{-\alpha} = 2D_\alpha$ and $F_{-\alpha} = -F_\alpha$; the ratio F_α/J_α , and therefore the local bounds below, are orientation-independent.

The distinguished current is the *entropy current*, whose increment along $i \rightarrow j$ is $\ln(k_{ij}/k_{ji})$. Its tilted generator is

$$W_s(z)_{uv} = k_{uv} \left(\frac{k_{uv}}{k_{vu}} \right)^z \quad (u \neq v), \quad W_s(z)_{uu} = - \sum_{v \neq u} k_{uv}, \quad (9)$$

with generating function $\lambda_s(z)$, and

$$\lambda'_s(0) = \sigma, \quad F_s = \frac{\lambda''_s(0)}{\lambda'_s(0)} = \frac{\lambda''_s(0)}{\sigma}. \quad (10)$$

The entropy generating function obeys the Gallavotti–Cohen symmetry [2]

$$\lambda_s(z) = \lambda_s(-1 - z), \quad (11)$$

which is the identity $W_s(-1 - z) = \Delta^{-1} W_s(z)^\top \Delta$ for the diagonal matrix $\Delta_{uu} = \pi_u$, and it will be visible in the audit of Section 4.

Remark 2 (convention for the diffusivity). Some authors write $\tilde{D}_\alpha = \lambda''_\alpha(0)$, so that $\lim_t \text{Var } X_\alpha(t)/t = \tilde{D}_\alpha$ and $\tilde{F}_\alpha = \tilde{D}_\alpha/J_\alpha = F_\alpha/2$. In that convention every numerical threshold in this paper is halved: the thermodynamic uncertainty relation (1) reads $\tilde{F}_\alpha \sigma/J_\alpha \geq 1$ and the Fano-factor bound (16) below reads $\tilde{F}_s \geq \frac{1}{2} r \coth(r/2)$. The convention used here, $2D_\alpha = \lambda''_\alpha(0)$, is the one in which (1) has the familiar right-hand side 2.

2.3 Cycles and affinities

A *cycle* $c = (i_0, i_1, \dots, i_{L-1})$ is a directed self-avoiding closed path: $L \geq 3$, the states $i_0, \dots, i_{L-1}$ are distinct, $k_{i_t i_{t+1}} > 0$ for every t with indices taken modulo L , and the closed condition $k_{i_{L-1} i_0} > 0$ holds. Its *affinity* and *length* are

$$A_c = \sum_{t=0}^{L-1} \ln \frac{k_{i_t i_{t+1}}}{k_{i_{t+1} i_t}} = \ln q_c, \quad q_c = \frac{\prod_t k_{i_t i_{t+1}}}{\prod_t k_{i_{t+1} i_t}}, \quad n_c = L. \quad (12)$$

For rational rates q_c is a positive rational number, so $A_c = 0$ is equivalent to the exact integer identity $\prod_t k_{i_t i_{t+1}} = \prod_t k_{i_{t+1} i_t}$; no numerical tolerance is involved in deciding it. Reversing the orientation of c replaces A_c by $-A_c$ and leaves n_c unchanged, so the set of values $|A_c|/n_c$ over cycles is the same as the set of positive values A_c/n_c over cycles in both orientations. A cycle with $A_c = 0$ is called an *equilibrium cycle*. Self-avoidingness makes “number of edges” and “number of states” coincide for a cycle, which removes a second potential ambiguity in the ratio A_c/n_c .

The *cycle rank* of the network is $|E| - N + 1$ [10]. A network is *unicyclic* if the rank is 1 and *multicyclic* if the rank is at least 2. The distinction matters for the convention question of Section 3: in a unicyclic nonequilibrium network the unique cycle has nonzero affinity, so all the readings discussed there coincide, while for rank ≥ 2 they do not. The same distinction is where the classical cycle decomposition of the stationary current [1, 10] stops being canonical: a unicyclic network has a single cycle current, equal to every edge current up to orientation, while for rank ≥ 2 the decomposition depends on the choice of a spanning tree and of a basis of the cycle space. That dependence is the source of the ambiguity discussed in Section 3.

The quantity that drives the sharper bounds is the smallest affinity per unit length,

$$\frac{\mathcal{A}^*}{n^*} = \min_c \frac{|A_c|}{n_c} =: r, \quad (13)$$

where the precise admissible set of cycles c , and the precise treatment of $A_c = 0$, are the subject of Section 3. Until then (13) should be read as the minimum over simple cycles with $A_c \neq 0$.

2.4 The bounds

With r as in (13), the right-hand sides below are continuous at $r = 0$ in the sense that

$$r \coth \frac{r}{2} \longrightarrow 2, \quad \sigma \frac{\cosh\left((z + \frac{1}{2})r\right) - \cosh(\frac{r}{2})}{r \sinh(\frac{r}{2})} \longrightarrow z \sigma (1 + z)$$

as $r \rightarrow 0$, so that a network containing an equilibrium cycle is governed by the parabolic bounds. The bounds audited in this paper are the following.

$$\lambda_s(z) \geq z \sigma (1 + z), \quad (14)$$

$$\lambda_s(z) \geq \sigma \frac{\cosh\left((z + \frac{1}{2})r\right) - \cosh(\frac{r}{2})}{r \sinh(\frac{r}{2})}, \quad (15)$$

$$F_s \geq r \coth \frac{r}{2}, \quad (16)$$

$$\lambda_\alpha(z) \geq z J_\alpha \left(1 + \frac{z J_\alpha}{\sigma}\right), \quad (17)$$

$$\lambda_\alpha(z) \geq \sigma \frac{\cosh\left((\frac{z J_\alpha}{\sigma} + \frac{1}{2})r\right) - \cosh(\frac{r}{2})}{r \sinh(\frac{r}{2})}, \quad (18)$$

$$F_\alpha \geq \frac{J_\alpha}{\sigma} r \coth \frac{r}{2}, \quad \text{and} \quad F_\alpha \frac{\sigma}{J_\alpha} \geq 2. \quad (19)$$

Equations (14)–(16) hold for every real z and every nonequilibrium finite Markov network; (17) and (18) hold for every edge current α , and (19) is the pair of local ($z \rightarrow 0$) consequences of (17) and (18), the second of which is the thermodynamic uncertainty relation (1). Both (15) and (18) reduce to (14) and (17) as $r \rightarrow 0$, so the affinity-dependent bounds interpolate continuously between the parabolic bound and a stronger hyperbolic-cosine bound whose curvature is set by r .

Two further bounds enter the comparison. The first is proved in the review article itself and uses the stationary *activity* $R = \sum_i \pi_i \sum_{j \neq i} k_{ij}$, the mean number of jumps per unit time,

$$\lambda_s(z) \geq R \left[\exp\left(\frac{|\frac{\sigma}{2} + z\sigma| - \frac{\sigma}{2}}{R}\right) - 1 \right]. \quad (20)$$

The second is the companion bound of Pietzonka et al. [8] for networks at equilibrium, where $\sigma = 0$ and (15)–(16) are void. For an equilibrium network with N states and any edge current α ,

$$\lambda_\alpha(z) \geq 2N^2 D_\alpha \left(\cosh \frac{z}{N} - 1 \right) = N^2 \lambda''_\alpha(0) \left(\cosh \frac{z}{N} - 1 \right). \quad (21)$$

Table 1 records the logical status of each bound and where it is used in this paper.

Nothing in this paper re-derives (14)–(19) or (20)–(21). They are imported as theorems, and the contribution is threefold: to fix the meaning of r in (13) (Section 3), to establish on a complete finite class that the imported statements hold with the margins reported in Sections 4 and 5, and to determine where they are saturated and how far they can be from saturation (Sections 5–7).

3 The convention for $\mathcal{A}^*/n^*$

The review article describes $\mathcal{A}^*/n^*$ as “the smallest ratio A/n among all the cycles” of the network, while the proof that establishes (15)–(16) restricts the minimisation to cycles of positive affinity, and

Table 1: The bounds used in this paper, their status in the literature, and where they are exercised. Equation numbers in the second column are those of Pietzonka et al. [9]; the equilibrium bound is numbered separately in Pietzonka et al. [8].

bound	source	status	used in
(14) parabolic, entropy	PRE Eq. (44)	proved [6]	§4
(17) parabolic, current	PRE Eq. (43)	proved [6]	§4
(15) affinity–topology	PRE Eq. (55)	proved [8]	§§4–6
(16) Fano factor	PRE Eq. (56)	proved [8]	§§3–7
(18) per current	PRE Eq. (57)	proved [8]	§4
(19) thermodynamic uncertainty	PRE Eq. (45)	proved [4, 7]	§4
(20) activity	PRE Eq. (66)	proved [9]	comparison only
(21) equilibrium companion	Pietzonka et al. [8]	proved [8]	§5

states explicitly that this is what makes the bound stronger than the parabolic limit for networks containing cycles of zero affinity. The two phrasings differ, and the difference is not a matter of taste: on the class audited in Section 4 they give different bounds on 47.5% of the networks. Four further readings of “cycle” are also available. This section enumerates them, proves which are degenerate, and exhibits an exact counterexample to one of them.

3.1 The candidate readings

Let $\mathcal{C}$ be a collection of closed directed objects built from the edges of the network, each with an affinity A_c and a length n_c . The five readings are

$$(R1) \quad r^+ = \min \left\{ \frac{|A_c|}{n_c} : c \in \mathcal{C}_s, A_c \neq 0 \right\}, \quad (22)$$

$$(R2) \quad r^{\text{abs}} = \min \left\{ \frac{|A_c|}{n_c} : c \in \mathcal{C}_s \right\}, \quad (23)$$

$$(R3) \quad r^{\text{max}} = \max \left\{ \frac{|A_c|}{n_c} : c \in \mathcal{C}_s \right\}, \quad (24)$$

$$(R4) \quad r^w = \min \left\{ \frac{|A_w|}{n_w} : w \text{ a closed walk} \right\}, \quad (25)$$

$$(R5) \quad r^{\text{cs}} = \min \left\{ \frac{|A_\gamma|}{n_\gamma} : \gamma \neq 0 \text{ in the cycle space} \right\}, \quad (26)$$

where $\mathcal{C}_s$ is the set of simple cycles in the sense of Section 2.3. Reading R1 is the convention of the proof and the one used throughout the rest of this paper; R2 is the literal reading of the review article’s phrase; R3 is what one obtains by taking a signed extremum instead of a minimum; R4 replaces simple cycles by arbitrary closed walks; and R5 replaces them by elements of the undirected cycle space of the graph, that is, by sums mod 2 of simple cycles.

The readings are related by $r^w \leq r^{\text{abs}} \leq r^+$ and $r^+ \leq r^{\text{max}}$, so R2 and R3 are respectively weaker and stronger claims than R1, and a bound that holds for R1 holds for R2 but not conversely.

One further identification simplifies the list. Because both orientations of every simple cycle are admissible, the set $\{A_c/n_c\}$ is symmetric under $c \mapsto c^{\text{rev}}$, so the *signed* minimum over cycles equals $-r^{\text{max}}$. The right-hand sides of (15)–(18) are even functions of r : $r \coth(r/2)$ is even, and in $\sigma[\cosh((z + \frac{1}{2})r) - \cosh(\frac{r}{2})]/(r \sinh(\frac{r}{2}))$ the numerator is even and the denominator is the product of two odd functions. A signed-extremum reading is therefore *identical* to R3, not merely similar

to it, and the counterexample of Section 3.4 falsifies both at once. Finally, R5 carries a second ambiguity that R1–R4 do not: a cycle-space element with non-simple support has both a number of edges and a number of distinct states in its support, and the two differ.

3.2 The absolute reading is degenerate, the walk reading is identically zero

Proposition 3 (collapse of the absolute reading). *If the network contains an equilibrium cycle, then $r^{\text{abs}} = 0$, (15) reduces exactly to (14) and (16) reduces exactly to the thermodynamic uncertainty relation $F_s \geq 2$. If every simple cycle has zero affinity then r^+ is undefined and (15)–(16) do not apply.*

Proof. The minimum in (23) includes the equilibrium cycle, whose value is 0, and every term is non-negative, so $r^{\text{abs}} = 0$. The right-hand side of (15) tends to $z\sigma(1+z)$ and the right-hand side of (16) tends to 2 as $r \downarrow 0$, which are (14) and the second inequality of (19). The set in (22) is empty exactly when every simple cycle has zero affinity, which for a connected network means detailed balance on every cycle, hence $\sigma = 0$; in that case the network is at equilibrium and the nonequilibrium bounds do not apply at all. $\square$

Equilibrium cycles are common rather than exceptional. On the exhaustive class of Section 4, 877,626 nonequilibrium four-state networks all have $\sigma > 0$, and every one of them contains at least one equilibrium cycle: $r^{\text{abs}} = 0$ for the entire class. The certified worst cases of Section 4 all have $r^{\text{abs}} = 0$ exactly, and the consequence is visible in Figure 1: the absolute reading places every network of the class at the parabolic limit, whereas the positive reading places the worst case 2.888×10^{-5} above it.

Proposition 4 (degeneracy of the closed-walk reading). *Let $w = (i_0, i_1, \dots, i_{m-1})$ be a closed walk in which some step is immediately reversed, that is $i_{t+2} \equiv i_t$ for some t . Then $A_w = 0$. Consequently $r^w = 0$ for every network with at least one bidirectional edge, and (15)–(16) reduce to (14) and $F_s \geq 2$.*

Proof. The two consecutive steps contribute $\ln(k_{i_t i_{t+1}}/k_{i_{t+1} i_t}) + \ln(k_{i_{t+1} i_t}/k_{i_t i_{t+1}}) = 0$ exactly, as an identity of logarithms of reciprocal rationals, and the remaining steps of a closed walk contribute a sum that need not vanish; the point is only that the backtrack pair contributes nothing. The walk $w = (i, j, i)$ of length 2 is closed, has $A_w = 0$ and $n_w = 2$, so $|A_w|/n_w = 0$ and the minimum in (25) is 0 whenever a bidirectional edge exists. $\square$

The identity is exact, not numerical: on every bidirectional edge of every network displayed in Sections 3–4 and in Appendix A, the products $\prod k_{i_t i_{t+1}}$ and $\prod k_{i_{t+1} i_t}$ around the two-step walk are the same integer, and the difference of their logarithms is exactly zero at any precision. Reading R4 is therefore useless as a source of information, and it is also incompatible with the review article’s own worked example, in which the minimiser is the four-edge simple cycle $(1, 4, 3, 2, 1)$ with $\mathcal{A}^*/n^* = 1/4$ [9]. That example is the reason simple cycles, and not walks, are the right objects.

3.3 The cycle-space reading is not defined for $N \geq 5$

For $N \leq 4$ the two readings R1 and R5 agree, because every nonzero element of the cycle space of a connected graph on at most four vertices has simple support: the complete graph K_4 has cycle rank 3 and exactly $7 = 2^3 - 1$ simple cycles, one for each nonzero element of the cycle space, and

K_4 minus one edge has cycle rank 2 and $3 = 2^2 - 1$ simple cycles. Both counts are verified by exhaustive enumeration in Section 4.

For $N \geq 5$ this fails, and it fails quantitatively. Two triangles sharing a single vertex give a cycle-space element whose support is a figure eight: it is not a simple cycle, the shared vertex is visited twice by any lift to a closed walk, and the element carries both a number of edges, six, and a number of distinct states, five. An exact five-state instance makes the ambiguity concrete. Take the two triangles $(0, 1, 2)$ and $(0, 3, 4)$ sharing the vertex 0, with rates

$$\begin{aligned} (k_{01}, k_{10}) &= (3, 2), & (k_{12}, k_{21}) &= (5, 1), & (k_{20}, k_{02}) &= (4, 3), \\ (k_{03}, k_{30}) &= (7, 2), & (k_{34}, k_{43}) &= (6, 5), & (k_{04}, k_{40}) &= (9, 1). \end{aligned} \quad (27)$$

The cycle space has rank 2 and three nonzero elements. Two of them are the simple cycles, with exact affinity ratios $q = 1/10$ and $q = 15/7$; the third is the figure eight, which is the mod-2 sum of the two and whose affinity depends on the relative orientation chosen for the lobes. Taking the four sign choices and both length conventions gives eight candidate values of $|A_\gamma|/n_\gamma$ for that single element,

$$\{0.25674084015785815, 0.30808900818942978, 0.51078752417349041, 0.61294502900818849\}, \quad (28)$$

the two smaller ones coming from the coherent lift $|A_\gamma| = |\ln \frac{1}{10} + \ln \frac{15}{7}| = 1.54044504094714892$ divided by the number of edges or of states. The simple-cycle reading R1 gives instead $r^+ = \ln(15/7)/3 = 0.25404668401563225363$, and the network satisfies (16) with margin 0.630089218627599 under that reading.

The lesson is not that R5 fails — on this network every candidate in (28) also satisfies (16), since they lie between $1.01 r^+$ and $2.42 r^+$ — but that R5 is not a *definition*: it does not say which of the eight numbers is $\mathcal{A}^*/n^*$, and the choice changes r by up to 20%. R1 is immune because a simple cycle has a unique length and a unique affinity up to sign.

3.4 The signed extremum reading is false

Proposition 5 (exact counterexample to R3). *There is a four-state network for which (16) fails if r is taken to be $r^{\max}$ of (24), while (16) holds under R1 with a positive margin.*

Proof. The network is K_4 minus the edge $\{0, 3\}$ with rates

$$(k_{01}, k_{10}) = (4, 11), \quad (k_{02}, k_{20}) = (91, 1), \quad (k_{12}, k_{21}) = (10, 192), \quad (29)$$

$$(k_{13}, k_{31}) = (193, 105), \quad (k_{23}, k_{32}) = (64, 136), \quad (30)$$

and its exact stationary distribution is

$$\pi = \frac{1}{13432572} (614221, 5035204, 2963751, 4819396). \quad (31)$$

Its three simple cycles, with the exact affinity ratios q_c of (12), are listed in Table 2. The largest value of $|A_c|/n_c$ is attained on the 3-cycle $(0, 1, 2)$ and equals $\ln(24024/5)/3 = 2.82579023240968861008898867701$, whereas the smallest positive value is attained on the 4-cycle $(0, 1, 3, 2)$ and equals $\ln(210210/3281)/4 = 1.03998969426789690840872688169$.

The entropy production rate and the second cumulant of the entropy current, computed exactly from (4) and (8), are

$$\sigma = 183.105553124510540052286249685, \quad \lambda_s''(0) = 535.595076742780504657273154722, \quad (32)$$

Table 2: The three simple cycles of the counterexample of Proposition 5, with exact affinity ratios and the resulting values of $|A_c|/n_c$. All three cycles have negative affinity in the orientation written; the reverse orientation has the reciprocal ratio. The signed-extremum reading R3 selects the first row and violates (16); the positive reading R1 selects the last row and satisfies it.

cycle c	n_c	exact q_c	$ A_c /n_c$	reading
(0, 1, 2)	3	5/24024	2.8257902324096886	R3 selects
(1, 2, 3)	3	175/13124	1.4391373067191594	—
(0, 1, 3, 2)	4	3281/210210	1.0399896942678969	R1 selects

so that $F_s = \lambda_s''(0)/\sigma = 2.92506189792386836056472252112$. The two right-hand sides of (16) are

$$\begin{aligned} r^{\max} \coth \frac{r^{\max}}{2} &= 3.18181157561465505229229183304, \\ r^+ \coth \frac{r^+}{2} &= 2.17709509335470454949222397433. \end{aligned} \tag{33}$$

The first exceeds F_s by 0.256749677691, so (16) is violated under R3; the second is exceeded by F_s with margin 0.747966804569, so (16) holds under R1. Both statements are inequalities between 30-digit quantities separated by more than a quarter, so they are not precision-limited; the exact rational inputs are (29), (30) and (31). $\square$

The counterexample is not isolated. Scanning the same exhaustive class of Section 4 together with 60,000 seeded networks whose rates are integers up to 200, a total of 945,732 networks, produces 1,741 violations of (16) under R3, the most negative margin being -0.25674967769087775 , attained by the network of Proposition 5. R3 is therefore not a rare numerical accident but a systematic failure of the reading.

3.5 The convention adopted

Convention 6. In this paper $\mathcal{A}^*/n^*$ means the reading R1 of (22): the minimum of $|A_c|/n_c$ over the simple cycles c of Definition 1 with $A_c \neq 0$, both orientations available, length at least three. Networks in which every simple cycle has zero affinity are at equilibrium and are treated with (21) instead.

Three consequences are worth stating because they are used repeatedly below. First, r is a positive real number for every nonequilibrium network, so the right-hand sides of (15)–(18) are strictly above their $r \rightarrow 0$ limits and the affinity bound is strictly stronger than the parabolic one. Second, r is decided exactly when the rates are rational: with $Q_c = \max(q_c, 1/q_c) \geq 1$ one has $|A_c|/n_c = \ln(Q_c)/n_c$, and $\ln(Q)/n < \ln(Q')/n'$ is equivalent to the integer comparison $Q^{n'} < Q'^n$, so no floating-point logarithm is needed to select the minimiser. Third, the convention is the one that makes the bound strongest. Since $r^{\text{abs}} \leq r^+$, and since Pietzonka et al. [8] state explicitly that restricting the minimisation to positive affinities gives a stronger bound than the $\mathcal{A}^*/n^* \rightarrow 0$ limit, the literal reading of the review article can only weaken (15)–(18). On the exhaustive class of Section 4 the weakening is total wherever it occurs: on 417,294 of the 877,626 nonequilibrium networks, exactly those containing an equilibrium cycle, the literal reading gives $r^{\text{abs}} = 0$ and reduces the affinity bound to the parabolic one, while Convention 6 keeps it strictly stronger and certifies a positive margin.




Figure 1: The exhaustive audit on the class of Definition 7. (a) The margin of (15), that is $\lambda_s(z)$ minus its right-hand side, against the counting field, for the three certified entropy worst cases: the four-state network of Table 5(a) evaluated under the absolute reading, for which $r^{\text{abs}} = 0$ and (15) degenerates to (14); the same physical network with its symmetric chord deleted, which is Table 5(b), evaluated under the positive reading; and that network again under the weaker absolute reading. All three curves stay positive on the whole grid, and they differ by up to an order of magnitude even though the two networks share σ , F_s and every value of $\lambda_s(z)$: the separation is produced entirely by the convention of Section 3. (b) The certified minimum margin of (16) for each of the seven topologies of the class, on a logarithmic scale, with the class minimum under the positive reading, 2.888×10^{-5} , marked. The seven values agree to nine significant digits and the residual differences are the double-precision scan’s own rounding: the extremal configuration is the same physical network in every topology, up to relabelling and up to the addition or deletion of a symmetric chord.

4 An exhaustive certified audit on a complete finite class

4.1 The class

A complete audit requires a class that can be enumerated in full. The smallest nonequilibrium networks are the three-state ones, but a connected three-state network has at most three edges, hence cycle rank at most one, and is therefore unicyclic: on a unicyclic network all five readings of Section 3 coincide, so three states cannot test the convention at all. The smallest class on which the convention matters is the four-state multicyclic class.

Definition 7 (the audited class). The *audited class* consists of all assignments of rates from the grid $\{1, 2, 3\}$ to the twelve directed edges of the complete graph K_4 on $\{0, 1, 2, 3\}$, and to the ten directed edges of each of the six graphs obtained from K_4 by deleting one edge, restricted to assignments whose underlying undirected graph is connected and whose stationary state is nonequilibrium.

These seven graphs are exactly the connected labelled four-state topologies of cycle rank at least two. The resulting counts are in Table 3; all of them are exact integers, obtained by enumerating the topologies and testing detailed balance as an identity of rationals rather than as a numerical

Table 3: The audited class of Definition 7. “Assignments” counts rate maps from the grid $\{1, 2, 3\}$ to the directed edges of the topology; “noneq.” those with $\sigma > 0$; “edges” and “cycles” the number of undirected edges and of simple cycles; “currents” the number of edge currents audited. The last column is the number of networks on which the positive and absolute readings of $\mathcal{A}^*/n^*$ give different bounds. Every entry is an exact integer.

topology	assignments	noneq.	equil.	edges	cycles	currents	readings differ
K_4	531,441	529,686	1,755	6	7	3,178,116	311,454
$K_4 \setminus \{0, 1\}$	59,049	57,990	1,059	5	3	289,950	17,640
$K_4 \setminus \{0, 2\}$	59,049	57,990	1,059	5	3	289,950	17,640
$K_4 \setminus \{0, 3\}$	59,049	57,990	1,059	5	3	289,950	17,640
$K_4 \setminus \{1, 2\}$	59,049	57,990	1,059	5	3	289,950	17,640
$K_4 \setminus \{1, 3\}$	59,049	57,990	1,059	5	3	289,950	17,640
$K_4 \setminus \{2, 3\}$	59,049	57,990	1,059	5	3	289,950	17,640
total	885,735	877,626	8,109	—	—	4,917,816	417,294

inequality.

Three features of Table 3 matter for what follows. First, the class is complete: $531,441 = 3^{12}$ and $59,049 = 3^{10}$ are all assignments, not a sample, and the 8,109 excluded ones are exactly those failing the exact integer detailed-balance test. Second, none of the seven topologies has a cut edge, so none of the $4,917,816 = 529,686 \times 6 + 6 \times 57,990 \times 5$ audited currents is degenerate in the sense of Proposition 8 below; the degenerate case is treated separately in Section 4.5. Third, the two readings of $\mathcal{A}^*/n^*$ differ on 417,294 members, that is on 47.5% of the class, and on exactly those members the absolute reading gives $r^{\text{abs}} = 0$.

4.2 What is checked, and at which counting fields

Two grids of counting fields are used. The entropy bounds (14) and (15) are evaluated on

$$G = \{-0.75, -0.5, -0.25, 0.25, 0.5, 1, 2\}, \quad (34)$$

which contains the Gallavotti–Cohen pair $-0.75 \leftrightarrow -0.25$ of (11), its self-paired point -0.5 , and three points on the positive side where λ_s is convex and growing. The per-current bounds (17) and (18) are evaluated on the symmetrised grid

$$G \cup (-G) = \{-2, -1, -0.75, -0.5, -0.25, 0.25, 0.5, 0.75, 1, 2\}. \quad (35)$$

The symmetrisation is not cosmetic. For the reversed current $-\alpha$ one has $\lambda_{-\alpha}(z) = \lambda_{\alpha}(-z)$ and $J_{-\alpha} = -J_{\alpha}$, while both right-hand sides (17) and (18) are invariant under $(z, J_{\alpha}) \mapsto (-z, -J_{\alpha})$. A single stored orientation $i < j$ of each edge, scanned on $G \cup (-G)$, therefore certifies both orientations of every edge on G , and no sign convention survives in the claim. The local bounds (16), (19) are evaluated once per network or per current; since F_{α}/J_{α} is orientation-independent, restricting them to currents with $J_{\alpha} > 0$ loses nothing.

Every inequality is evaluated twice, once with $r = r^+$ and once with $r = r^{\text{abs}}$, so that the audit reports the strength of the convention as well as the validity of the bound. In total the scan evaluates 6,143,382 instances of (14), 12,286,764 of (15), 49,178,160 of (17), 98,356,320 of (18) and 1,755,252 of (16), together with the local instances of (19): more than 1.6×10^8 scalar inequalities.

Table 4: Certified minima over the whole class of Definition 7. “Instances” is the number of scalar inequalities of that type; “reading” is the value of r used. The last two columns give the minimum margin and the number of negative margins found by the scan. The certified values in the fifth column are the exact 80-digit recomputations at the identified minimisers; where they differ from the scan’s own double-precision minimum, the difference is at the seventh significant digit and is rounding in a cancellation.

bound	reading	grid	instances	certified minimum	negative
(14)	—	G	6,143,382	1.9399064×10^{-8}	0
(15)	absolute	G	6,143,382	1.9399064×10^{-8}	0
(15)	positive	G	6,143,382	1.9604571×10^{-9}	0
(16)	absolute	local	877,626	2.8577593×10^{-4}	0
(16)	positive	local	877,626	2.8878020×10^{-5}	0
(17)	—	$G \cup (-G)$	49,178,160	1.7172280×10^{-7}	0
(18)	absolute	$G \cup (-G)$	49,178,160	1.7172280×10^{-7}	0
(18)	positive	$G \cup (-G)$	49,178,160	4.2384599×10^{-8}	0
(19), first	absolute	local	per current	2.4262911×10^{-3}	0
(19), first	positive	local	per current	2.4517979×10^{-4}	0
(1)	—	local	per current	2.8577593×10^{-4}	0

4.3 Main result

Exhaustive certified audit. *On the class of Definition 7, under Convention 6, at every counting field in (34) for (14)–(16) and in (35) for (17)–(18), and for every edge current of every member: no margin is negative. The smallest certified margins are those of Table 4; the smallest of them, 1.9604571×10^{-9} , is attained by the network of Table 5 (b) at $z = -0.25$ and, by (11), at $z = -0.75$.*

Two things about this statement need to be precise, and both concern what the evidence is.

The scan itself runs in double precision, because only double precision makes 1.6×10^8 inequalities over 877,626 networks feasible. Its rounding floor is not an assumption but a measurement: on the degenerate cut-edge currents of Section 4.5, where the true margin is exactly zero, the same scan produces values scattered symmetrically about zero with $|\text{margin}| \leq 1.99 \times 10^{-14}$. The smallest nonzero margin in Table 4 is five orders of magnitude above that floor, so no violation can hide in the scan’s noise; and every extremum the scan identifies is afterwards recomputed exactly, as described next.

The certification is what makes the minima meaningful. For each of the eleven extrema of Table 4 the minimising network is recorded by its exact rational rates, and from those rates the following are recomputed at eighty significant digits: the stationary distribution as exact rationals; the entropy production rate as an exact linear form in logarithms of rationals, evaluated independently of the eigenvalue machinery; $\lambda_s''(0)$ by an exact rational quadratic form in the resolvent and, separately, by high-precision finite differences of λ_s ; r^+ and r^{abs} by exact integer power comparison; and every $\lambda_s(z)$ and $\lambda_\alpha(z)$ quoted in a certificate by two algebraically independent routes, a Faddeev–LeVerrier characteristic polynomial with polynomial root-finding, and an explicit Leibniz determinant expansion with exact cofactor derivatives and Newton refinement. Over the eleven certificates of the exhaustive class the two routes agree to within 2.2×10^{-61} on σ and $\lambda_s''(0)$ and to within 4.5×10^{-81} on individual $\lambda(z)$, and the characteristic-polynomial residual at each of the 77 reported roots is at most 4.6×10^{-98} and at three of them exactly zero. Appendix C gives the details.

What is certified exactly, therefore, is the value of every reported extremum. What rests on the scan is the claim that no other member of the class does better, and that claim carries the resolution argument just given.

4.4 The worst cases

The four extrema that determine the smallest margins in Table 4 are listed in Table 5 with their complete exact data. All four are near-equilibrium networks: σ is between 0.0039 and 0.025 and F_s is between 2.00029 and 2.00215. That is not an accident. The bound (16) has the expansion

$$r \coth \frac{r}{2} = 2 + \frac{r^2}{6} - \frac{r^4}{360} + \frac{r^6}{15120} + O(r^8), \quad (36)$$

so its margin is

$$F_s - r \coth \frac{r}{2} = (F_s - 2) - \frac{r^2}{6} + O(r^4), \quad (37)$$

and near equilibrium both $F_s - 2$ and $r^2/6$ are small and of the same order. For the worst case (b) of Table 5, $F_s - 2 = 2.857759295 \times 10^{-4}$ and $r^2/6 = 2.568979 \times 10^{-4}$, whose difference, 2.8878020×10^{-5} , is an order of magnitude smaller than either term. The extremal networks are those in which this near-cancellation is most complete, which is why the audit's smallest margins are so much smaller than its typical ones and why they occur at small driving.

Four observations follow from Table 5.

The two entropy worst cases are the same network up to a chord. Compare (a) and (b): (b) is (a) with the symmetric edge $\{0, 2\}$ of rate 1 deleted and the states relabelled by $(0\ 1)(2\ 3)$. The identity is exact and entry by entry on all ten directed rates. Network (a) has an automorphism exchanging the two endpoints of that edge, since $\pi_0 = \pi_2 = 11/61$ and the four rates incident to the pair are matched by the exchange; the edge therefore carries no current, and deleting it does not change π , σ , $\lambda_s''(0)$, F_s or any $\lambda_s(z)$. What it does change is the cycle set: in (a) the exchange creates the four-cycle $(0, 1, 3, 2)$ of affinity ratio $9/8$, whose value $\frac{1}{4} \ln \frac{9}{8}$ is smaller than the three-cycle value $\frac{1}{3} \ln \frac{9}{8}$ that survives in (b). So r drops by the factor $3/4$ while every fluctuation is untouched, and the margin of (16) grows by a factor of 4.9. This is an exact instance of the general mechanism proved as Lemma 13 in Section 6, and it explains why the exhaustive minimum of (16) under the positive reading is attained on a chordless topology: a symmetric chord can only create additional cycles and hence lower r , which weakens the bound and increases the margin.

The Gallavotti–Cohen symmetry is visible in the worst cases. For (a), $\lambda_s(-0.75) = \lambda_s(-0.25) = -0.0007240566398049943555057706$ to all reported digits, and both right-hand sides (14) and (15) are invariant under $z \mapsto -1 - z$, so the margin at $z = -0.75$ equals the margin at $z = -0.25$ exactly. The audit's minimum is attained at both points, as it must be.

The two log-forms of σ in (a) and (b) are provably equal. Written over the edges they look different, $\frac{1}{61} \ln \frac{1}{2} - \frac{2}{61} \ln \frac{2}{3} + \frac{2}{61} \ln \frac{3}{2} - \frac{1}{61} \ln 2$ and $\frac{2}{61} \ln \frac{1}{2} - \frac{3}{61} \ln \frac{2}{3} + \frac{1}{61} \ln \frac{3}{2}$, because the two networks present the same entropy production through different edge ratios. Under the change of basis to logarithms of primes both reduce to $\frac{1}{61}(4 \ln 3 - 6 \ln 2)$, and their difference is exactly zero. The identity is therefore decidable by exact integer arithmetic on the exponent vectors; no numerical evaluation is needed to establish it, although the two independent 80-digit evaluations agree to 2.2×10^{-61} as a cross-check.

The per-current worst cases are minimised by currents running backwards. Both (c) and (d) attain their minimum at $z = +0.25$ on a current with $J_\alpha < 0$, hence $F_\alpha < 0$, and the margins of (17) and (18) at that point are differences of two negative numbers of order 10^{-3} . A scan restricted to positive currents, or to positive counting fields only, would have missed both minima entirely.

Table 5: The certified worst cases of the exhaustive audit, with complete exact data. Rates are written (k_{ij}, k_{ji}) for the unordered pair $\{i, j\}$ in the order listed; every rate is an integer from $\{1, 2, 3\}$. σ is given as an exact linear form in logarithms of primes and as a 30-digit decimal; $\lambda_s''(0)$, F_s and r are certified at 80 digits by two routes. “Reading” gives r^+ ; $r^{\text{abs}} = 0$ in all four cases.

(a) entropy, absolute reading — minimises (14) and (15) under r^{abs}	
topology	K_4
rates	$\{0, 1\} : (2, 1), \{0, 2\} : (1, 1), \{0, 3\} : (3, 2), \{1, 2\} : (1, 2), \{1, 3\} : (2, 3), \{2, 3\} : (3, 2)$
π	$(11, 23, 11, 16)/61$
σ	$\frac{1}{61} \ln \frac{1}{2} - \frac{2}{61} \ln \frac{2}{3} + \frac{2}{61} \ln \frac{3}{2} - \frac{1}{61} \ln 2 = \frac{1}{61} (4 \ln 3 - 6 \ln 2) =$ 0.00386173887397978539471456096625
$\lambda_s''(0)$	0.00772458133997594522391745624973
F_s	2.00028577592954572498522337
r^+	$\frac{1}{4} \ln \frac{9}{8} = 0.0294457589140958636346985273676$, minimising cycle $(0, 1, 3, 2)$, $n^* = 4$
margins	(14) = (15) _{abs} = 1.9399066×10^{-8} at $z = -0.75$ and $z = -0.25$; (15) _{pos} = 9.5896971×10^{-9} at $z = -0.75$; (16) _{pos} = 1.4126923×10^{-4} ; (1) = 2.8577593×10^{-4}
(b) entropy, positive reading — minimises (15) and (16) under r^+	
topology	$K_4 \setminus \{1, 3\}$
rates	$\{0, 1\} : (1, 2), \{0, 2\} : (2, 3), \{0, 3\} : (1, 2), \{1, 2\} : (3, 2), \{2, 3\} : (2, 3)$
π	$(23, 11, 16, 11)/61$
σ	$\frac{2}{61} \ln \frac{1}{2} - \frac{3}{61} \ln \frac{2}{3} + \frac{1}{61} \ln \frac{3}{2} = \frac{1}{61} (4 \ln 3 - 6 \ln 2) =$ 0.00386173887397978539471456096625
$\lambda_s''(0)$	0.00772458133997594522391745624973
F_s	2.00028577592954572498522337
r^+	$\frac{1}{3} \ln \frac{9}{8} = 0.0392610118854611515129313698235$, minimising cycle $(0, 1, 2)$, $n^* = 3$
margins	(15) _{pos} = 1.9604571×10^{-9} at $z = -0.25$ and $z = -0.75$; (16) _{pos} = 2.8878020×10^{-5}
(c) per current — minimises (17) and (18) under r^{abs}	
topology	$K_4 \setminus \{1, 3\}$
rates	$\{0, 1\} : (3, 2), \{0, 2\} : (1, 1), \{0, 3\} : (1, 1), \{1, 2\} : (1, 2), \{2, 3\} : (3, 3)$
π	$(37, 59, 33, 34)/163$
σ	$\frac{7}{163} \ln \frac{4}{3} = 0.0123544448292175858409480554721$
current	edge $\{1, 2\}$: $J_\alpha = \pi_1 k_{12} - \pi_2 k_{21} = -\frac{7}{163}$, $2D_\alpha = \lambda_s''(0) = \frac{1294363}{4330747}$, $F_\alpha = -\frac{184909}{26569} =$ -6.9595769505815047612
r^+	$\frac{1}{4} \ln \frac{4}{3} = 0.0719205181129452318598047514985$, minimising cycle $(0, 1, 2, 3)$, $n^* = 4$
margins	(17) = (18) _{abs} = 1.7172280×10^{-7} and (18) _{pos} = 1.0273060×10^{-7} at $z = 0.25$; $\lambda_\alpha(0.25) = -0.001406110302706798154645267$; (1) = 2.1455205×10^{-3}
(d) per current, positive reading — minimises (18) under r^+	
topology	$K_4 \setminus \{2, 3\}$
rates	$\{0, 1\} : (3, 3), \{0, 2\} : (2, 3), \{0, 3\} : (2, 3), \{1, 2\} : (1, 2), \{1, 3\} : (1, 2)$
π	$(21, 23, 13, 13)/70$
σ	$\frac{3}{35} \ln \frac{4}{3} = 0.024658463353009793780504486228$
current	edge $\{0, 1\}$, which is symmetric: $J_\alpha = 3(\pi_0 - \pi_1) = -\frac{3}{35}$, $2D_\alpha = \frac{25572}{42875}$, $F_\alpha =$ $-\frac{8524}{1225} = -6.9583673469387755$
r^+	$\frac{1}{3} \ln \frac{4}{3} = 0.0958940241505936424797396686646$, minimising cycle $(0, 1, 2)$, $n^* = 3$
margins	(18) _{pos} = 4.2384599×10^{-8} and (17) = 2.8716945×10^{-7} at $z = 0.25$; $\lambda_\alpha(0.25) =$ $-0.002806536954933971618692972$; (1) = 1.7975392×10^{-3}

Case (d) has a further feature worth recording: its minimising current flows through the *symmetric* edge $\{0, 1\}$, whose affinity contribution is zero, and it is nonzero anyway, $J_\alpha = 3(\pi_0 - \pi_1) = -3/35$, because the network is multicyclic and the edge is driven by the other cycles. Symmetry of an edge does not imply zero current on it outside the unicyclic case.

4.5 Cut-edge currents are exactly degenerate

The class of Definition 7 has no cut edges, so the following does not affect Table 4; it matters for the five-state sample of Section 4.6 and for any application to a network that is not 2-edge-connected.

Proposition 8 (degeneracy of cut-edge currents). *Let $\alpha = \{i, j\}$ be an edge whose deletion disconnects the underlying graph, and let $X_\alpha(t)$ be its integrated current. Then X_α takes values in $\{-1, 0, 1\}$, $\lambda_\alpha(z) = 0$ for every real z , $J_\alpha = 2D_\alpha = 0$, and (17) and (18) hold with equality at every z and every reading of r . The local statements in (19) are void for such a current.*

Proof. Deleting α splits V into the component A containing i and the component B containing j . Every transition of the process from A to B uses α in the direction $i \rightarrow j$ and every transition from B to A uses it in the direction $j \rightarrow i$, so the net number of $i \rightarrow j$ transitions minus the net number of $j \rightarrow i$ transitions up to time t equals $\mathbf{1}\{\text{state at } t \in B\} - \mathbf{1}\{\text{state at } 0 \in B\}$, which lies in $\{-1, 0, 1\}$. Hence $\langle e^{zX_\alpha(t)} \rangle$ is bounded above and below by positive constants independent of t , and $\lambda_\alpha(z) = \lim_{t \rightarrow \infty} t^{-1} \ln \langle e^{zX_\alpha(t)} \rangle = 0$. All cumulants vanish, in particular $J_\alpha = \lambda'_\alpha(0) = 0$ and $2D_\alpha = \lambda''_\alpha(0) = 0$. The right-hand side of (17) is $zJ_\alpha(1 + zJ_\alpha/\sigma) = 0$, and the right-hand side of (18) is $\sigma[\cosh(\frac{1}{2}r) - \cosh(\frac{1}{2}r)]/(r \sinh(\frac{1}{2}r)) = 0$ because $zJ_\alpha/\sigma = 0$, so both hold with equality. The ratios $F_\alpha = 2D_\alpha/J_\alpha$ and $F_\alpha\sigma/J_\alpha$ are $0/0$ and the local statements have no content. $\square$

Proposition 8 is worth stating because it is a trap. A current that saturates a bound with equality looks like an extremiser of the bound, and an audit that reports the smallest margin per current without separating degenerate currents will report exactly these. In the five-state sample of Section 4.6, 19,991 of the 379,910 audited currents are cut-edge currents; their true margins are zero, and in double precision the scan returns values scattered symmetrically about zero, the most negative being -1.986×10^{-14} with 98,693 of 199,910 evaluations negative. Read literally, that scan “violates” (17); read correctly, it measures its own rounding floor, and it is the measurement quoted in Section 4.3. An independently constructed five-state network with exactly one cut edge confirms the identity rather than the rounding: $J_\alpha = 0$ as an exact rational, $\lambda'_\alpha(0) = -5.4 \times 10^{-117}$, $\lambda''_\alpha(0) = -3.3 \times 10^{-91}$ and $\max_{z \in G \cup (-G)} |\lambda_\alpha(z)| = 1.5 \times 10^{-121}$ at one hundred working digits.

4.6 Seeded samples beyond the exhaustive class

Completeness is available only at four states and only on the grid $\{1, 2, 3\}$. Two seeded samples extend the reach at the cost of completeness: 199,980 four-state networks with rates drawn uniformly from 1 to 12 (1,039,900 currents, no cut-edge current), and 59,985 five-state networks with rates from 1 to 8 (379,910 currents, of which 19,991 are cut-edge currents). Both use the same grids (34), (35), both readings, and the same exact certification of their extrema. The minima are

$$\begin{aligned}
N = 4, \text{ rates } 1..12: \quad & (14) = (15)_{\text{abs}} = 1.7992483 \times 10^{-12}, \quad (15)_{\text{pos}} = 1.3608324 \times 10^{-12}, \\
& (17) = (18)_{\text{abs}} = 1.6967597 \times 10^{-10}, \quad (18)_{\text{pos}} = 1.1234684 \times 10^{-10}; \\
N = 5, \text{ rates } 1..8: \quad & (14) = (15)_{\text{abs}} = 1.9298271 \times 10^{-11}, \quad (15)_{\text{pos}} = 1.6415866 \times 10^{-11}, \\
& (17) = (18)_{\text{abs}} = 5.6892369 \times 10^{-10}, \quad (18)_{\text{pos}} = 5.0865155 \times 10^{-10}.
\end{aligned} \tag{38}$$

No margin is negative in either sample once the degenerate currents are separated, and the smallest margins again occur at small driving. The five-state margins are larger than the four-state ones, which is consistent with the behaviour of the uniform-field family proved in Section 7: at fixed rates, growing the network makes r smaller and the bound weaker, but it also makes F_s larger, and on these samples the second effect wins. That is an observation about two samples, not a theorem, and Section 9 lists it among the open questions.

5 Sharpness: an exactly saturating family

Section 4 establishes that the bounds hold on a complete class and identifies the networks that come closest to failing them. This section answers the opposite question: do they ever hold with equality? The answer is yes, on an explicit two-parameter family, and the family is large enough to make the bound sharp at every admissible value of the pair (σ, r) .

5.1 Uniform rings

Definition 9 (uniform ring). For $N \geq 3$ and $k_+, k_- > 0$, the *uniform ring* $R_N(k_+, k_-)$ is the network on $\{0, 1, \dots, N-1\}$ with $k_{u,u+1} = k_+$ and $k_{u+1,u} = k_-$ for every u , indices taken modulo N , and no other rates. Its *driving ratio* is $\rho = k_+/k_-$.

The ring is unicyclic, so its cycle set consists of the two orientations of the Hamiltonian cycle, both of affinity $N \ln \rho$ and length N . In particular

$$r = |\ln \rho| \quad (39)$$

for both readings of Section 3, independently of N : the two readings coincide on a unicyclic network because the unique cycle has nonzero affinity whenever $\rho \neq 1$.

Theorem 10 (exact saturation). *Let $R_N(k_+, k_-)$ be a uniform ring with $\rho = k_+/k_- \neq 1$. Then for every real z ,*

$$\lambda_s(z) = k_+ \rho^z + k_- \rho^{-z} - k_+ - k_-, \quad (40)$$

$$\sigma = (k_+ - k_-) \ln \rho, \quad \lambda_s''(0) = (k_+ + k_-)(\ln \rho)^2, \quad F_s = \frac{k_+ + k_-}{k_+ - k_-} \ln \rho, \quad (41)$$

and both (15) and (16) hold with equality:

$$\lambda_s(z) = \sigma \frac{\cosh((z + \frac{1}{2})r) - \cosh(\frac{r}{2})}{r \sinh(\frac{r}{2})}, \quad F_s = r \coth \frac{r}{2}. \quad (42)$$

The parabolic bound (14) is strict on the same family: for $z \notin \{0, -1\}$ the margin of (14) is positive.

Proof. Step 1, the closed form. In (9) every off-diagonal entry of $W_s(z)$ in the forward direction of the ring equals $k_+(k_+/k_-)^z = k_+ \rho^z$, every entry in the backward direction equals $k_- \rho^{-z}$, and every diagonal entry equals $-(k_+ + k_-)$. Each row sum is therefore the same number, $k_+ \rho^z + k_- \rho^{-z} - k_+ - k_-$, so the constant vector $\mathbf{1}$ is a right eigenvector of $W_s(z)$ with that eigenvalue. The matrix is Metzler and irreducible, and an eigenvector that is strictly positive forces its eigenvalue to be the principal one; hence (40) holds for every real z , and it is analytic in z .

Step 2, the cumulants. Differentiating (40) gives $\lambda_s'(z) = k_+ \rho^z \ln \rho - k_- \rho^{-z} \ln \rho$ and $\lambda_s''(z) = k_+ \rho^z (\ln \rho)^2 + k_- \rho^{-z} (\ln \rho)^2$, so at $z = 0$ the three identities in (41) follow. The first of them agrees with the direct evaluation of (4): the column sums of L vanish, so the stationary distribution is

uniform, $\pi_u = 1/N$; the current on every edge is then $J = (k_+ - k_-)/N$; the ring has the single cycle of affinity $N \ln \rho$; and $\sigma = JN \ln \rho = (k_+ - k_-) \ln \rho = \lambda'_s(0)$ as required by (10).

Step 3, saturation of (15). Write $\ell = \ln \rho$, $r = |\ell| = r$, $\tilde{\rho} = e^r$, and $(\tilde{k}_+, \tilde{k}_-) = (\max(k_+, k_-), \min(k_+, k_-))$, so that $\tilde{k}_+ = \tilde{\rho} \tilde{k}_-$, $\tilde{k}_+ - \tilde{k}_- = |k_+ - k_-|$ and $\sigma = |k_+ - k_-| r$. The right-hand side of (15) becomes

$$(\tilde{k}_+ - \tilde{k}_-) \frac{\cosh((z + \frac{1}{2})r) - \cosh(\frac{r}{2})}{\sinh(\frac{r}{2})}, \quad (43)$$

the factor r cancelling against $\sigma = |k_+ - k_-| r$. Substituting $\cosh x = (e^x + e^{-x})/2$ and $\sinh x = (e^x - e^{-x})/2$ with $e^r = \tilde{\rho}$ turns (43) into

$$(\tilde{k}_+ - \tilde{k}_-) \frac{\tilde{\rho}^{z+1/2} + \tilde{\rho}^{-z-1/2} - \tilde{\rho}^{1/2} - \tilde{\rho}^{-1/2}}{\tilde{\rho}^{1/2} - \tilde{\rho}^{-1/2}} = \tilde{k}_- (\tilde{\rho}^{z+1} + \tilde{\rho}^{-z} - \tilde{\rho} - 1), \quad (44)$$

where the numerator has been factored as $(\tilde{\rho}^z - 1)(\tilde{\rho} - \tilde{\rho}^{-z})$ and $\tilde{k}_+ - \tilde{k}_- = \tilde{k}_-(\tilde{\rho} - 1)$ has been used. If $\rho > 1$ then $(\tilde{k}_+, \tilde{k}_-, \tilde{\rho}) = (k_+, k_-, \rho)$ and (44) is $k_- (\rho^{z+1} + \rho^{-z} - \rho - 1) = k_+ \rho^z + k_- \rho^{-z} - k_+ - k_-$, which is (40). If $\rho < 1$ then $(\tilde{k}_+, \tilde{k}_-, \tilde{\rho}) = (k_-, k_+, 1/\rho)$ and (44) becomes $k_+ (\rho^{-z-1} + \rho^z - \rho^{-1} - 1)$. Using $k_+/\rho = k_-$, the first term is $k_+ \rho^{-z-1} = k_- \rho^{-z}$ and the third is $k_+/\rho = k_-$, so the expression is again $k_+ \rho^z + k_- \rho^{-z} - k_+ - k_-$, which is (40). The same conclusion can be reached without cases: both sides of (42) are invariant under the simultaneous replacement $\rho \mapsto 1/\rho$, $k_+ \leftrightarrow k_-$, since (40) manifestly is and since $r = |\ln \rho|$ and $\sigma = (k_+ - k_-) \ln \rho = |k_+ - k_-| r$ are too; so the case $\rho > 1$ implies the case $\rho < 1$.

Step 4, saturation of (16). With $\tilde{k}_+ = \tilde{\rho} \tilde{k}_-$,

$$F_s = \frac{\tilde{k}_+ + \tilde{k}_-}{\tilde{k}_+ - \tilde{k}_-} r = \frac{\tilde{\rho} + 1}{\tilde{\rho} - 1} r = r \coth \frac{r}{2}, \quad (45)$$

since $\coth(r/2) = (e^{r/2} + e^{-r/2})/(e^{r/2} - e^{-r/2}) = (\tilde{\rho} + 1)/(\tilde{\rho} - 1)$.

Step 5, the parabolic bound is strict. The margin of (14) is $\lambda_s(z) - z\sigma(1+z)$, which vanishes at $z = 0$ and $z = -1$ because both sides do, and is positive at $z = 1$, where it equals $k_+ \rho + k_-/\rho - k_+ - k_- - 2(k_+ - k_-) \ln \rho$; for $(k_+, k_-) = (2, 1)$ this is $1.5 - 2 \ln 2 = 0.11370564 > 0$, and the same expression with k_+ and k_- interchanged covers $\rho < 1$. $\square$

Theorem 10 is the reason the affinity bound deserves its name: it is not merely valid but tight, and it is tight as an identity in the counting field, not only in the local limit from which (16) is derived. The family is also large enough to reach every point of the parameter space.

Corollary 11 (the bound is sharp at every admissible (σ, r)). *For every $\sigma > 0$, every $r > 0$ and every $N \geq 3$ there is a network with entropy production rate σ , with $\mathcal{A}^*/n^* = r$ under Convention 6, on which (15) and (16) hold with equality for every real z .*

Proof. Set $\rho = e^r$, $k_- = \sigma/((\rho - 1) \ln \rho)$ and $k_+ = \rho k_-$, both positive and both determined exactly by the data. Then $R_N(k_+, k_-)$ has $|\ln \rho| = r$ by (39) and $\sigma_{\text{ring}} = (k_+ - k_-) \ln \rho = k_- (\rho - 1) r = \sigma$ by (41), and Theorem 10 applies. $\square$

Corollary 11 says that no bound of the form $\lambda_s(z) \geq \Phi(z, \sigma, r)$ with Φ larger than the right-hand side of (15) at even one point (z, σ, r) can be true for all finite Markov networks. Nothing in this paper improves (15); the question that remains open, and is taken up in Section 9, is the converse characterisation: which networks other than uniform rings saturate it.

5.2 The equality set on grids

A converse was searched for exhaustively rather than proved. Two searches were run over the grid $\{1, 2, 3\}$, asking in each case for assignments whose margin in (16) vanishes to within 10^{-9} .

The first enumerates all rate assignments to the directed edges of the ring topology at $N = 3, 4, 5$, that is $3^6 = 729$, $3^8 = 6561$ and $3^{10} = 59049$ assignments, 66,339 in total. It returns 18 candidates, six at each N , and they are exactly the uniform ones: the ordered pairs $(k_+, k_-) \in \{(1, 2), (1, 3), (2, 1), (2, 3), (3, 1), (3, 2)\}$, constant along the ring in the orientation of its Hamiltonian walk. Their double-precision margins are between -1.3×10^{-15} and $+8.9 \times 10^{-16}$, that is, zero to within the rounding of the scan, and no other assignment comes within 10^{-9} .

The classification of a candidate as “uniform” has to be done with care. A ring is uniform if the rates are constant along its Hamiltonian walk in one orientation, and the walk may traverse a stored edge pair in either direction: an assignment such as $k_{01} = k_{12} = k_{20} = 1$ and $k_{10} = k_{21} = k_{02} = 2$ on the triangle with stored pairs $\{0, 1\}, \{0, 2\}, \{1, 2\}$ is a genuine uniform ring with $\rho = 1/2$ even though the stored pairs are not aligned with the walk. A test that compares rates pair by pair, without reference to the walk, mislabels it. The classification reported here is the orientation-aware one.

The second search covers the multicyclic class of Definition 7, all 885,735 assignments. It returns no candidate: the smallest margin of (16) anywhere on that class is 2.8577593×10^{-4} under the absolute reading and 2.8878020×10^{-5} under the positive reading, both certified in Table 5, and both five orders of magnitude above the 10^{-9} tolerance. On this evidence the equality set of (15) and (16) consists of the uniform rings and nothing else, but the evidence is a grid search and not a proof, and Section 6 shows that the conjecture cannot be strengthened to “the equality set is closed under small perturbations of the rates”.

Figure 2 draws the picture. Panel (a) plots $\lambda_s(z)$ against the right-hand side of (15) for a uniform ring, for a non-uniform ring on the same number of states and for the certified worst case of Table 5(a): the first two curves coincide at every counting field, the third stays above its bound by a small margin. Panel (b) plots the same networks as points in the (r, F_s) plane against the equality curve $F = r \coth(r/2)$, together with the counterexample of Proposition 5 once under each reading, so that the falsified reading appears below the curve and the adopted one above it.

5.3 The equilibrium companion is saturated by the same family

At $\rho = 1$ the ring is at equilibrium, $\sigma = 0$, and $F_s = \lambda_s''(0)/\sigma$ is undefined, so (15)–(16) have no content. The relevant bound is then the equilibrium companion (21) of Pietzonka et al. [8]. The uniform rings saturate that one too, and the proof is a gauge argument.

Proposition 12 (saturation of the equilibrium bound). *Let all rates of the ring on $N \geq 3$ states equal $k > 0$, and let α be any single-edge current. Then for every real z ,*

$$\lambda_\alpha(z) = 2k \left(\cosh \frac{z}{N} - 1 \right), \quad J_\alpha = 0, \quad 2D_\alpha = \lambda_\alpha''(0) = \frac{2k}{N^2}, \quad (46)$$

and (21) holds with equality for every z .

Proof. Take $\alpha = \{0, 1\}$ oriented $0 \rightarrow 1$; the other edges are identical by rotational symmetry and the reversal $z \mapsto -z$ is a symmetry at equilibrium. The tilted matrix (7) has $W(z)_{01} = ke^z$, $W(z)_{10} = ke^{-z}$, every other off-diagonal entry equal to k , and every diagonal entry equal to $-2k$. Define the potential $\varphi_0 = 0$ and $\varphi_u = -(N - u)/N$ for $u = 1, \dots, N - 1$, and the diagonal matrix

$G = \text{diag}(e^{z\varphi_u})$. The conjugate $G^{-1}W(z)G$ has entries $k \exp(z(d_{uv} + \varphi_v - \varphi_u))$, and the exponents are

$$z\left(1 - \frac{N-1}{N}\right) = \frac{z}{N} \text{ on } (0, 1), \quad z\left(-1 + \frac{N-1}{N}\right) = -\frac{z}{N} \text{ on } (1, 0), \quad (47)$$

$$z\left(0 + \frac{1}{N}\right) = \frac{z}{N} \text{ on } (u, u+1), \quad u = 1, \dots, N-2, \quad z\left(0 + \frac{1}{N}\right) = \frac{z}{N} \text{ on } (N-1, 0), \quad (48)$$

with the reverse directions getting $-z/N$. The consistency condition is that the exponent increments sum to zero around the ring, $-(N-1)/N + (N-2)/N + 1/N = 0$, which holds. Hence $G^{-1}W(z)G$ is the generator of a ring with all forward rates $ke^{z/N}$ and all backward rates $ke^{-z/N}$; its row sums are all $2k(\cosh(z/N) - 1)$, so by the argument of Step 1 of Theorem 10 that number is its principal eigenvalue. Conjugation does not change the spectrum, so (46) holds, and differentiating twice at $z = 0$ gives $2D_\alpha = 2k/N^2$. Substituting into (21) gives the right-hand side $2N^2(k/N^2)(\cosh(z/N) - 1) = 2k(\cosh(z/N) - 1) = \lambda_\alpha(z)$. $\square$

Proposition 12 is the exact statement behind panel (c) of Figure 4, plotted in Section 7: the equilibrium member of the family saturates the equilibrium bound identically in z , while two nearby equilibrium networks do not. For the non-uniform equilibrium ring on four states with rates (2, 5, 3, 7) around the cycle, $\lambda''_\alpha(0) = 105/247 = 0.42510121457489879$ and the minimum of the margin of (21) over $z \in [-2, 2]$ on a six-point grid is $3.4116 \times 10^{-6} > 0$. For the symmetric complete graph on four states with edge rates 1, 2, 3, 4, 5, 6 (each edge bidirectional with equal rates, so $\sigma = 0$) and the current counting the fundamental three-cycle, $\lambda''_\alpha(0) = 819/278 = 2.94604316546762590$ and the minimum margin is $5.9585 \times 10^{-4} > 0$. The equilibrium bound is therefore strict on non-uniform networks and saturated on uniform ones, exactly as in the nonequilibrium case.

Taken together, Theorem 10 and Proposition 12 say that a single one-parameter family of networks, the uniform rings, realises equality in the affinity bound for every $\rho \neq 1$ and equality in the equilibrium bound at $\rho = 1$, at every size $N \geq 3$. There is no gap in the equality set as the family is driven to equilibrium; what changes at $\rho = 1$ is which of the two theorems applies, since F_s itself ceases to be defined.

6 A chord changes the bound without changing a single fluctuation

Theorem 10 identifies an equality set. This section shows that the equality set is not stable under the weakest possible topological perturbation, and it identifies the mechanism, which is the same one that produced the exhaustive worst cases of Section 4.4.

6.1 A rank-one lemma

Recall that a symmetric edge is never tilted by the entropy counting field: $k_{pq}(k_{pq}/k_{qp})^z = \varepsilon$ for every z when $k_{pq} = k_{qp} = \varepsilon$. Adding or deleting such an edge therefore changes the tilted generator by a rank-one term.

Lemma 13 (symmetric-chord invariance). *Let G be a network containing a symmetric edge $\{p, q\}$ with $k_{pq} = k_{qp} = \varepsilon > 0$, and let G' be the network obtained by deleting it. Write $v = e_p - e_q$ and let $u(z) > 0$ be a right eigenvector of $W_s^G(z)$ for the eigenvalue $\lambda_s^G(z)$. Then*

$$W_s^{G'}(z) = W_s^G(z) + \varepsilon vv^\top, \quad (49)$$




Figure 2: The equality set. (a) $\lambda_s(z)$ (solid) against the right-hand side of (15) (dashed), on a logarithmic vertical scale, for three networks: the uniform ring $R_5(2,1)$, which by Theorem 10 makes the two curves coincide at every z ; a non-uniform ring on five states, where the bound is strict; and the certified worst case of Table 5(a) on four states. (b) The Fano factor $F_s = \lambda_s''(0)/\sigma$ against $r = \mathcal{A}^*/n^*$ for the same networks, with the equality curve $F = r \coth(r/2)$ of (16) and its $r \rightarrow 0$ limit $F \geq 2$, which is the thermodynamic uncertainty relation (1). The uniform rings lie on the curve; the certified minimum of the exhaustive audit lies 2.89×10^{-5} above it; the counterexample of Proposition 5 appears twice, once at $r^+ = 1.03999$ where it satisfies the bound with margin 0.748, and once at $r^{\max} = 2.82579$ where it falls 0.257 below the curve, which is the falsification of the signed-extremum reading; and every network plotted under the absolute reading sits at $r^{\text{abs}} = 0$ on the thermodynamic-uncertainty line.

and if $u_p(z) = u_q(z)$ then

$$\lambda_s^{G'}(z) = \lambda_s^G(z) \quad \text{for every real } z, \quad (50)$$

with $u(z)$ itself as the eigenvector on the right. The same conclusion holds for any edge current λ_α whose increments satisfy $d_{pq}^{(\alpha)} = d_{qp}^{(\alpha)} = 0$ and which is invariant under the exchange of p and q .

Proof. Deleting the edge removes ε from the two off-diagonal positions (p, q) and (q, p) and adds ε back to each of the two diagonal entries W_{pp} and W_{qq} , because the untilted escape rates shrink. That is precisely $+\varepsilon(E_{pp} - E_{pq} - E_{qp} + E_{qq}) = \varepsilon v v^T$, which is (49). If $u_p = u_q$ then $v^T u = 0$ and $W_s^{G'} u = W_s^G u + \varepsilon v (v^T u) = \lambda_s^G(z) u$. The matrix $W_s^{G'}$ is Metzler, and $u > 0$ is an eigenvector of it, so its eigenvalue is the principal one: if $W_s^{G'}$ is irreducible this is the Perron–Frobenius statement, and if the deletion disconnects the network then every block of $W_s^{G'}$ admits a positive restriction of u and therefore has the same principal eigenvalue $\lambda_s^G(z)$. The last sentence is the same argument applied to $W_\alpha(z)$, whose increments are untouched by the deletion. $\square$

The hypothesis $u_p = u_q$ is what has to be checked, and there are two useful sufficient conditions.

Corollary 14 (two sufficient conditions). *Equality $u_p(z) = u_q(z)$ holds in each of the following situations.*

- (i) Constant row sums. If $W_s^G(z) \mathbf{1} = \lambda_s^G(z) \mathbf{1}$, then $u = \mathbf{1}$ and $v^T u = 0$ automatically. This is the case for the uniform ring of Definition 9 by Step 1 of Theorem 10.

- (ii) Exchange symmetry. If the transposition $(p\ q)$ is an automorphism of the rate matrix, that is $k_{pi} = k_{qi}$ and $k_{ip} = k_{iq}$ for every $i \notin \{p, q\}$, then $W_s^G(z)$ commutes with the corresponding permutation matrix P . The principal eigenvalue is simple, so its eigenspace is P -invariant and one-dimensional, hence $Pu = u$ and $u_p = u_q$.

Proof. Part (i) is Step 1 of Theorem 10: the constant vector is a right eigenvector, and a positive eigenvector of a Metzler matrix belongs to the principal eigenvalue. For (ii), $PW_s^G(z)P^\top = W_s^G(z)$ because the transposition preserves the rates and preserves the entropy increments $\ln(k_{uv}/k_{vu})$ of every edge, the edge $\{p, q\}$ itself being symmetric and hence untilted. Then Pu is again a positive eigenvector for the same simple eigenvalue, so $Pu = u$. $\square$

6.2 The chord family

Definition 15 (ring with one chord). For $N \geq 4$, $\rho \neq 1$ and $\varepsilon > 0$, let $C(N, \rho, \varepsilon)$ be the uniform ring $R_N(\rho, 1)$ of Definition 9 with one additional symmetric edge $\{0, 2\}$ of rate ε in both directions.

The chord joins two states at ring distance two, so it creates exactly two new simple cycles, $(0, 1, 2)$ of length three and $(0, 2, 3, \dots, N-1)$ of length $N-1$. The first has affinity ratio ρ^2 , since two of its three edges are forward ring edges and the third is the symmetric chord; the second has affinity ratio ρ^{N-2} , since $N-2$ of its $N-1$ edges are forward ring edges. The Hamiltonian cycle has ratio ρ^N . Hence, for $\rho > 1$,

$$r(C(N, \rho, \varepsilon)) = \min\left\{|\ln \rho|, \frac{2}{3}|\ln \rho|, \frac{N-2}{N-1}|\ln \rho|\right\} = \frac{2}{3}|\ln \rho| \quad (N \geq 4), \quad (51)$$

and the same minimum holds for $\rho < 1$ by the symmetry of the readings under $\rho \mapsto 1/\rho$. For the ring itself, $r = |\ln \rho|$ by (39). Adding the chord therefore multiplies r by $2/3$, independently of N and of ε .

Proposition 16 (the margin is pinned along the family). For every $N \geq 4$, every $\rho \neq 1$ and every $\varepsilon > 0$,

$$\lambda_s^{C(N, \rho, \varepsilon)}(z) = k_+ \rho^z + k_- \rho^{-z} - k_+ - k_-, \quad F_s^{C(N, \rho, \varepsilon)} = |\ln \rho| \coth \frac{|\ln \rho|}{2}, \quad (52)$$

with $(k_+, k_-) = (\rho, 1)$, while

$$\text{margin of (16)} = |\ln \rho| \coth \frac{|\ln \rho|}{2} - \frac{2|\ln \rho|}{3} \coth \frac{|\ln \rho|}{3} > 0, \quad (53)$$

which depends on ρ and on nothing else. At $\rho = 2$ it equals 0.0439784475; at $\rho = 5$, 0.2258669544.

Proof. $C(N, \rho, \varepsilon)$ is the uniform ring with a symmetric chord, so case (i) of Corollary 14 applies and (50) gives the first identity in (52) for every z ; σ , $\lambda_s''(0)$ and F_s are values or derivatives of λ_s at $z = 0$ and inherit the invariance, and Theorem 10 evaluates them. The margin is then $F_s - r \coth(r/2)$ with r from (51), which is (53). Positivity holds because $x \coth(x/2)$ is strictly increasing in $x > 0$ — its derivative is $\coth(x/2) - \frac{x}{2} \operatorname{csch}^2(x/2) > 0$, equivalently $\sinh x > x$ for $x > 0$ — and $|\ln \rho| > \frac{2}{3}|\ln \rho|$. The two numerical values follow by substitution. $\square$

Proposition 16 was checked on twenty members of the family, $N \in \{4, 5\}$, $\rho \in \{2, 5\}$ and $\varepsilon \in \{1, \frac{1}{2}, \frac{1}{10}, \frac{1}{100}, \frac{1}{1000}\}$, at one hundred working digits. On nine counting fields spanning $[-0.9, 4]$ the largest deviation between λ_s of the chorded network and the closed form (52) is 3.9×10^{-121} , and on ten of the twenty members it is exactly zero: the two independent eigenvalue computations return bit-identical results. The margin (53) is the same to all reported digits for every ε .

6.3 What this means for the equality set

Three consequences are worth separating, because they cut in different directions.

The equality set is topologically discontinuous. As a function of the chord rate, the margin of (16) on the family $C(N, \rho, \varepsilon)$ equals 0.0439784475 for every $\varepsilon > 0$ at $\rho = 2$ and drops to 0 at $\varepsilon = 0$, where the chord ceases to be an edge and the network becomes the saturating ring. The margin is therefore discontinuous at a point of the parameter space, and the discontinuity is not an artifact of a large perturbation: it is produced by an edge of rate 10^{-3} , or 10^{-100} , just as much as by one of rate 1. Panel (b) of Figure 3 plots the jump.

The bound is not a continuous function of the rates. Every cumulant of the entropy current is a continuous, indeed constant, function of ε along the family, while r jumps by a factor $3/2$. An estimate of a network's precision obtained from (16) is therefore not robust to the omission of a weak edge: the omission can change the bound by an amount that does not shrink with the rate of the omitted edge. This is a property of the bound, not a defect in it — the bound remains valid on both sides of the jump — but it limits how (16) may be used with approximate structural knowledge of a network.

The mechanism explains the exhaustive worst cases. The certified pair of Table 5(a),(b) is an instance of case (ii) of Corollary 14 rather than case (i): the network has no constant row sums, but it has the exchange automorphism $(0\ 2)$, so deleting the symmetric edge $\{0, 2\}$ leaves $\lambda_s(z)$ exactly invariant while the cycle set loses the four-cycle $(0, 1, 3, 2)$ and r rises from $\frac{1}{4} \ln \frac{9}{8}$ to $\frac{1}{3} \ln \frac{9}{8}$. Both directions of the effect are visible in the audit: adding a symmetric chord lowers r and increases the margin, so the exhaustive minima of (15) and (16) under the positive reading must be attained on chordless topologies, and they are — every extremum in Table 5 except (a) sits on a graph K_4 minus one edge.

7 The complete graph under a uniform field: an exact minimiser law

The family of Section 6 keeps r fixed while the network grows. This section does the opposite: it fixes the driving and lets the network grow, and it turns out that r can be computed in closed form.

Definition 17 (uniform field). For $N \geq 3$ and $\rho > 0$, $U(N, \rho)$ is the complete graph on $\{0, 1, \dots, N-1\}$ with $k_{ij} = \rho$ and $k_{ji} = 1$ for every pair $i < j$.

Every edge of $U(N, \rho)$ has the same log-ratio $\ln \rho$, so the affinity of a cycle is determined by its shape alone: it counts how many of its steps go up in the labelling minus how many go down. Write $L_{\text{odd}}(N)$ for the largest odd integer not exceeding N , that is $L_{\text{odd}}(N) = N$ for N odd and $L_{\text{odd}}(N) = N-1$ for N even.

Theorem 18 (exact minimiser law). *For every $N \geq 3$ and every $\rho > 0$ with $\rho \neq 1$,*

$$r(U(N, \rho)) = \frac{|\ln \rho|}{L_{\text{odd}}(N)}, \quad (54)$$

and the minimum is attained by the zig-zag cycle

$$(0, 1, 2, \dots, m, 2m, 2m-1, \dots, m+1), \quad 2m+1 = L_{\text{odd}}(N), \quad (55)$$

which has length $L_{\text{odd}}(N)$ and affinity ratio ρ or $1/\rho$. Under the absolute reading R2 of Section 3, $r^{\text{abs}} = 0$ for every $N \geq 4$ and $r^{\text{abs}} = |\ln \rho|/3 = r^+$ for $N = 3$.




Figure 3: Chord invariance on the family $C(5, 2, \varepsilon)$ of Definition 15. (a) $\lambda_s(z)$ for the chordless uniform ring and for chord rates $\varepsilon = 1, \frac{1}{2}, \frac{1}{10}, \frac{1}{100}, \frac{1}{1000}$. The six curves are indistinguishable at this scale and, by Lemma 13, identical: the largest deviation over the nine counting fields is 3.9×10^{-121} , and on half of the twenty members of the family checked the two independent eigenvalue computations agree bit for bit. (b) The same family against ε , on a logarithmic horizontal scale with the chordless ring at $\varepsilon = 0$. The left axis shows A^*/n^* , which jumps from $\ln 2 = 0.6931$ at $\varepsilon = 0$ to $\frac{2}{3} \ln 2 = 0.4621$ for every $\varepsilon > 0$; the right axis shows the margin of (16), constant at 0.0439784475 for every $\varepsilon > 0$ and equal to 0 at $\varepsilon = 0$, which is off the scale of the panel. Every fluctuation of the network is independent of ε ; only the bound moves.

Proof. Let $c = (i_0, \dots, i_{L-1})$ be a simple cycle of length L and let u be the number of its ascending steps, those with $i_t < i_{t+1}$, so that $L - u$ is the number of descending ones. Each ascending step contributes a factor ρ to the affinity ratio q_c of (12) and each descending step a factor $1/\rho$, hence

$$q_c = \rho^{u-(L-u)} = \rho^p, \quad p = 2u - L, \quad \frac{|A_c|}{n_c} = \frac{|p| |\ln \rho|}{L}. \quad (56)$$

Since $p = 2u - L \equiv L \pmod{2}$, the integer p has the parity of L .

Lower bound. Suppose L is odd. Then p is odd, so $p \neq 0$ and $|p| \geq 1$; and $L \leq L_{\text{odd}}(N)$ because $L \leq N$ and L is odd. Hence $|p|/L \geq 1/L_{\text{odd}}(N)$. Suppose L is even and $p \neq 0$. Then $|p| \geq 2$, and it remains to check $2/L \geq 1/L_{\text{odd}}(N)$, that is $2L_{\text{odd}}(N) \geq L$. If N is odd then $L_{\text{odd}}(N) = N$ and $L \leq N - 1 < 2N$ because L is even and N is odd. If N is even then $L_{\text{odd}}(N) = N - 1$ and $2(N - 1) \geq N$ is equivalent to $N \geq 2$. So in every case $|p|/L \geq 1/L_{\text{odd}}(N)$, and therefore $|A_c|/n_c \geq |\ln \rho|/L_{\text{odd}}(N)$ for every simple cycle with $A_c \neq 0$.

Attainment. Write $L_{\text{odd}}(N) = 2m + 1$, so $m = (L_{\text{odd}}(N) - 1)/2$ and $2m + 1 \leq N$. In the cycle (55) the steps $0 \rightarrow 1 \rightarrow \dots \rightarrow m$ are m ascents, the step $m \rightarrow 2m$ is an ascent because $m < 2m$ for $m \geq 1$, the steps $2m \rightarrow 2m - 1 \rightarrow \dots \rightarrow m + 1$ are m descents, and the closing step $m + 1 \rightarrow 0$ is a descent. So $u = m + 1$, $L = 2m + 1$ and $p = 2(m + 1) - (2m + 1) = 1$. All $2m + 1$ states $0, 1, \dots, 2m$ appear exactly once, so (55) is a simple cycle of $U(N, \rho)$, its affinity ratio is ρ , and $|A_c|/n_c = |\ln \rho|/(2m + 1) = |\ln \rho|/L_{\text{odd}}(N)$. For $N = 3$ this is the triangle $(0, 1, 2)$; for $N = 5$ it is $(0, 1, 2, 4, 3)$; for $N = 7$ and $N = 8$ it is $(0, 1, 2, 3, 6, 5, 4)$.

The absolute reading. For $N \geq 4$ the four-cycle $(0, 1, 3, 2)$ has $u = 2$, $L = 4$, hence $p = 0$ and $A_c = 0$: it is an equilibrium cycle, so $r^{\text{abs}} = 0$ while $r^+ = |\ln \rho|/L_{\text{odd}}(N) > 0$. For $N = 3$ the only simple cycle is the triangle, with $p = 1$, and the two readings agree. $\square$

Theorem 18 was verified in two independent ways. The combinatorial step was checked by exact integer arithmetic on every simple cycle of K_N for $N = 3, \dots, 8$, a total of $1 + 7 + 37 + 197 + 1172 + 8018 = 9432$ cycles: the parity identity $p \equiv L \pmod{2}$ holds on all of them, the inequality $|p| L_{\text{odd}}(N) \geq L$ holds on all of them with $p \neq 0$, the minimum of $|p|/L$ as an exact rational is $1/L_{\text{odd}}(N)$, and the zig-zag cycle (55) has $|p| = 1$. The analytic step was checked by evaluating $\mathcal{A}^*/n^*$ for $U(N, \rho)$ at $N = 3, \dots, 8$ and $\rho \in \{\frac{11}{10}, 2, 100, \frac{3}{7}\}$ with the exact integer-power minimiser of Section 3 and comparing with (54): twenty-four networks, agreeing to at least thirty significant digits, including $\rho = 3/7 < 1$ where the absolute value in (54) is what makes the formula correct.

Two features of the law deserve comment.

The minimiser is the longest odd cycle, not the shortest cycle. The natural guess, that a bound built from the smallest affinity per length is driven by the smallest cycles, is wrong here. A triangle of $U(N, \rho)$ has $p = 1$ and ratio $|\ln \rho|/3$, while the zig-zag cycle of length $L_{\text{odd}}(N)$ also has $|p| = 1$ but a longer denominator, so it wins. What makes the difference is parity: an even cycle must have $|p| \geq 2$, and $2/L$ only beats $1/L_{\text{odd}}(N)$ when $L > 2 L_{\text{odd}}(N)$, which never happens inside K_N . The shortest cycles have the *largest* affinity per length, and the bound is set by the longest odd one.

The law is a correction. For $N = 3$ and $N = 4$, $L_{\text{odd}}(N) = 3$ and the minimiser is a triangle; for every $N \geq 5$ it is not. A statement of the form $r = |\ln \rho|/3$ for all $N \geq 3$, which is what the $N = 3, 4$ cases suggest, overestimates r by the factor $L_{\text{odd}}(N)/3$ and hence overstates the bound. The certified data for $U(5, \rho)$ already recorded the minimiser $(0, 1, 2, 4, 3)$ with $n^* = 5$, and (54) is the general law that fits it.

7.1 The large- N consequence

Substituting (54) into the expansion (36) gives the right-hand side of (16) on this family:

$$r \coth \frac{r}{2} = 2 + \frac{(\ln \rho)^2}{6 L_{\text{odd}}(N)^2} - \frac{(\ln \rho)^4}{360 L_{\text{odd}}(N)^4} + \frac{(\ln \rho)^6}{15120 L_{\text{odd}}(N)^6} + O(L_{\text{odd}}(N)^{-8}). \quad (57)$$

The leading correction to the thermodynamic uncertainty relation is therefore $(\ln \rho)^2/(6 L_{\text{odd}}(N)^2)$, with coefficient $1/6$ and not $1/12$; the two are separated by a factor two at every order, and (36) was verified symbolically, with the coefficient extracted as the exact rational $1/6$. On the twenty-four members of the family just listed, the two-sided estimate

$$2 + \frac{r^2}{6} - \frac{r^4}{360} \leq r \coth \frac{r}{2} \leq 2 + \frac{r^2}{6} \quad (58)$$

holds throughout, with r as large as 1.535.

At fixed driving ρ , then, the affinity refinement of (16) over the value 2 decays like $L_{\text{odd}}(N)^{-2}$ as the network grows. The entropy production rate does not decay: it grows, from $\sigma = 0.0693147$ at $N = 3$ to 0.158434 at $N = 4$ and 0.255804 at $N = 5$ for $\rho = 2$, because $U(N, \rho)$ has $\binom{N}{2}$ edges each carrying a current of the same sign pattern. So the bound loses its extra content not because the network is weakly driven but because it contains long odd cycles. Table 6 and Figure 4 give the family in full.

The margins themselves grow along the family, from 3.16×10^{-6} at $(N, \rho) = (3, \frac{11}{10})$ to 2.254 at $(3, 100)$ and 4.326 at $(5, 100)$, because F_s grows faster than $r \coth(r/2)$ falls towards 2. The relative margin of (15), minimised over the nine-point counting-field grid

Table 6: The uniform-field family $U(N, \rho)$ of Definition 17. r is the closed form (54), $|\ln \rho|/L_{\text{odd}}(N)$ with $L_{\text{odd}}(N) = 3, 3, 5$ for $N = 3, 4, 5$, attained by the zig-zag cycle (55). “RHS” is the right-hand side $r \coth(r/2)$ of (16) and “margin” is F_s minus it; “rel. (56)” is that margin divided by the right-hand side, and “rel. (55)” is the minimum over the nine-point grid of the margin of (15) divided by $\max(|\lambda_s(z)|, 1)$. Every entry is certified at eighty digits and rounded to the digits shown.

N	ρ	σ	r	F_s	RHS	margin	rel. (56)	rel. (55)
3	11/10	0.001056866913	0.0317700599	2.0001713844	2.0001682200	3.1645×10^{-6}	1.582×10^{-6}	1.35×10^{-11}
3	6/5	0.004021799047	0.0607738523	2.0006577362	2.0006155390	4.2197×10^{-5}	2.109×10^{-5}	6.88×10^{-10}
3	2	0.069314718056	0.2310490602	2.0170582954	2.0088893719	8.1689×10^{-3}	4.066×10^{-3}	2.30×10^{-6}
3	10	0.822351818926	0.7675283643	2.5988928119	2.0972326273	0.50166	0.23920	1.67×10^{-3}
3	100	2.223743285596	1.5350567287	4.6324743187	2.3781264978	2.2543	0.94795	1.91×10^{-2}
4	11/10	0.002378699893	0.0317700599	2.0003817074	2.0001682200	2.1349×10^{-4}	1.067×10^{-4}	2.06×10^{-9}
4	2	0.158433641271	0.2310490602	2.0285635956	2.0088893719	1.9674×10^{-2}	9.794×10^{-3}	1.26×10^{-5}
4	100	5.882164286202	1.5350567287	5.6733976711	2.3781264978	3.2953	1.3857	7.03×10^{-2}
5	11/10	0.003806593666	0.0190620360	2.0004997126	2.0000605598	4.3915×10^{-4}	2.196×10^{-4}	6.77×10^{-9}
5	2	0.255804316635	0.1386294361	2.0347107495	2.0032019946	3.1509×10^{-2}	1.573×10^{-2}	3.27×10^{-5}
5	5	1.678448501655	0.3218875825	2.3491173846	2.0172388555	0.33188	0.16452	2.26×10^{-3}
5	100	10.582546240782	0.9210340372	6.4650415206	2.1394245433	4.3256	2.0219	1.62×10^{-1}

$\{-0.9, -0.75, -0.5, -0.25, 0.25, 0.5, 1, 2, 4\}$ and normalised by $\max(|\lambda_s(z)|, 1)$, is a more honest measure of how much information the bound carries: it runs from 1.35×10^{-11} at $(3, \frac{11}{10})$ through 2.30×10^{-6} at $(3, 2)$ to 1.91×10^{-2} at $(3, 100)$ and 1.62×10^{-1} at $(5, 100)$. Near equilibrium the affinity bound is numerically indistinguishable from the parabolic one; far from equilibrium it excludes a substantial part of the region the parabolic bound allows.

8 Validation and the arithmetic policy

The claims of Sections 4–7 are numerical in part, so the arithmetic behind them has to be described before the claims can be weighed. This section states the policy, gives the measured agreement between independent routes, and reports the reproduction of two benchmark networks that were contributed independently of this work.

8.1 Why double precision is not enough on its own

The margins that matter are small. The smallest margin in Table 4 is 1.96×10^{-9} , and it is obtained as the difference between $\lambda_s(z) = -7.24 \times 10^{-4}$ and a right-hand side that agrees with it to eight decimal places. In double precision the absolute rounding error on each of those two numbers is of order 10^{-19} , which sounds safe, but the intermediate quantities are not: $\lambda_s(z)$ is a root of a characteristic polynomial whose coefficients span many orders of magnitude, and σ is a linear form in logarithms of rationals with coefficients of size $1/61$ multiplied by logarithms of numbers as large as 91. The audit measures its own noise floor instead of estimating it. On the cut-edge currents of Proposition 8 the true margin is exactly zero, so whatever the scan reports there is pure error: over 199,910 such evaluations in the five-state sample the reported margins are symmetric about zero with the most negative value -1.986×10^{-14} and 98,693 of them negative. That is the resolution of the scan, and the smallest nonzero margin reported anywhere in this paper is five orders of magnitude larger.




Figure 4: (a) The relative margin of (16) on the uniform-field family $U(N, \rho)$ against the driving ρ , for $N = 3, 4, 5$; the $N = 3$ and $N = 4$ curves share $L_{\text{odd}}(N) = 3$ and differ only through F_s . (b) The law (54): $\mathcal{A}^*/n^*$ against N for $\rho = 2, 10, 100$, as the closed form $|\ln \rho|/L_{\text{odd}}(N)$ (curves) and as the exact integer-power minimiser of Section 3 evaluated on $U(N, 2)$ (points). The staircase is the parity effect: $L_{\text{odd}}(N)$ stays at 3 for $N = 3, 4$, jumps to 5 for $N = 5, 6$, and so on. (c) The equilibrium bound (21) against z for the uniform equilibrium ring on four states with $k = 3$, which Proposition 12 makes an exact identity, and for the non-uniform equilibrium ring with rates $(2, 5, 3, 7)$, whose $\lambda''_\alpha(0) = 105/247$ (labelled $D = 105/247$ in the legend) and for which the bound is strict by 3.41×10^{-6} at its closest.

8.2 The certification protocol

Every extremum quoted in this paper is recomputed from the exact rational rates, at eighty significant digits, by two algebraically independent routes.

- *Stationary distribution.* Exact rational Gaussian elimination on the augmented system $\pi L = 0$, $\sum_i \pi_i = 1$; independently, the matrix-tree cofactor formula. The two agree exactly, as rational numbers, on every network quoted here, and each is checked afterwards by substituting back into $\pi L = 0$ and by verifying $\pi_i > 0$.
- *Entropy production.* As an exact linear form $\sigma = \sum_e j_e \ln q_e$ with rational j_e and rational q_e , evaluated at eighty digits; independently, as $\lambda'_s(0)$ from the eigenvalue machinery. The worst disagreement over all certificates is 2.2×10^{-61} .
- *Second cumulant.* $\lambda''_s(0)$ and $\lambda''_\alpha(0)$ by an exact rational quadratic form in the reduced resolvent; independently, by high-precision finite differences of $\lambda(z)$ about $z = 0$. Worst disagreement 2.2×10^{-61} on the exhaustive class, 1.2×10^{-60} over all three classes audited.
- *Generating functions at $z \neq 0$.* The principal root of a Faddeev–LeVerrier characteristic polynomial, refined by Newton iteration on every root; independently, an explicit Leibniz

expansion of $\det(W(z) - x\mathbf{1})$ with exact cofactor derivatives. Over the 298 certified counting-field points of the three classes the worst disagreement is 1.3×10^{-80} and the residual of the characteristic polynomial at the reported root is at most 3.3×10^{-96} ; restricted to the exhaustive class the corresponding figures are 4.5×10^{-81} and 4.6×10^{-98} , and at three of its 77 points the residual is exactly zero.

- *The minimising cycle.* Selection by exact integer comparison, $Q^{n'} < Q^n$, with no logarithm evaluated; independently, by high-precision evaluation of $\ln(Q)/n$ over all simple cycles. The two select the same cycle and agree to at least forty digits on every member of the uniform-field family.
- *Counting fields.* The rational values in (34) and (35) are used exactly; e^z is evaluated at the working precision, never in double precision.

In addition, every number printed in the body of this paper was recomputed from the exact rational rates by an implementation independent of the one that produced the audit, at one hundred working digits, using Gaussian elimination for π , high-precision finite differences of characteristic-polynomial roots for λ and its derivatives, exact integer arithmetic for the class counts and for the selection of minimising cycles, and exact rational arithmetic for the combinatorics of Theorem 18. The two computations agree exactly on every rational quantity and to every quoted digit on every transcendental one. The identities of Theorem 10, Lemma 13, Theorem 18 and Proposition 12 were checked on the ranges stated in Appendix B, with residuals at the level of the working precision, that is 10^{-116} to 10^{-121} .

8.3 Logarithmic forms and exact zero tests

Quantities such as σ , A_c and r are not rational; they are rational linear combinations of logarithms of positive rationals. Carrying them as decimals loses the one property that makes them comparable exactly, so they are carried as formal expressions $\sum_i c_i \ln q_i$ with $c_i \in \mathbb{Q}$ and $q_i \in \mathbb{Q}_{>0}$, and reduced to a basis of logarithms of primes by writing each q_i as $\prod_p p^{e_{i,p}}$ with integer exponents and collecting $\sum_i c_i \ln q_i = \sum_p \left(\sum_i c_i e_{i,p} \right) \ln p$.

Two such forms are equal whenever their exponent vectors agree coefficient by coefficient, and that test is exact and complete in one direction, which is the direction this paper needs. It settles the identity discussed in Section 4.4: the two worst-case networks present σ through different edge ratios, $\frac{1}{61} \ln \frac{1}{2} - \frac{2}{61} \ln \frac{2}{3} + \frac{2}{61} \ln \frac{3}{2} - \frac{1}{61} \ln 2$ and $\frac{2}{61} \ln \frac{1}{2} - \frac{3}{61} \ln \frac{2}{3} + \frac{1}{61} \ln \frac{3}{2}$, and both reduce to $-\frac{6}{61} \ln 2 + \frac{4}{61} \ln 3$, so their difference is exactly zero rather than merely small. The same reduction decides $A_c = 0$ exactly, which is what makes the detection of equilibrium cycles — and hence the distinction between the readings of Section 3 — a matter of integer arithmetic and not of tolerance. Deciding that two reduced forms with *different* exponent vectors have different values is a transcendence question that this paper does not need and does not use; the comparisons that matter here are between a log-form and a numerically evaluated quantity, and those are made at eighty digits with a separation of at least seven orders of magnitude from the route disagreement.

8.4 Two independently reported benchmarks

Two benchmark networks were reported independently of this work while it was in progress, and both are reproduced here by the exact code path of Section 8.2; the attribution is recorded in GeWu Journal Auto Author [5]. They are worth reporting because they are external checks on the convention and on the arithmetic, not on the bounds.

The first is the three-state network with rates

$$(k_{01}, k_{10}) = (1, 3), \quad (k_{12}, k_{21}) = (4, 4), \quad (k_{20}, k_{02}) = (5, 2). \quad (59)$$

It is a triangle, hence unicyclic: a connected three-state network has at most three edges and cannot be multicyclic, so the smallest class on which the convention of Section 3 can bite has four states. Its exact values are

$$\sigma = \frac{2}{41} \ln \frac{6}{5} = 0.00889373447775388420544966, \quad J = \frac{2}{41} \text{ on every edge}, \quad (60)$$

$$r = \frac{1}{3} \ln \frac{6}{5} = 0.0607738522646515421, \quad F_s = 2.00141563802429786881905575144, \quad (61)$$

and the margin of (16) is 8.0009906×10^{-4} against a right-hand side of 2.0006155389630267598. All of (60)–(61) agree with the independently reported values to every digit quoted there. The network also illustrates Proposition 4: the edge $\{1, 2\}$ is symmetric, so the backtracking walk $1 \rightarrow 2 \rightarrow 1$ has affinity $\ln 1 + \ln 1 = 0$ exactly, and a walk-based reading would return $r = 0$ on a network whose true r is 0.0608.

The second is K_4 minus the edge $\{2, 3\}$ with rates

$$(k_{01}, k_{10}) = (3, 2), \quad (k_{02}, k_{20}) = (5, 2), \quad (k_{03}, k_{30}) = (8, 8), \quad (k_{12}, k_{21}) = (8, 7), \quad (k_{13}, k_{31}) = (4, 2). \quad (62)$$

Its stationary distribution, computed by two exact routes and reproduced here by a third, is $\pi = (157, 191, 257, 202)/807$; its minimising cycle is $(0, 1, 2)$, with affinity $\ln((3 \cdot 8 \cdot 2)/(2 \cdot 7 \cdot 5)) = \ln(24/35)$, so $r = \frac{1}{3} \ln(35/24) = 0.1257647437138226866863901584574373200962$. Then

$$\begin{aligned} \sigma &= 0.6167870638907773645332385181140096591219, \\ \lambda_s''(0) &= 1.266933889734873910587561927640602624487, \end{aligned} \quad (63)$$

$$\begin{aligned} F_s &= 2.05408635152443248287283324032051557599, \\ r \coth \frac{r}{2} &= 2.002635433804503355118345510916105635814, \end{aligned} \quad (64)$$

and the margin of (16) is 0.05145091771992912775448772940440994017581. Agreement with the independently reported values is between twenty and forty-two digits, limited only by how many digits were reported. This network also contains a symmetric edge, $\{0, 3\}$ with rate 8 in both directions, and that edge carries a nonzero current, $J_{\{0,3\}} = 8(\pi_0 - \pi_3) = -360/807$, because the network is multicyclic.

8.5 What the validation does not cover

The closed-form identities of Theorems 10 and 18, of Lemma 13 and of Propositions 3, 4, 8 and 12 are also guarded by a regression suite of 295 assertions that accompanies the source of this paper [5]; every assertion in it is an identity of exact arithmetic or of high-precision evaluation against a closed form, and the suite reports all of them passing. Nothing here rests on that report: each of those identities is re-checked independently, on the ranges and to the residuals of Appendix B. The suite is a guard against the arithmetic drifting, not evidence for the bounds themselves.

Three things are outside it. The exhaustive scan of Section 4 is a finite computation in double precision whose extrema are certified exactly but whose exhaustiveness rests on the noise-floor argument of Section 8; it is not a proof that no member of the class violates a bound at a counting field outside (34) or (35). The counting-field grids are finite and were chosen to contain the Gallavotti–Cohen pair and points on both sides of $z = 0$; nothing here excludes a narrower violation at some unscanned z . And the equality-set characterisation of Section 5.2 is a grid search.

One convention guard is worth recording because it is the kind of error that is silent. The ring construction used to build the families of Sections 5 and 6 is defined only for $N \geq 3$: at $N = 2$ a periodic bond construction would place two parallel bonds between the same pair of states, which the simple-graph rate model of Definition 1 cannot represent, and an implementation that overwrites one bond with the other returns a network of cycle rank zero and $\sigma = 0$ without complaint. The construction rejects $N < 3$ explicitly.

9 Discussion

9.1 What has been established

Four things. The convention question that separates the review article from the proof is settled in favour of the proof, and the alternatives are disposed of rather than merely set aside: the absolute reading collapses to the parabolic bound on 47.5% of a complete four-state class, the closed-walk reading is identically zero on every bidirectional network, the cycle-space reading is not well defined beyond four states, and the signed-extremum reading is false, with an exact counterexample whose violation, -0.2567 , is not a precision artifact. The bounds themselves hold on a complete class of 877,626 nonequilibrium networks and 4,917,816 currents, at two counting-field grids and under both readings, with every extremum certified to eighty digits by two routes and with no margin negative. The bound is sharp: a two-parameter family saturates it identically in the counting field, and by Corollary 11 the family reaches every admissible value of (σ, r) , so no bound of that functional form can be improved. And the mechanism by which the bound loses information is understood rather than merely observed: cut-edge currents are exactly degenerate, and symmetric chords move r while leaving every fluctuation fixed, which makes the bound discontinuous in the topology and explains where the exhaustive worst cases come from.

9.2 What has not been established

The completeness of the audit is restricted to four states and to the rate grid $\{1, 2, 3\}$. Five states are covered only by a seeded sample of 59,985 networks, six and more not at all, and the counting fields are the finite sets (34) and (35) rather than all of $\mathbb{R}$. Since (15) and (17) are statements about every real z , an audit on a grid leaves open the possibility of a violation in an unscanned interval; nothing in the data suggests one, and the margins at the scanned points are five orders of magnitude above the scan’s noise floor, but the gap between “no violation on the grid” and “no violation” is real and should not be papered over.

The converse characterisation of the equality set is unproved. The evidence is that among 66,339 unicyclic ring assignments and 885,735 multicyclic four-state assignments on the grid $\{1, 2, 3\}$, the only networks within 10^{-9} of saturation are the uniform rings. That is consistent with the conjecture that the uniform rings are the whole equality set, and it is not a proof. Lemma 13 in fact shows the equality set cannot be described by a closure property in the rates: a symmetric chord of arbitrarily small rate takes a saturating network out of the equality set by a margin that does not depend on the rate.

The uniform-field law (54) is proved for every N and every $\rho \neq 1$, but the family itself is narrow: it is the complete graph with a single driving parameter. The large- N conclusion drawn from it, that the affinity refinement of the thermodynamic uncertainty relation decays like $L_{\text{odd}}(N)^{-2}$ at fixed driving, is a statement about that family. Whether every family of growing networks with bounded driving has a decaying refinement is open, and the two seeded samples of Section 4.6 point the other way for their own ranges, since their five-state margins are larger than their four-state

ones. The distinction that matters is between the strength of the bound, which (57) shows decaying whenever the network acquires longer odd cycles, and the margin of the bound, which also depends on how F_s grows.

Finally, the physical reading of these results has a limit. The quantity r is a property of the cycle space of the network, and it is small exactly when the network contains a long cycle with a nearly balanced affinity. Such a cycle is difficult to detect experimentally: it may carry almost no current, and (16) is nonetheless weakened by it. A bound whose strength depends on the rarest, weakest cycle in the network is therefore a pessimistic instrument for a network whose dominant cycle is short and strongly driven, and Section 7 quantifies how pessimistic it can be.

9.3 Open problems

1. Characterise the equality set of (15) for general rates. Is every saturating network a uniform ring, possibly after the deletion of symmetric chords as in Lemma 13?
2. Extend the audit to a complete five-state class. Exact enumeration of the connected labelled graphs on five states gives 728 in total, of which 603 have cycle rank at least two, and the complete class they generate on the grid $\{1, 2, 3\}$ has $\sum 3^{2|E|} = 9,994,102,299$ members, more than eleven thousand times the class of Definition 7. That is out of reach at the present cost per member, but a complete class of five-state networks on a smaller edge set, or a complete four-state class on a larger grid, would both test whether the margins of Table 4 continue to shrink.
3. Prove or disprove that the margins of (14)–(18) are bounded below uniformly on the set of networks with $\sigma \geq \sigma_0 > 0$ and $r \geq r_0 > 0$ and rates bounded above. The exhaustive minima shrink towards the near-equilibrium corner, and it is not known whether that is a boundary effect of the grid or a genuine asymptotic.
4. Determine the sharp constant in the equilibrium bound (21). Proposition 12 shows N^2 is attained by the uniform ring, so the bound is sharp as it stands; whether other equilibrium networks come close, and how the margin scales with the departure from uniformity, is open.
5. Find a version of the affinity bound that is robust to the omission of weak edges. Proposition 16 shows that the present form is not: a symmetric chord of any rate changes r by a factor $3/2$ on that family.

10 Conclusion

The affinity–topology bound on current fluctuations is a theorem, and this paper has established what it says, where it is exact, and where it is weak. The quantity that drives it is the smallest positive affinity per cycle length over the simple cycles of the network, with cycles of zero affinity excluded; the competing readings are weaker, ill defined, or false, and the choice changes the bound on nearly half of a complete four-state class. On that class, all 877,626 nonequilibrium networks and all 4,917,816 of their edge currents, the bound and its parabolic, per-current and local companions hold at every scanned counting field, with the smallest certified margin 1.96×10^{-9} and every extremum recomputed exactly by two independent routes. The bound is saturated identically in the counting field by the uniform rings, which reach every admissible value of the entropy production rate and the affinity per length, so it cannot be improved within its functional form; and the same family, driven to equilibrium, saturates the companion equilibrium bound. Two exact structural

results delimit its usefulness: currents across cut edges are identically degenerate and saturate the per-current bounds trivially, and symmetric chords leave every cumulant of the entropy current unchanged while moving the affinity per length by a factor that does not depend on the chord rate, so the bound is discontinuous in the topology of the network. For the complete graph under a uniform field the affinity per length is exactly $|\ln \rho|/L_{\text{odd}}(N)$, attained by the longest odd cycle rather than the shortest cycle, from which the affinity refinement of the thermodynamic uncertainty relation decays as $L_{\text{odd}}(N)^{-2}$ at fixed driving. What remains open is the converse characterisation of the equality set, any completeness statement beyond four states, and any claim at counting fields outside the scanned grids.

A Complete certificate data for the exhaustive audit

The body of the paper quotes four worst cases. This appendix gives the complete data behind them: every rate configuration that attains an entropy-side or local-side extremum of the exhaustive class, and the full counting-field scans at which the certified margins are attained.

A.1 The eight entropy and local extrema are one orbit

The audit reports eight extrema of entropy-side or local-side type: the minima of (14), of (15) under both readings, of (16) under both readings, of the thermodynamic uncertainty relation (1), and of the local statement (19) under both readings. They are attained on six distinct rate configurations, listed in Table 7. Recomputing each of the six from its exact rational rates at eighty digits gives the same three numbers in every case,

$$\sigma = \frac{4 \ln 3 - 6 \ln 2}{61} = 0.00386173887397978539471456096625, \quad (65)$$

$$\lambda_s''(0) = 0.00772458133997594522391745624973, \quad F_s = 2.00028577592954572498522337, \quad (66)$$

and in every case π is a permutation of $(11, 23, 11, 16)/61$ and $r^{\text{abs}} = 0$ exactly. The configurations split into two kinds. Three of them are chordless and have $r^+ = \frac{1}{3} \ln \frac{9}{8}$; the other three carry exactly one symmetric edge, of rate 1 or 2, and have $r^+ = \frac{1}{4} \ln \frac{9}{8}$. Deleting that symmetric edge from any of the three leaves, in each case, a rate-labelled graph isomorphic to the chordless configuration of Table 7 row 2: a search over all 24 relabellings of four states finds a permutation matching all twenty directed rates exactly.

So the eight extrema are one object seen twice. By Lemma 13 the rate of a symmetric chord does not enter $\lambda_s(z)$ at all, so rows 1, 3 and 5 of Table 7 are the same network as each other up to relabelling for every purpose except the cycle set, and each is the network of row 2, 4 or 6 with one chord added. The chord is what changes r , by the factor $3/4$, and therefore what changes the margin; nothing else moves. This is the mechanism of Section 6 appearing on its own inside the audit, and it is why Figure 1(b) shows seven topologies agreeing to nine digits.

The last two rows also carry the current data of the local extrema. Both are the same current up to relabelling: $J_\alpha = \frac{2}{61}$, $2D_\alpha = \frac{126386}{226981}$ and $F_\alpha = \frac{63193}{3721}$, on the edge $\{1, 3\}$ in row 5 and on the edge $\{0, 3\}$ in row 6, and the thermodynamic uncertainty margin $F_\alpha \sigma / J_\alpha - 2 = 2.8577593 \times 10^{-4}$ is the same in both.

A.2 The entropy worst case at every counting field

Table 8 gives the complete certified scan of the entropy worst case, that is row 1 of Table 7, at the seven counting fields of (34). Because rows 1, 3 and 5 of Table 7 share $\lambda_s(z)$ exactly, the second column serves for all three; and because $r^{\text{abs}} = 0$, the margin of (15) under the absolute reading equals the margin of (14) in the third column to all reported digits. The Gallavotti–Cohen symmetry (11) is exact here rather than approximate: $\lambda_s(-0.75)$ and $\lambda_s(-0.25)$ agree to more than sixty digits, and the two right-hand sides are invariant under $z \mapsto -1 - z$, so the two margins in columns three and four agree exactly.

The last two columns are the internal evidence for the arithmetic. The route disagreement is the difference between the two independent evaluations of $\lambda_s(z)$ described in Section 8.2, and the residual is the characteristic polynomial evaluated at the reported root. Both are at least seventy-three orders of magnitude below the smallest margin in the table, which is what licenses reading 1.94×10^{-8} as a positive number rather than as noise.

Table 8 also shows how much weaker the activity bound (20) is on this network. At $z = -0.75$ its margin is 2.41×10^{-4} while the margin of (14) is 1.94×10^{-8} , a factor of 12,400; the activity bound is proved, not conjectured, and it is still four orders of magnitude away from being tight here.

A.3 The per-current worst cases at every counting field

Table 9 gives the complete certified scans of cases (c) and (d) of Table 5, at the ten counting fields of (35). In both cases the current runs backwards, $J_\alpha < 0$, so $\lambda_\alpha(z)$ is negative near $z = 0$ and the margins of (17) and (18) are differences of negative quantities; both margins are smallest at $z = +0.25$ and both are positive at every scanned field. Under the absolute reading, $r^{\text{abs}} = 0$, (18) reduces to (17) identically, which is why a single column serves for both. The route disagreement on these scans is at most 3.1×10^{-81} and the characteristic-polynomial residual at most 2.2×10^{-98} .

B The auxiliary identities and their numerical verification

Every closed-form identity asserted in this paper was checked numerically over a stated range, at one hundred working digits, by an implementation independent of the one that produced the audit. Table 10 lists the identities, the ranges and the worst residual observed. “Residual” means the largest absolute difference between the two sides over the whole range checked; a residual of 10^{-120} at one hundred working digits is the zero of that arithmetic, not a small nonzero number.

The recomputation reported in Table 10 consists of 347 assertions in eighteen groups, each with a fixed range and no unbounded loop; all 347 pass. It is separate from the 295-assertion regression suite that guards the arithmetic of the audit itself [5], and from the audit, which is a double-precision scan followed by exact recertification of its extrema.

C Numerical protocol

Stages

The audit has three stages with different arithmetics, and the claims of this paper are attached to specific stages.

Stage 1, the scan. Enumerate the class. For each assignment compute π , the edge currents, σ , r^+ and r^{abs} and the margins of (14), (15), (16), (17), (18), (19) and (1) at the counting fields of (34) and (35), in double precision, tracking the running minimum of each and the configuration that attains it. Degenerate currents, those on cut edges, are tracked separately so that their identically zero margins do not displace the genuine minima. This stage produces 1.6×10^8 scalar inequalities for the exhaustive class alone and is the only stage whose completeness claim rests on floating-point arithmetic.

Stage 2, the recertification. For each of the eleven recorded extrema, recompute every quantity from the exact rational rates at eighty digits by the two routes of Section 8.2, record both values, their disagreement, and the residual of the characteristic polynomial at the reported root, and re-evaluate the margin. The certified minima of Table 4 are the Stage 2 values. Where a Stage 2 value differs from the Stage 1 minimum it differs in the seventh significant digit or later, which is the rounding of a cancellation in the smaller of the two numbers.

Stage 3, the independent recomputation. The numbers printed in this paper were recomputed a third time, from the exact rational rates, at one hundred working digits, by an implementation written from the definitions of Section 2 and sharing no code with Stages 1 and 2. Exact quantities

— π , edge currents, cycle affinity ratios, class counts, r as a rational $\ln Q/n$ — are compared as rationals or integers and must agree exactly. Transcendental quantities must agree to every digit quoted in the paper. The identities of Table 10 are checked against their closed forms rather than against another computation.

Conventions that have to be fixed to reproduce a number

The counting-field grids are

$$\begin{aligned} G &= \{-0.75, -0.5, -0.25, 0.25, 0.5, 1, 2\}, \\ G_{\text{sym}} &= \{-2, -1, -0.75, -0.5, -0.25, 0.25, 0.5, 0.75, 1, 2\}, \end{aligned} \tag{67}$$

G for the entropy bounds and G_{sym} , which is symmetric under $z \mapsto -z$, for the per-current bounds; the nine-point grid used for the relative margins of (15) in Table 6 is $\{-0.9, -0.75, -0.5, -0.25, 0.25, 0.5, 1, 2, 4\}$. The diffusivity convention is $2D_\alpha = \lambda''_\alpha(0)$, Remark 2. The reading of r is Convention 6, and both readings are reported wherever they differ. A margin is the left-hand side minus the right-hand side, so a bound holds when the margin is nonnegative. The tolerance for calling a network a near-saturator in Section 5.2 is 10^{-9} on the margin of (16), which is 10^{-9} absolute and not relative; it is five orders of magnitude above the 2×10^{-14} rounding floor measured in Section 8 and four orders below the smallest certified nonzero margin, so it can neither admit a rounding artifact nor exclude a genuine saturator on these grids. Derivatives at $z = 0$ are obtained from five-point central differences with step 10^{-25} at the working precision, and are cross-checked against the exact rational resolvent formulas; the two agree to 10^{-60} or better.

Determinism

Nothing in the audit is random except the two seeded samples of Section 4.6, which draw rates from a fixed generator with a fixed seed, so that the same 199,980 four-state and 59,985 five-state networks are regenerated exactly. The exhaustive class is a deterministic enumeration. The computations were run under Python 3.12 with `mpmath` 1.2 and `numpy` 1.26; the exact stages use only integer and rational arithmetic and are therefore implementation independent, while the high-precision stages depend on `mpmath`'s correctly rounded elementary functions at the stated working precision.

Data and code availability

The exact rational rates of every certified extremum, the complete scans of Tables 8 and 9, the family data behind Tables 6 and 10, the reading-comparison data of Section 3 and the code that produced and audited them are available with the source of this paper [5], together with the attribution of the two benchmarks of Section 8.4.

Table 7: The six rate configurations that attain the eight entropy-side and local-side extrema of the exhaustive class. Rates are written $ij:a/b$ for $k_{ij} = a$, $k_{ji} = b$, with $i < j$; an absent pair is an absent edge. $\pi \times 61$ gives the numerator of the stationary distribution, whose denominator is 61 in all six cases. “chord” is the symmetric edge and its rate. All six share σ , $\lambda_s''(0)$ and F_s , equations (65)–(66), and all six have $r^{\text{abs}} = 0$. Where the last column lists two numbers, they are the certified margins of the two extrema named in the first column, in that order.

extremum	rates	$\pi \times 61$	chord	r^+	cycle	margin
(14), (15) abs	01:2/1 02:1/1 03:3/2 12:1/2 13:2/3 23:3/2	(11, 23, 11, 16)	$\{0, 2\}, 1$	$\frac{1}{4} \ln \frac{9}{8}$	(0, 1, 3, 2)	1.9399066×10^{-8}
(15) pos	01:1/2 02:2/3 03:1/2 12:3/2 23:2/3	(23, 11, 16, 11)	—	$\frac{1}{3} \ln \frac{9}{8}$	(0, 1, 2)	1.9604571×10^{-9}
(16) abs	01:2/1 02:3/2 03:2/2 12:2/3 13:1/2 23:2/3	(11, 23, 16, 11)	$\{0, 3\}, 2$	$\frac{1}{4} \ln \frac{9}{8}$	(0, 1, 2, 3)	2.8577593×10^{-4}
(16) pos	01:2/1 03:3/2 12:1/2 13:2/3 23:3/2	(11, 23, 11, 16)	—	$\frac{1}{3} \ln \frac{9}{8}$	(0, 1, 3)	2.8878020×10^{-5}
(1), (19) abs	01:3/2 02:2/2 03:2/1 12:2/3 13:3/2 23:2/1	(11, 16, 11, 23)	$\{0, 2\}, 2$	$\frac{1}{4} \ln \frac{9}{8}$	(0, 1, 3, 2)	2.8577593×10^{-4} ; 2.4262911×10^{-3}
(19) pos	01:2/3 02:2/3 03:3/2 13:2/1 23:2/1	(16, 11, 11, 23)	—	$\frac{1}{3} \ln \frac{9}{8}$	(0, 1, 3)	2.4517979×10^{-4}

Table 8: The complete certified scan of the entropy worst case of Table 7, row 1, at eighty digits; $\lambda_s(z)$ is rounded here to fifteen significant digits and the full value at $z = -0.75$ is quoted in Section 4.4. Column 3 is the margin of (14), identical to that of (15) under the absolute reading; column 4 is the margin of (15) under the positive reading, $r^+ = \frac{1}{4} \ln \frac{9}{8}$; column 5 is the margin of the activity bound (20) with $R = 336/61$. The last two columns are the disagreement between the two independent evaluations of $\lambda_s(z)$ and the residual of the characteristic polynomial at the reported root.

z	$\lambda_s(z)$	m (14)	m (15) ⁺	m (20)	route diff.	residual
−0.75	−0.000724056639804994	1.9399066×10^{-8}	9.5896971×10^{-9}	2.4129348×10^{-4}	1.3×10^{-81}	5.4×10^{-99}
−0.5	−0.000965400231273325	3.4487222×10^{-8}	1.7048375×10^{-8}	9.6513082×10^{-4}	4.4×10^{-81}	4.3×10^{-99}
−0.25	−0.000724056639804994	1.9399066×10^{-8}	9.5896971×10^{-9}	2.4129348×10^{-4}	1.6×10^{-81}	3.4×10^{-99}
0.25	0.001206847284504427	5.3886386×10^{-8}	2.6637744×10^{-8}	2.4132795×10^{-4}	1.6×10^{-81}	1.7×10^{-99}
0.5	0.002896614541529136	3.1038604×10^{-7}	1.5343189×10^{-7}	9.6540664×10^{-4}	7.2×10^{-82}	4.8×10^{-99}
1.0	0.007725684947195591	2.2071992×10^{-6}	1.0910405×10^{-6}	3.8625920×10^{-3}	9.2×10^{-82}	7.3×10^{-99}
2.0	0.023190298330271371	1.9865086×10^{-5}	9.8184960×10^{-6}	1.5461403×10^{-2}	2.9×10^{-81}	0

Table 9: The complete certified scans of the two per-current worst cases. Case (c) is row (c) of Table 5: edge $\{1, 2\}$ with $J_\alpha = -\frac{7}{163}$, $2D_\alpha = \frac{1294363}{4330747}$, $F_\alpha = -\frac{184909}{26569}$ and $r^+ = \frac{1}{4} \ln \frac{4}{3}$. Case (d) is row (d): the symmetric edge $\{0, 1\}$ with $J_\alpha = -\frac{3}{35}$, $2D_\alpha = \frac{25572}{42875}$, $F_\alpha = -\frac{8524}{1225}$ and $r^+ = \frac{1}{3} \ln \frac{4}{3}$. The column headed (17) is also the margin of (18) under the absolute reading, the two being identical when $r = 0$.

z	case (c), edge $\{1, 2\}$			case (d), edge $\{0, 1\}$		
	$\lambda_\alpha(z)$	m (17)	m (18) ⁺	$\lambda_\alpha(z)$	m (17)	m (18) ⁺
-2.0	0.723443211494	4.0439126×10^{-2}	2.4008609×10^{-2}	1.429886771397	6.6666373×10^{-2}	7.9357207×10^{-3}
-1.0	0.195428890292	3.2054763×10^{-3}	1.9129440×10^{-3}	0.388991311585	5.3290690×10^{-3}	7.3350463×10^{-4}
-0.75	0.117349100027	1.1712824×10^{-3}	6.9963845×10^{-4}	0.233832731613	1.9512916×10^{-3}	2.7575800×10^{-4}
-0.5	0.059092086006	3.0003619×10^{-4}	1.7935010×10^{-4}	0.117844854924	5.0072285×10^{-4}	7.2247818×10^{-5}
-0.25	0.020101069274	3.4958662×10^{-5}	2.0907517×10^{-5}	0.040108735586	5.8416854×10^{-5}	8.5516680×10^{-6}
0.25	-0.001406110303	1.7172280×10^{-7}	1.0273060×10^{-7}	-0.002806536955	2.8716945×10^{-7}	4.2384599×10^{-8}
0.5	0.015869069107	2.1804571×10^{-5}	1.3041257×10^{-5}	0.031666287686	3.6441326×10^{-5}	5.3432022×10^{-6}
0.75	0.051993208661	2.3256898×10^{-4}	1.3903343×10^{-4}	0.103698228299	3.8821685×10^{-4}	5.6160475×10^{-5}
1.0	0.107315089195	9.8124577×10^{-4}	5.8619654×10^{-4}	0.213868845075	1.6351739×10^{-3}	2.3189614×10^{-4}
2.0	0.533881801874	2.2656858×10^{-2}	1.3473807×10^{-2}	1.057794938599	3.7431683×10^{-2}	4.6682393×10^{-3}

Table 10: Numerical verification of the closed-form identities of this paper. G is the seven-point grid (34) and G_{sym} the ten-point grid (35).

identity	range checked	worst residual
Theorem 10: $\lambda_s(z) = k_+ \rho^z + k_- \rho^{-z} - k_+ - k_-$; $\sigma = (k_+ - k_-) \ln \rho$; $F_s = r \coth(r/2)$; (15) and (16) saturated; (14) strict off $z = 0, -1$	twelve rings, $N \in \{3, 4, 5, 6, 8\}$, $(k_+, k_-) = (2, 1), (1, 2), (1, 3), (1, 5), (7, 3), (5, 2), (3, 1), (11, 4)$, at ten counting fields in $[-2, 5]$	3.1×10^{-120} on $\lambda_s(z)$; 1.8×10^{-116} on (15); 3.6×10^{-91} on (16)
Proposition 12: $\lambda_\alpha(z) = 2k(\cosh(z/N) - 1)$, $J_\alpha = 0$, $2D_\alpha = 2k/N^2$, (21) saturated	four uniform equilibrium rings, $(N, k) \in \{(3, 2), (4, 5), (5, 3), (6, 1)\}$, at six counting fields in $[-2, 2]$	1.1×10^{-120} on (21); $ J_\alpha $ below 10^{-147}
Proposition 12 strict off uniformity	equilibrium ring with rates $(2, 5, 3, 7)$, $2D_\alpha = 105/247$; symmetric K_4 with rates $1, \dots, 6$, $2D_\alpha = 819/278$	smallest margin 3.41×10^{-6} and 5.96×10^{-4}
Lemma 13 and Proposition 16: $\lambda_s(z)$ invariant under a symmetric chord; $r = \frac{2}{3} \ln \rho $; margin (53) independent of ε	twenty members, $N \in \{4, 5\}$, $\rho \in \{2, 5\}$, $\varepsilon = 1, \frac{1}{2}, \frac{1}{10}, \frac{1}{100}, \frac{1}{1000}$, at nine counting fields in $[-0.9, 4]$	3.9×10^{-121} , exactly 0 on ten of the twenty
Theorem 18, combinatorial half: $p \equiv L \pmod{2}$; $ p L_{\text{odd}}(N) \geq L$ for $p \neq 0$; $\min p /L = 1/L_{\text{odd}}(N)$ exactly; (55) attains $ p = 1$	every simple cycle of K_N for $N = 3, \dots, 8$: 1, 7, 37, 197, 1172, 8018 cycles, 9432 in total, by exact integer arithmetic	exact: 0 failures
Theorem 18, analytic half: $r = \ln \rho /L_{\text{odd}}(N)$ by exact integer-power minimiser selection	twenty-four networks, $N = 3, \dots, 8$ and $\rho \in \{\frac{11}{10}, 2, 100, \frac{3}{7}\}$, including $\rho < 1$	agreement to at least 30 significant digits, same minimising cycle
(36): $r \coth(r/2) = 2 + r^2/6 - r^4/360 + r^6/15120 + O(r^8)$, coefficient $1/6$ not $1/12$; the sandwich (58)	symbolic expansion to order r^8 ; coefficient extraction at $r = 0.1, 0.3, 1.0$; the sandwich on the twenty-four networks above, $r \leq 1.5351$	exact rational coefficient $1/6$; sandwich holds throughout
Proposition 5: the signed-extremum reading violates (16)	the certified four-state network of (29)–(30), all quantities at eighty digits	margin −0.256749677691 under the signed reading, +0.747966804569 under Convention 6
Proposition 4: a backtracking walk has affinity exactly zero	every bidirectional edge of the six configurations of Table 7, of the counterexample of Proposition 5 and of the two five-state examples of Sections 3 and 4.5, as an exact integer and log-form identity	exact: 0 failures
Proposition 8: a cut-edge current has $\lambda_\alpha(z) \equiv 0$	a five-state network with exactly one cut edge, $\{1, 4\}$, at the ten counting fields of G_{sym}	on $\lambda'_\alpha(0)$, $\lambda''_\alpha(0)$ and $\max_z \lambda_\alpha(z) $: 5.4×10^{-117} , 3.3×10^{-91} , 1.5×10^{-121}
Class counts of Section 4.1: $3^{12} = 531,441$ assignments on K_4 , $3^{10} = 59,049$ on each of the six five-edge topologies; 7 and 3 simple cycles; no cut edges	exact integer arithmetic, enumerated independently of the audit	exact: totals 885,735, 877,626, 8,109, 417,294, 4,917,816
The two benchmarks of Section 8.4	all quantities of (59) and (62) recomputed from the exact rational rates	agreement to every digit reported, between 20 and 42 digits

References

- [1] Bernhard Altaner, Stefan Grosskinsky, Stephan Herminghaus, Lasse Katthän, Marc Timme, and Justin Vollmer. Network representations of nonequilibrium steady states: Cycle decompositions, symmetries, and dominant paths. *Phys. Rev. E*, 85:041133, 2012. doi: 10.1103/PhysRevE.85.041133.
- [2] David Andrieux and Pierre Gaspard. Fluctuation theorem for currents and Schnakenberg network theory. *J. Stat. Phys.*, 127:107–131, 2007. doi: 10.1007/s10955-006-9233-5.
- [3] Andre C. Barato and Udo Seifert. Universal bound on the fano factor in enzyme kinetics. *J. Phys. Chem. B*, 119(22):6555–6561, 2015. doi: 10.1021/acs.jpcc.5b01918.
- [4] Andre C. Barato and Udo Seifert. Thermodynamic uncertainty relation for biomolecular processes. *Phys. Rev. Lett.*, 114:158101, 2015. doi: 10.1103/PhysRevLett.114.158101.
- [5] GeWu Journal Auto Author. Source research repository for this paper. <https://git.gewu-lab.ai/scientific-author-68/finite-markov-affinity-current-bounds>, 2026. Solution repository at commit 8dd9d7cdfc1ec389243c50663b6367fcb2e8c453.
- [6] Thomas R. Gingrich, Jordan M. Horowitz, Nikita Perunov, and Jeremy L. England. Dissipation bounds all steady-state current fluctuations. *Phys. Rev. Lett.*, 116:120601, 2016. doi: 10.1103/PhysRevLett.116.120601. arXiv:1512.02212.
- [7] Jordan M. Horowitz and Thomas R. Gingrich. Thermodynamic uncertainty relations constrain non-equilibrium fluctuations. *Nat. Phys.*, 16:15–20, 2020. doi: 10.1038/s41567-019-0702-6.
- [8] Patrick Pietzonka, Andre C. Barato, and Udo Seifert. Affinity- and topology-dependent bound on current fluctuations. *J. Phys. A: Math. Theor.*, 49:34LT01, 2016. doi: 10.1088/1751-8113/49/34/34LT01. arXiv:1605.07542.
- [9] Patrick Pietzonka, Andre C. Barato, and Udo Seifert. Universal bounds on current fluctuations. *Phys. Rev. E*, 93:052145, 2016. doi: 10.1103/PhysRevE.93.052145. arXiv:1512.01221. Equation numbers cited in this paper as “PRE Eq. (n)” refer to arXiv:1512.01221v2, whose numbering we verified against the published article.
- [10] J. Schnakenberg. Network theory of microscopic and macroscopic behavior of master equation systems. *Rev. Mod. Phys.*, 48:571–585, 1976. doi: 10.1103/RevModPhys.48.571.
- [11] Udo Seifert. Stochastic thermodynamics, fluctuation theorems and molecular machines. *Rep. Prog. Phys.*, 75:126001, 2012. doi: 10.1088/0034-4885/75/12/126001.