RESEARCH MANUSCRIPT

A sharp universal skewness bound for unicyclic currents

Dev D. Goyal

29 October 2026; revised 1 November 2026

Verification scope. The Lean target is the finite-matrix derivative interval and existence of a local branch. The identification with long-time Markov-current cumulants and the sharpness limits are conventional mathematical proofs. Fresh reproduction status is recorded in the companion verification manifest.

Abstract

For a finite continuous-time Markov chain on a bidirectional ring, let C_k denote the kth cumulant per unit time of the net crossing count of a fixed edge. Ptaszyński [1] proposed the dimension-independent interval $-1/16 \le C_3/C_1 \le 1$ for such winding currents. We prove it for every ring with at least three states, arbitrary strictly positive rates and nonzero affinity; both constants are sharp. The proof combines the known positive midpoint spectrum of the tilted generator with a finite-ring proof of the thermodynamic uncertainty relation. The lower bound then reduces to a two-variable polynomial inequality, certified by exact positive integer coefficients on three rational intervals, which are listed in full. The conventional two-state and one-state cases are included as supplements. A Lean 4 development verifies the interval for the actual finite-rate tilted matrix, including existence of its local eigenvalue branch; the identification with long-time stochastic cumulants and the sharpness of the constants rest on conventional proofs given here.

Keywords. Markov currents; unicyclic networks; cumulant bounds; thermodynamic uncertainty; exact positivity certificate; Lean 4.

1 The problem and the contribution

Ptaszyński [1], Eq. (21), proposes the universal interval

$$-\frac{1}{16} \le \frac{C_3}{C_1} \le 1 \tag{1.1}$$

for winding currents in unicyclic Markov networks. The nearest-neighbour ring and the counted edge are specified in Eqs. (72)–(74) of that paper. Here C_k is a cumulant per unit time; the ratio C_3/C_1 is the normalized third cumulant used in that work, rather than the usual standardized third central moment.

The original paper numerically optimizes over rates and studies an irreversible family with one slow transition. That family already shows that the proposed universal constants are optimal. A proof of the interval for arbitrary bidirectional rates was not supplied. The result here proves (1.1) for every ring with $n \ge 3$ states, strictly positive forward and reverse rates, and nonzero affinity. The separate kurtosis, joint-cumulant, and sharper fixed-size optimization claims of [1] are outside the theorem.

The determinant separation and the real-rooted midpoint shift used below are established prior work: they appear in Appendix C of Pietzonka, Barato and Seifert [5]. The thermodynamic uncertainty relation is also established prior work [2]. The contribution is the sharp third/first cumulant interval obtained by combining these ingredients with a new exact scalar inequality. Finite constructions for the stationary law and correction terms derive these ingredients from the rate assumptions in the formal development.

The proposed contribution is a proof of (1.1). Bibliographic priority is not established by the formal proof. First-passage skewness bounds [3, 4] concern different normalizations and do not directly supply this result. The novelty assessment remains subject to further literature comparison.

The Lean development checks the finite-matrix form of the theorem (Theorem 2.2); the identification with stochastic cumulants (Section 3) and sharpness (Section 8) are conventional proofs. Section 9 states the formal scope exactly.

Earlier public work (note added 1 November 2026). On 31 October 2026 a check of public problem records found that the Omega Institute's public repository records a proof of Ptaszyński's numerically conjectured minimum skewness for *unidirectional* cycles (issue #22468), committed on 26 October 2026, before the first version of this manuscript [6]. It gives the exact minimum for every number of states and $-1/16$ as the infimum over all state counts. That is a different question: the theorem here concerns bidirectional rings with positive forward and reverse rates and proves both sides of (1.1). The two results agree on the lower constant $-1/16$. No originality or priority claim is made for the unidirectional case.

2 Statement and conventions

Read indices modulo n. Write $u_i > 0$ for the rate $i \to i+1$ and $v_i > 0$ for the rate $i+1 \to i$. Set

$$U = \prod_{i=0}^{n-1} u_i, \qquad V = \prod_{i=0}^{n-1} v_i, \qquad A = \log(U/V). \tag{2.1}$$

Let N_t count forward minus reverse crossings of any fixed edge during the time interval $[0, t]$. For an arbitrary initial probability distribution, define

$$C_k = \lim_{t\to\infty} \frac{\mathrm{cum}_k(N_t)}{t}. \tag{2.2}$$

Theorem 2.1 (Sharp universal current bound)**.** *Let $n \geq 3$, let every u_i, v_i be strictly positive, and suppose $U \neq V$. The limits in (2.2) exist for $k = 1, 2, 3$, $C_1 \neq 0$, and*

$$\boxed{-\frac{1}{16} \leq \frac{C_3}{C_1} \leq 1.}$$

Both constants are sharp over strictly positive rates and arbitrary ring sizes.

Let W_s be the column generator tilted at the counted edge: multiply its forward off-diagonal rate by e^s and its reverse rate by e^{-s}, leaving the diagonal escape rates unchanged. In particular, $(W_0)_{i+1,i} = u_i$, $(W_0)_{i,i+1} = v_i$, and $(W_0)_{ii} = -(u_i + v_{i-1})$.

Theorem 2.2 (Finite-matrix form)**.** *Under the rate assumptions of Theorem 2.1, there exists a local C^3 real branch λ satisfying*

$$\lambda(0) = 0, \qquad \det(\lambda(s) I - W_s) = 0$$

for all s sufficiently close to zero. Every such local branch satisfies

$$\lambda'(0) \neq 0, \qquad -\frac{1}{16} \leq \frac{\lambda'''(0)}{\lambda'(0)} \leq 1.$$

The formal target is Theorem 2.2, together with its branch existence assertion. The stationary distribution, correction terms, midpoint spectral positivity, uncertainty inequality, and characteristic coefficient formulas are derived from the rates; they are not extra hypotheses. Section 3 gives the ordinary mathematical identification $C_k = \lambda^{(k)}(0)$.

Remark 2.3 (Scope). At $U = V$ the mean winding current vanishes, so the literal quotient C_3/C_1 is undefined. The theorem is about the ring specified above and does not add dangling trees or additional cycles. Two states require two distinguished transition channels; the one-state counted-event convention is a difference of Poisson counts. These boundary models are treated in Appendix B.

3 The stochastic and finite-matrix formulations

For a chain with state X_t, define the marked generating vector by $F_i(t,s) = \mathbb{E}[e^{sN_t}\mathbf{1}_{\{X_t=i\}}]$. Conditioning over the next infinitesimal jump gives

$$\partial_t F(t,s) = W_s F(t,s), \qquad F(0,s) = p_{\text{initial}}.$$

The holding rates remain unchanged, while a marked forward or reverse jump contributes e^s or e^{-s}. Consequently,

$$\mathbb{E}[e^{sN_t}] = \mathbf{1}^T e^{tW_s} p_{\text{initial}}. \tag{3.1}$$

For real s, W_s is an irreducible Metzler matrix. Adding a sufficiently large multiple of the identity produces an irreducible nonnegative matrix with positive diagonal. Perron–Frobenius theory gives a simple real eigenvalue $\lambda(s)$ strictly larger in real part than the other eigenvalues, with positive left and right eigenvectors. At $s = 0$ it is zero. Simple-root perturbation gives an analytic continuation near zero. Let $\Pi(s)$ be its analytic spectral projection. In a sufficiently small neighbourhood of zero there is a decomposition

$$e^{tW_s} = e^{t\lambda(s)}\Pi(s) + R_t(s), \qquad a(s) = \mathbf{1}^T\Pi(s)p_{\text{initial}}.$$

At zero, $\Pi(0) = \pi\mathbf{1}^T$, so $a(0) = 1$ for every initial probability distribution. The complementary spectrum has a uniform gap. Analytic spectral projection and finite-dimensional semigroup estimates therefore bound the first three s-derivatives of $e^{-t\lambda(s)}R_t(s)$ by a polynomial in t times a uniformly decaying exponential. The same estimates follow by analytic continuation to a smaller complex neighbourhood and Cauchy bounds. Since $a(s)$ is analytic and nonzero near zero, (3.1) yields

$$\log \mathbb{E}[e^{sN_t}] = t\lambda(s) + O(1), \tag{3.2}$$

where the remainder and its first three s-derivatives are bounded locally. This derivative-uniform estimate justifies differentiating (3.2) at zero and dividing by t, which proves existence of the limits and

$$C_k = \lambda^{(k)}(0), \qquad k = 1, 2, 3. \tag{3.3}$$

This is the standard path-to-matrix identification used for finite-state current cumulants. It is supplied here as a mathematical argument. The Lean development does not construct a probability space of continuous-time sample paths.

4 The positive midpoint spectrum

For now orient the ring so that $U > V$. Define

$$P(z) = \det(zI - W_0), \qquad \delta = (\sqrt{U} - \sqrt{V})^2 > 0. \tag{4.1}$$

Lemma 4.1 (Counted-edge separation)**.** *For every real s, z,*

$$\det(zI - W_s) = P(z) - U(e^s - 1) - V(e^{-s} - 1). \tag{4.2}$$

Proof. Expand the determinant over permutations. A supported permutation of the cycle consists of fixed vertices, adjacent transpositions, and possibly one full forward or reverse cycle. The two counting factors cancel in an adjacent transposition. Only the full cycles retain a counting factor; their signed contributions are $-Ue^s$ and $-Ve^{-s}$. Subtracting the untilted determinant proves the identity. □

Define the cyclic gradient matrix B by

$$(Bx)_i = \sqrt{u_i}\, x_i - \sqrt{v_i}\, x_{i+1}, \qquad H = B^T B. \tag{4.3}$$

Its diagonal is $u_i + v_{i-1}$, and each neighbouring off-diagonal is $-\sqrt{u_i v_i}$. The determinant of B is $\sqrt{U} - \sqrt{V}$, so B is invertible, H is positive definite, and $\det H = \delta$. The same permutation expansion gives

$$\det(zI + H) = P(z) + \delta. \tag{4.4}$$

Equivalently, H is minus a diagonal conjugate of $W_{-A/2}$. These midpoint facts are the known reduction of [5, Appendix C]; they do not require reversibility of W_0.

Let $\theta_i > 0$ be the eigenvalues of H, and put

$$T = \sum_i \frac{1}{\theta_i}, \qquad x_i = \frac{1}{T\theta_i}, \qquad p_2 = \sum_i x_i^2, \qquad p_3 = \sum_i x_i^3. \tag{4.5}$$

The vector x is a positive probability vector. Also define

$$c = \frac{\sqrt{U} + \sqrt{V}}{\sqrt{U} - \sqrt{V}} = \coth(A/4) > 1. \tag{4.6}$$

If a_j denotes the coefficient of z^j in P, expansion of (4.4) and the elementary symmetric identities give

$$a_1 = \delta T, \quad a_2 = \frac{\delta T^2}{2}(1 - p_2), \quad a_3 = \frac{\delta T^3}{6}(1 - 3p_2 + 2p_3). \tag{4.7}$$

In particular $a_1 > 0$, which also gives the local root branch by the implicit function theorem. Differentiating $P(\lambda(s)) = U(e^s - 1) + V(e^{-s} - 1)$ at zero yields

$$\begin{aligned} a_1 C_1 &= U - V, \\ a_1 C_2 + 2a_2 C_1^2 &= U + V, \\ a_1 C_3 + 6a_2 C_1 C_2 + 6a_3 C_1^3 &= U - V. \end{aligned} \tag{4.8}$$

Use $U - V = \delta c$ and $U + V = \delta(c^2 + 1)/2$ to obtain

$$\begin{aligned} TC_1 &= c, \\ TC_2 &= c^2 p_2 - \frac{c^2 - 1}{2}, \\ S := \frac{C_3}{C_1} &= c^2(3p_2^2 - 2p_3) + \frac{c^2 - 1}{2}(1 - 3p_2). \end{aligned} \tag{4.9}$$

Thus the remaining problem is an inequality for a probability vector and the affinity parameter.

5 A finite proof of the variance–affinity constraint

The proof uses the thermodynamic uncertainty inequality

$$\frac{C_2}{C_1} \geq \frac{2}{A}. \tag{5.1}$$

It is an established result [2]. The following finite-ring proof supplies all of its ingredients directly from the rates.

5.1 Stationarity and correctors

Let π be the stationary probability vector and write

$$a_i = \pi_i u_i, \qquad b_i = \pi_{i+1} v_i.$$

Stationarity means $a_i - b_i = J$, independently of i. An explicit positive construction is

$$\begin{aligned} \rho_i &= v_i/u_i, \\ w_i &= \sum_{k=0}^{n-1} \frac{1}{u_{i+k}} \prod_{j=0}^{k-1} \rho_{i+j}, \\ Z &= \sum_i w_i, \qquad \pi_i = w_i/Z, \qquad J = (1 - V/U)/Z. \end{aligned} \tag{5.2}$$

The empty product is one. The geometric recurrence gives $u_i w_i - v_i w_{i+1} = 1 - V/U$, proving positivity, normalization, and the common flux identity. In particular $J > 0$.

Let d_i be the increment assigned to a forward crossing of edge i, with $\sum_i d_i = 1$. For a single counted edge, d_i is its indicator. The backward generator is

$$Qf_i = u_i(f_{i+1} - f_i) + v_{i-1}(f_{i-1} - f_i).$$

Solve the centered Poisson equation

$$Qh_i + u_i d_i - v_{i-1} d_{i-1} = J, \qquad \sum_i \pi_i h_i = 0. \tag{5.3}$$

The forcing has stationary mean zero. The maximum principle gives $\ker Q = \operatorname{span}\{\mathbf{1}\}$: at a maximum of a function in the kernel, positivity of both neighbouring rates forces its neighbours to have the same value, and this propagates around the ring. Thus Q has rank $n-1$. Since $\pi^T Q = 0$, its image is precisely $\{b : \pi^T b = 0\}$. The Poisson equation therefore has a solution, unique after imposing the centering condition. One can also construct it by solving a finite cyclic recurrence for the gradient and integrating a zero-sum increment around the cycle. For an explicit cyclic integration, if $\sum_i r_i = 0$, then

$$h_i = -\frac{1}{n} \sum_{k=0}^{n-1} \sum_{j=0}^{k-1} r_{i+j} \quad \text{satisfies} \quad h_{i+1} - h_i = r_i.$$

Subtracting the π-mean centers h. These finite recurrences, including their solvability for signed forcing, are proved in the accompanying Lean construction.

Set $g_i = d_i + h_{i+1} - h_i$, so $\sum_i g_i = 1$. The first characteristic derivative is $C_1 = J$. Differentiating the centered eigenvector equation twice gives

$$C_2 = \sum_i \pi_i \left[u_i d_i^2 + v_{i-1} d_{i-1}^2 + 2u_i d_i h_{i+1} - 2v_{i-1} d_{i-1} h_{i-1} \right]. \tag{5.4}$$

The same identity can be obtained without an analytic eigenvector: solve a second finite Poisson equation for k and substitute the quadratic approximate eigenpair $\lambda(s) = Js + C_2 s^2/2$, $r_i(s) = 1 + sh_i + s^2 k_i/2$. Its tilted-matrix residual vanishes to order two. Multiplication by the adjugate transfers this vanishing to the determinant, identifying J, C_2 with the first two equations in (4.8). This is the route used by the formal proof.

Lemma 5.1 (Corrected-increment variance)**.** *The second cumulant has the exact quadratic representation*

$$C_2 = \sum_i (a_i + b_i) g_i^2. \tag{5.5}$$

Proof. Stationarity applied to h^2 yields the Dirichlet identity

$$\sum_i (a_i + b_i)(h_{i+1} - h_i)^2 = -2 \sum_i \pi_i h_i Q h_i.$$

Insert (5.3) and use $\sum_i \pi_i h_i = 0$. The right side becomes $2\sum_i (a_i h_i - b_i h_{i+1}) d_i$. Expanding the square in (5.5) then gives

$$\sum_i (a_i + b_i) d_i^2 + 2 \sum_i d_i (a_i h_{i+1} - b_i h_i),$$

which is (5.4) after shifting an index. □

5.2 Cauchy–Schwarz and the logarithm

Weighted Cauchy–Schwarz and $\sum_i g_i = 1$ give

$$C_2 \sum_i \frac{1}{a_i + b_i} \geq 1. \tag{5.6}$$

For $a > b > 0$,

$$\log(a/b) \geq \frac{2(a-b)}{a+b}.$$

Indeed the derivative of $\log y - 2(y-1)/(y+1)$ is $(y-1)^2/[y(y+1)^2] \geq 0$ for $y \geq 1$, and its value at one is zero. Apply this with $a_i - b_i = J > 0$. Telescoping the stationary factors gives

$$A = \sum_i \log(a_i/b_i) \geq 2J \sum_i \frac{1}{a_i + b_i}. \tag{5.7}$$

Combining (5.6) and (5.7) yields $AC_2 \geq 2J$, proving (5.1) because $C_1 = J$.

6 Reduction of the lower bound to exact algebra

Put $r = 1/c \in (0,1)$, $z = r^2$, and $t = \sqrt{p_2} \in (0,1]$. Each $x_i \leq t$, so

$$p_3 \leq t^3. \tag{6.1}$$

Since $A = 2\log((1+r)/(1-r))$, equations (4.9) and (5.1) imply

$$t^2 \geq \frac{1-r^2}{2} + \frac{r}{\log((1+r)/(1-r))}. \tag{6.2}$$

Lemma 6.1 (Rational logarithm bound)**.** *For* $0 \le r < 1$*,*

$$\log \frac{1+r}{1-r} \le 2r + \frac{2r^3}{3} + \frac{2r^5}{5} + \frac{2r^7}{7(1-r^2)}. \tag{6.3}$$

Proof. The difference between the right and left sides is zero at zero. Its derivative is $4r^8/[7(1-r^2)^2] \ge 0$. □

Define

$$D(z) = 105 - 70z - 14z^2 - 6z^3, \tag{6.4}$$

$$R(t,z) = 2D(z)t^2 - (1-z)(D(z)+105), \tag{6.5}$$

$$F(t,z) = 48t^4 - 32t^3 - 24t^2 + 8 + z(24t^2 - 7). \tag{6.6}$$

On $[0,1]$ we have $D(z) \ge 15$. The right side of (6.3) equals $2rD(z)/[105(1-z)]$. Taking reciprocals in (6.2), all denominators being positive, proves

$$R(t,z) \ge 0. \tag{6.7}$$

The key scalar statement is the following.

Lemma 6.2 (Exact scalar inequality)**.** *If* $0 \le t, z \le 1$ *and* $R(t,z) \ge 0$*, then* $F(t,z) \ge 0$*.*

Before proving the lemma, observe that it completes the lower-bound argument. Equation (4.9) gives the exact identity

$$16z\left(S + \frac{1}{16}\right) = F(t,z) + 32(t^3 - p_3) \ge 0. \tag{6.8}$$

Since $z > 0$, this is $S \ge -1/16$.

7 An exact positivity certificate

Proof of Lemma 6.2. At $z = 1$,

$$F(t,1) = (2t-1)^2(12t^2 + 4t + 1) \ge 0. \tag{7.1}$$

Suppose, for a contradiction, that $F(t,z) < 0$. Write $B = 24t^2 - 7$. The affine dependence on z and (7.1) imply $B > 0$. Set

$$q = -\frac{48t^4 - 32t^3 - 24t^2 + 8}{B}.$$

Then $z < q \le 1$. Also $B > 0$ implies $t^2 > 7/24 > (27/50)^2$, hence $t \ge 27/50$. Let

$$E(z,q) = 15 + (1-z)(1-q)(20 + 6(z+q)) > 0.$$

Direct polynomial expansion gives

$$D(z)R(t,q) - D(q)R(t,z) = (q-z)\big[D(z)D(q) + 105E(z,q)\big] > 0. \tag{7.2}$$

Since $D(z), D(q) > 0$ and $R(t,z) \ge 0$, we conclude that $R(t,q) > 0$. A second exact identity is

$$B^4 R(t,q) = 2(t-1)P_{14}(t), \tag{7.3}$$

where

$$\begin{aligned} P_{14}(t) = {} & 10616832t^{14} - 15261696t^{13} - 20127744t^{12} \\ & + 20336640t^{11} + 53514240t^{10} - 47832576t^{9} \\ & - 29349376t^{8} + 30811392t^{7} + 5745408t^{6} \\ & - 7539040t^{5} - 671776t^{4} + 645104t^{3} \\ & + 179568t^{2} - 17623t - 17623. \end{aligned} \tag{7.4}$$

The integer certificates in Appendix A prove $P_{14}(t) \geq 0$ on $[27/50, 1]$. Thus the right side of (7.3) is nonpositive, because $t \leq 1$, while its left side is positive, because $B > 0$ and $R(t, q) > 0$. This contradiction proves the lemma. □

For completeness, the certificates cover the three closed intervals

$$[27/50, 541/1000], \qquad [541/1000, 3/5], \qquad [3/5, 1].$$

For an interval $[a/q_0, b/q_0]$ they have the form

$$L\,P_{14}(t) = \sum_{j=0}^{14} c_j (q_0 t - a)^j (b - q_0 t)^{14-j}, \qquad L > 0, \quad c_j > 0. \tag{7.5}$$

All quantities L, c_j, a, b, q_0 in the appendix are exact integers. Each displayed identity is an identity of polynomials, and every basis factor is nonnegative on its indicated interval. Consequently the certificate is directly verifiable by expansion and positivity; it is not a numerical sign test. The identities are literally expanded and checked by the `ring` tactic in `ScalarCertificate.lean`.

One can also reproduce the calculation from the ordinary Bernstein basis. Write $P_{14}(t) = \sum_{m=0}^{14} p_m t^m$. On an interval $[a, b]$ its standard Bernstein coefficients are

$$\beta_k = \sum_{j=0}^{k} \frac{\binom{k}{j}}{\binom{14}{j}} (b-a)^j \sum_{m=j}^{14} p_m \binom{m}{j} a^{m-j}. \tag{7.6}$$

They satisfy $P_{14}(a + (b-a)y) = \sum_k \beta_k \binom{14}{k} y^k (1-y)^{14-k}$. The exact minimum β_k on the three intervals is, respectively,

$$\frac{1672420971724577170464 3}{14916062355041503906250000000},$$
$$\frac{22891488470631833136239701625237 07}{15258789062500000000000000000000000000},$$
$$\frac{42140182568}{6103515625}.$$

These positive rational numbers summarize, but are not substitutes for, the full exact identities in the appendix.

8 The upper bound, reversal, and sharpness

Let

$$e_2 = \frac{1 - p_2}{2}, \qquad e_3 = \frac{1 - 3p_2 + 2p_3}{6}.$$

These are nonnegative elementary symmetric sums of the probability vector x. Equation (4.9) rewrites as

$$S = 1 - 3(c^2 + 1)e_2 + 12c^2e_2^2 - 6c^2e_3. \tag{8.1}$$

The nonnegativity of C_2 gives $4c^2e_2 \le c^2 + 1$. Multiplication by $3e_2 \ge 0$ shows that the positive quadratic term in (8.1) is at most $3(c^2 + 1)e_2$. Thus $S \le 1$.

Reversing the ring orientation replaces C_1, C_3 by $-C_1, -C_3$, leaving their ratio unchanged. At the matrix level, if σ is the successor permutation, reversal replaces it by σ^{-1} and sets

$$u'_i = v_{\sigma^{-1}i}, \qquad v'_i = u_{\sigma^{-1}i}.$$

The untilted generator is unchanged, and U, V are exchanged. This supplies the case $U < V$ and completes the proof of the two inequalities.

The extremizing family was already identified in [1]. In an irreversible ring, take one forward rate $1/(n-1)$ and all others equal to one. At $V = 0$, the inverse-rate weights are

$$x_1 = \frac{1}{2}, \qquad x_2 = \cdots = x_n = \frac{1}{2(n-1)},$$

and the third/first cumulant ratio is $3p_2^2 - 2p_3$. Direct substitution gives

$$S = -\frac{n^2 - 8n + 8}{16(n-1)^2} \longrightarrow -\frac{1}{16}. \tag{8.2}$$

For each fixed n, sufficiently small positive backward rates approximate this value. The simple zero eigenvalue and the coefficient formulas vary continuously with the rates, with $C_1 \neq 0$ at the irreversible limit. Taking n large proves sharpness over strictly positive rates. A single increasingly slow forward stage makes its normalized waiting-time weight tend to one and hence gives $S \to 1$; the same positive-rate approximation applies. No better dimension-independent constants are possible.

9 Formal verification and reproducibility

9.1 Exact formal scope

The complete source development contains 34 local modules. The companion verification manifest records the exact source hashes, pinned environment, compiler exit, and axiom reports for the fresh check. Compilation alone does not establish bibliographic novelty or formalize stochastic sample paths.

The principal checked declarations are

```
UnicyclicWinding.winding_skewness_interval
UnicyclicWinding.exists_winding_branch_and_interval.
```

Their assumptions specify a finite index type, a permutation forming one cycle on the whole type, cardinality at least three, strictly positive forward and reverse rates, unequal rate products, and a marked edge. The first declaration quantifies over every local C^3 zero-eigenvalue branch of the actual tilted characteristic matrix; the second constructs such a branch from the rates and proves the interval. Neither theorem assumes a stationary law, a spectral decomposition, a variance formula, or an uncertainty bound.

The algebraic proof works with a positive midpoint Gram matrix and its characteristic polynomial. Finite Green sums construct the stationary law and both correctors. A

second-order residual is transferred to the determinant through $\mathrm{adj}(M)M = \det(M)I$. Polynomial derivative identities then connect the actual branch derivatives to the certified scalar bound. The small-state supplements have separate declarations in `TwoStateWinding.lean` and `PoissonWinding.lean`.

The formal theorem does not introduce continuous-time sample paths or formalize the probability-space limit in (2.2). Its stochastic interpretation is the standard proof in Section 3. Sharpness is proved in Section 8; the principal Lean interval theorems do not themselves assert the optimizing limiting sequence.

9.2 Environment and build

The development uses Lean 4.19.0, compiler commit `6caaee842e94`, and mathlib revision

`c44e0c8ee63ca166450922a373c7409c5d26b00b`.

The proof uses the ordinary Lean/mathlib foundations `propext`, `Classical.choice`, and `Quot.sound`, with no added problem-specific axioms. No declaration containing `sorry`, `admit`, or an unproved substitute is acceptable as verification.

The accompanying `verify.py` assembles every local source into one isolated compilation, excludes local object files from its imports, checks the pinned compiler and mathlib revision, and records hashes and compiler results. It rejects placeholder proofs and any nonstandard axiom dependency. After installing the stated toolchain and mathlib dependencies, run:

```
python3 verify.py --mathlib /path/to/mathlib4 \
  --lean-bin /path/to/lean-4.19.0-linux/bin --output verification
python3 check_certificate.py
```

A successful run returns zero and records `all_passed=true` in `verification.json` in the `verification` directory. Its source hashes identify the exact checked files. The independent certificate check uses Python and SymPy. It expands every integer certificate and checks the polynomial reduction identities. The Lean proof checks these identities independently with kernel-checked tactics.

10 Assessment and limitations

The result resolves the dimension-independent interval proposed in [1], within its positive-rate ring model and for the specified winding observable. Violating the interval therefore excludes that model for this observable. It does not settle the paper's other cumulant conjectures or the optimal bounds at each fixed ring size.

The argument does not settle the separate kurtosis, joint-cumulant, or fixed-size optimization questions. Extending the method to other higher cumulants remains an open direction.

11 Research process and AI disclosure

This work is fully autonomous agentic research. AI systems autonomously generated and developed the research, proofs, exact certificate, formalization, corrections and earlier versions of this manuscript. Opus 5.5 by Anthropic used with Claude Code revised the presentation of the manuscript on 1 November 2026: it rewrote the front matter and abstract, added the

note on earlier public work and consolidated this disclosure. The revision does not change the mathematics. Dev D. Goyal is responsible for publication. No expert human validation or peer review is claimed; the work has not been externally peer reviewed. The formal scope is stated in Section 9 and the remaining limits in Section 10.

References

[1] K. Ptaszyński, *Bounds on skewness and kurtosis of steady-state currents*, Physical Review E **106**, 024119 (2022). doi:10.1103/PhysRevE.106.024119. arXiv:2204.08421.

[2] T. R. Gingrich, J. M. Horowitz, N. Perunov, and J. L. England, *Dissipation Bounds All Steady-State Current Fluctuations*, Physical Review Letters **116**, 120601 (2016). doi:10.1103/PhysRevLett.116.120601. arXiv:1512.02212.

[3] A. C. Barato and U. Seifert, *Skewness and Kurtosis in Statistical Kinetics*, Physical Review Letters **115**, 188103 (2015). doi:10.1103/PhysRevLett.115.188103.

[4] T. Wampler and A. C. Barato, *Skewness and Kurtosis in Stochastic Thermodynamics*, Journal of Physics A: Mathematical and Theoretical **55**, 014002 (2022). arXiv:2110.14540.

[5] P. Pietzonka, A. C. Barato, and U. Seifert, *Universal bounds on current fluctuations*, Physical Review E **93**, 052145 (2016), Appendix C. doi:10.1103/PhysRevE.93.052145. arXiv:1512.01221.

[6] The Omega Institute, problem dossier on the minimum skewness of unidirectional cycles (issue #22468), public repository `the-omega-institute/trureturing`, committed 26 October 2026. `https://github.com/the-omega-institute/trureturing/blob/036cdd4840c6bd823af165fa3c6db3ea6a790929/Problems/ptaszynski-2022-unicyclic-skewness-minimum.md`.

A Complete integer positivity certificate

For each of the following three intervals, the entries satisfy the exact identity

$$L\,P_{14}(t) = \sum_{j=0}^{14} c_j (q_0 t - a)^j (b - q_0 t)^{14-j}.$$

The polynomial P_{14} is given in (7.4). Each L and each c_j is positive. On $[a/q_0, b/q_0]$ every term is nonnegative, proving $P_{14}(t) \geq 0$. These tables reproduce all coefficients from the companion `certificate.json`; they can be checked by exact integer polynomial arithmetic independently of the formal proof.

A.1 Interval $[\frac{27}{50}, \frac{541}{1000}]$

$$q_0 = 1000, \qquad a = 540, \qquad b = 541.$$

The positive scaling integer is

`152587890625000000000000000000000000000`.

j	Exact integer c_j
0	`411473622932365107200000000000000`
1	`433270014446302252032000000000000000`
2	`212666107695373464248320000000000000`
3	`67964008530644141722931200000000000`
4	`17125590750459670225554432000000000000`
5	`3790826898379534575429314560000000000`
6	`732847244576579907131200473600000000`
7	`115880696336563855822957791936000000000`
8	`14194130141757552314240540077760000000`
9	`1304416217743499635244179600673600000`
10	`87764680806209783845344451154720000`
11	`4188673901560280156128117021872320000`
12	`134303335865043552944379353846235700`
13	`259635736453792961815266983186472 45`
14	`228914884706318331362397016252 3707`

A.2 Interval $[\frac{541}{1000}, \frac{3}{5}]$

$$q_0 = 1000, \qquad a = 541, \qquad b = 600.$$

The positive scaling integer is

945035738051118986124078521728515625000000000000000000000000000.

j	Exact integer c_j
0	228914884706318331362397016252370 7
1	391034186455695508734059735719726425
2	225914152837602510687384993573666700 00
3	418892676188928948795121221957240500000
4	300357413965852537663494147833670000000 0
5	1178467472996552135491764021220250000000 0
6	2936836257820210201171508501600000000000 0
7	5012559767533165076942265465000000000000 0
8	6104524983079457903007381000000000000000 0
9	5415921496641430296616525000000000000000 0
10	3517544280094465578570000000000000000000 0
11	1650504067474299845500000000000000000000 0
12	5365076547094004500000000000000000000000 0
13	1091838300715612500000000000000000000000 0
14	1053504564200000000000000000000000000000 0

A.3 Interval $[\frac{3}{5}, 1]$

$$q_0 = 5, \qquad a = 3, \qquad b = 5.$$

The positive scaling integer is

50000000000000.

j	Exact integer c_j
0	21070091284
1	814395636165
2	15638522649125
3	197875680587750
4	1738751880296250
5	10499141885084375
6	43305986699484375
7	123317604936562500
8	245694438175000000
9	344833392701171875
10	339366668701171875
11	229364686621093750
12	101562059082031250
13	26618551025390625
14	3148590087890625

B The small-state conventions

B.1 Two states with two distinguished channels

For $n = 2$, the two oppositely traversed edges become parallel transition channels. They must remain distinguished when counting. Put

$$U = u_0u_1, \quad V = v_0v_1, \quad a = u_0 + u_1 + v_0 + v_1,$$

with all four rates positive and $U \neq V$. The correct tilted matrix is

$$zI - W_s = \begin{pmatrix} z + u_0 + v_1 & -(u_1 + v_0e^{-s}) \\ -(u_0e^s + v_1) & z + u_1 + v_0 \end{pmatrix}. \tag{B.1}$$

Only the marked channel is tilted; multiplying a whole combined off-diagonal entry by a counting factor would describe a different observable. Direct expansion gives

$$\det(zI - W_s) = z^2 + az - U(e^s - 1) - V(e^{-s} - 1). \tag{B.2}$$

The same characteristic differentiation therefore yields

$$S = 1 - \frac{6(U + V)}{a^2} + \frac{12(U - V)^2}{a^4}. \tag{B.3}$$

Let $\alpha = \sqrt{U}$, $\beta = \sqrt{V}$, and set

$$r = \frac{\alpha - \beta}{\alpha + \beta}, \qquad k = \frac{4(\alpha + \beta)^2}{a^2}.$$

Arithmetic–geometric mean gives $0 < k \leq 1$, and $r^2 \leq 1$. Equation (B.3) becomes

$$S = 1 - \frac{3k}{4}(1 + r^2(1 - k)), \qquad S - \frac{1}{4} = \frac{3}{4}(1 - k)(1 - kr^2).$$

Both factors in the latter expression are nonnegative, and the former expression is at most one. Hence the stronger two-state interval is

$$\frac{1}{4} \leq \frac{C_3}{C_1} \leq 1. \tag{B.4}$$

The actual two-channel determinant, local branch existence, and this derivative interval are the subject of `TwoStateWinding.lean`.

B.2 One state with counted events

In the conventional one-state counted-event model, forward and reverse events form independent Poisson counts with positive rates u, v. Their difference has cumulant generating function per unit time

$$\lambda(s) = u(e^s - 1) + v(e^{-s} - 1).$$

Thus $\lambda'(0) = \lambda'''(0) = u - v$. If $u \neq v$, then

$$\frac{C_3}{C_1} = 1.$$

The explicit derivative identity is formalized in `PoissonWinding.lean`. These small-state calculations are conventional boundary cases and are not additional novelty claims.