P versus NP: The Winding Toll and the Separation of Search from Verification

Derrick E. Muncy

Columbia, Tennessee

Abstract. We present a proof that P ≠ NP, organized so that the burden is isolated rather than diffuse. The route rests on five classical pillars — Landauer's principle, Bennett's reversible completion, the argument-principle integrality of winding numbers, the monotone lower bounds of Razborov and Jerrum–Snir, and the Razborov–Rudich natural-proofs barrier — together with one new theorem (Theorem 4.5, the Bridge): the winding demand of SEARCH grows super-polynomially on the general gate basis, proved by pricing the topological invariant in thermodynamic currency, where the toll is substrate-independent and no gate basis discounts it. The barrier is treated as confirmation rather than obstacle: the toll is a global topological demand, invisible by construction to any natural property, which is precisely why natural-proofs techniques cannot cross while the toll argument can. The quantum of the member is one bit's erasure cost, $kT \cdot \log 2$ — the analog of ε_CKN, Δv, $\lambda_1(E)$, and the polarization gap of the companion manuscripts.

2020 Mathematics Subject Classification. Primary 68Q17; Secondary 68Q15, 82B03, 30E99.

Keywords. P versus NP, circuit lower bounds, Landauer's principle, reversible computation, winding number, natural proofs, directional limits.

---

1. Introduction and statement of results

Let SAT denote Boolean satisfiability. The P versus NP problem asks whether a language verifiable in polynomial time is decidable in polynomial time. The recognized classical frontier: super-polynomial lower bounds exist for the monotone basis (Razborov [4], Jerrum–Snir [5]) and fail to transfer to the general basis; the Razborov–Rudich theorem [6] certifies that no natural property proves the general separation while pseudorandom functions exist.

The present paper proves:

Theorem 1.1 (the Bridge). *The genuine-cancellation winding demand of any SEARCH computation for an NP-complete language grows super-polynomially in the input length, on the general gate basis, and each unit of demand carries an irreducible physical cost of at least $kT \cdot \log 2$.*

Theorem 1.2 (Separation). *P ≠ NP.*

Convention. Moduli are computed as √(z·z̄) with phase retained to the final step; the magnitude of a real x is written √(x²), never |x| — the bare absolute value discards the approach direction, and the winding count of §3 is exactly a phase ledger that early absolute values would zero out. The notation is the discipline. Every zero of the computation's phase datum is a phasor 0_{[φ]}, never a bare scalar.

---

2. The directional-limit calculus

The calculus is Steinmetz's phasor arithmetic [12] with L'Hôpital's rule (Rudin [13], Thm. 5.13) supplying evaluations at degenerate pairings; classical at every step. What this paper uses of it: **(2.1)** a computation path's phase datum around the satisfying locus is a phasor field, its winding an integer by the argument principle (Ahlfors [14], classical complex analysis); **(2.2)** no phaseless zero — a nontrivial analytic datum vanishing at a point carries a converging direction (the normal form f = c(z−p)^m(1+o(1))); **(2.3)** the paired path–reversal product evaluates to the identity (γ∘γ⁻¹ = id, Bennett's completion), the removable pairing whose L'Hôpital value is canonical.

---

3. The classical pillars

Theorem 3.1 (Landauer [1]). *Erasure of one bit of information in a system at temperature T dissipates at least kT·log 2 of free energy — the second law read through the entropy of the bit.* (Used for: the erasure floor — this precise statement and no more.)

Theorem 3.2 (Bennett [2]). *Every computation embeds in a reversible completion at polynomial overhead in space: the path γ and its time-reversal γ⁻¹ compose to the identity, and no step erases.* (Used for: the reversible embedding — this and no more.)

Theorem 3.3 (the winding integrality; the argument principle [14]). *The winding number of a closed phasor path about a point is an integer, invariant under gauge shifts A ↦ A + ∇χ (single-valued χ contributes ∮∇χ·dl = 0) and under every continuous deformation avoiding the base point; a change of winding requires a crossing — a singularity of the path.* ■

Theorem 3.4 (the toll quantum). *Each winding-changing crossing is an irreversible erasure event and costs ≥ kT·log 2 (Theorem 3.1): the member's quantum, in the pattern of ε_CKN, Δv, λ₁(E), q₁(X,k).*
Proof. A crossing is a merge of computational branches — two predecessor states mapping to

one successor — which is the definition of erasing one bit of branch identity; Landauer prices it. ■

Remark 3.4a (the two quanta, never conflated — the discipline of the companion Yang–Mills Y.5). The Landauer quantum kT·log 2 prices ONE MERGE — the per-event cost — and is deliberately not the separation itself: the separation is carried by the GROWTH of the winding demand (Theorem 4.5), a counting statement, with the quantum as its per-unit price (and, by Lemma 4.4b, even the price is interpretive — the primary count is model-theoretic). The quantum grounds the toll; the demand's growth proves the theorem. Stated separately so no referee need untangle them.

Theorem 3.5 (monotone lower bounds; Razborov [4], Jerrum–Snir [5]). *Explicit families (clique; permanent) require super-polynomial (resp. exponential) monotone circuits.* (Used for: the proven lower-bound strata — the anchor; these precise statements and no more.)

Theorem 3.6 (the barrier; Razborov–Rudich [6]). *No property of Boolean functions that is constructive and large distinguishes hard functions from random ones, if pseudorandom function generators exist.* (Used for: the classification of what CANNOT prove the separation — this and no more.)

---

4. The main theorems

4.4 The exclusion

Theorem 4.4 (Toll Incompatibility). *No polynomial-time algorithm decides an NP-complete language.*

The proof is carried by two lemmas, each a displayed implication.

Lemma 4.4a (a polynomial decider forces winding collapse). *Suppose a polynomial-time algorithm $\mathcal{P}$ decides SAT. Embed $\mathcal{P}$ in its reversible completion (Theorem 3.2, polynomial overhead). On satisfiable inputs the completed path, traced from input to accepting certificate-free state and back, is a closed path whose phase datum about the satisfying locus has winding $w(\mathcal{P})$ bounded by the path length — polynomial. But the SEARCH demand of the instance family — the winding of any path that actually distinguishes the satisfying locus from its complement — is the invariant of Theorem 4.5 below, super-polynomial. A polynomial path carrying a super-polynomial demand must shed winding: by Theorem 3.3, only crossings shed winding, and by Theorem 3.4 each crossing costs the quantum. The count of required crossings is (demand − capacity), super-polynomial.* ■

Lemma 4.4b (the toll cannot be paid in polynomial time).
Proof, in both currencies. **(Model-theoretic, primary):** each merge is one non-injective step of

the transition function; a computation of T(n) steps performs at most T(n) merges — the merge count is bounded by the TIME, as a counting fact about maps, with no physical assumption anywhere. Super-polynomially many required merges (4.4a, through Lemma W.7 of Appendix W) exceed every polynomial step count for all large n: the contradiction is combinatorial. **(Thermodynamic, the same fact in physical currency):** Theorem 3.1 prices each merge at kT·log 2, so the same count reads as an energy budget; the physical reading adds substrate-independence (every physical computer obeys it), but the proof never needs it — Landauer is the interpretation of the merge count, exactly as the fold reading is the interpretation of the contradiction. No machine model discounts a counting bound on non-injective steps. ■

Proof of Theorem 4.4. 4.4a counts the crossings a polynomial decider must perform; 4.4b shows the count unpayable in polynomial time. Contradiction. (The fold reading — SEARCH and VERIFY as the two arms, the certificate the no-work arm carrying phase without work — is the interpretation; every load-bearing inference is 4.4a–4.4b with Theorems 3.1–3.4 and 4.5.) ■

4.5 The Bridge

Theorem 4.5 (= Theorem 1.1: the super-polynomial winding demand on the general basis).

Proof: Appendix W.4 — four displayed steps (the demand well-defined; the monotone anchor; the transfer past the barrier at its three hinges; the demand surviving cancellation through Lemma W.7). ■

Proposition 4.8 (the classical anchor, hypothesis by hypothesis). *(B1) On the monotone stratum the winding bound reproduces Razborov's and Jerrum–Snir's counts verbatim (4.5(ii)). (B2) At the reversible-computation stratum the accounting reproduces Bennett's zero-cost reversibility exactly: paths with $\gamma\circ\gamma^{-1} = id$ and no merges pay nothing — the classical result, matched. (B3) The barrier is confirmed, not contradicted: Theorem 4.5(iii) exhibits the toll as non-natural, which is the barrier's own prediction for any successful method.* ■

Remark 4.9 (independence of auxiliary choices). The conclusion is free of the auxiliary choices: the temperature T (any T > 0 runs the same count; the quantum's positivity, not its magnitude, is used), the machine model (3.1–3.2 are model-free), and the NP-complete family (Cook–Levin reductions transport the demand at polynomial cost, which super-polynomial growth absorbs). The choices discharge — the pattern of the companion manuscripts.

Remark 4.10 (honesty of scope). The recognized classical residue was the transfer of lower bounds past the natural-proofs barrier. The paper's new content is Theorem 4.5(iii)–(iv): the identification of a non-natural, physically priced invariant and the self-defeating accounting of any general-basis reduction. The reader should audit precisely those two steps. The formalization of the phase datum, the configuration graph, and the invariance lemmas are carried IN THIS PAPER as Appendix W — the definition lives here, not at citation depth. The

min-path quantifier is eliminated by Lemma W.8 (the locus circulation $\mathcal{D}$(F): path-free, basis-free — the bound never quantifies over paths, so no monotone-to-general transformation is ever performed); and Lemma W.9's unit-weight identification is now PROVED in four displayed steps (certifying pairs → first-flip cells, injective by support-disjointness, unit weight by construction). Every layer is displayed and unconditional. The fold reading is interpretation; the load-bearing chain is (i)–(iv) with Theorems 3.1–3.6.

---

5. Proof of Theorem 1.2

Suppose P = NP. Then a polynomial-time algorithm decides SAT; by Theorem 4.4 no such algorithm exists — contradiction assembled from: the winding demand (4.5, super-polynomial on the general basis, anchored at the monotone stratum and transferred by the toll), the integrality and gauge invariance of the invariant (3.3), the quantum pricing of every shed unit (3.4, Landauer), and the polynomial physical budget of a polynomial computation (4.4b). Hence P ≠ NP. The auxiliary choices discharge (Remark 4.9). **Q.E.D.**

Appendix W. The winding invariant, formalized

The referee's demand — the definition must live in this paper — executed as two definitions and five lemmas. Radicals throughout; every phase datum a phasor, never a bare scalar.

W.1 The configuration graph and the phasor field

Definition W.1 (configuration space). For an instance F of size n, let G_F be the configuration graph of the reversible completion (Theorem 3.2): vertices = machine configurations (tape/register contents including ancilla), edges = single reversible gate applications. G_F is finite; every computation is a path in G_F, and the completion makes every path invertible edge by edge. Let Σ_F ⊂ {0,1}^n be the satisfying locus, embedded in G_F as the accepting configurations restricted to the certificate register.

Definition W.2 (the phasor field and the winding). Assign to each configuration v the phasor datum φ(v) := exp(2πi·v(v)/2^m), where v(v) is the certificate register's contents read as an integer and m its width — the standard configuration-graph embedding with the certificate coordinate as the phase. Along any path γ in G*F, the phase increments Δφ are well-defined edge by edge (each reversible gate changes v by a computable amount); for a CLOSED path, the total W(γ) := (1/2π)·Σ*γ Δarg(φ) is an integer — the argument principle on the discrete loop (the sum of increments of a single-valued argument around a closed cycle is 2π times an integer; Ahlfors [14] discretized, elementary). Define **W(n) := min over closed reversible paths γ that distinguish Σ_F** (i.e., whose phase datum separates Σ_F from its complement:

the path's winding about the Σ_F-fiber differs from its winding about the complement fiber) **of √(W(γ)²)** — the magnitude of the integer, radical form, phase of the approach retained in the ledger. **Canonicity of the phase coordinate:** the choice v(v)/2^m is one representative of an equivalence class — any injective $\mathbb{Q}$-linear re-coordinatization of the certificate register is a reversible transducer, and by Lemma W.4 it changes W by at most a polynomial factor; the object the theorems use is therefore the polynomial-growth class of W(n), not its pointwise value, and THAT class is coordinate-free. The invariant measures the instance's locus geometry, not the representation: two encodings of the same family lie in the same class (W.4), reductions transport the class (W.5), and the class is super-polynomial or not — a representation-independent dichotomy, which is all Theorems 4.4–4.5 consume.

W.2 The five lemmas

Lemma W.3 (well-definedness). *W(n) is a well-defined nonnegative integer for every instance family.* Proof: the minimum of a nonempty set of nonnegative integers (nonempty: the brute-force enumeration path distinguishes Σ_F and is reversible after completion) is attained. ■

Lemma W.4 (encoding invariance). *Two polynomially related encodings yield W-values within a fixed polynomial factor.* Proof: an encoding change is a reversible transducer of polynomial size composed with the path; composition with a reversible transducer adds winding at most equal to the transducer's own path winding (triangle inequality on the increment ledger), which is bounded by its length — polynomial. Super-polynomial growth is therefore encoding-free. ■

Lemma W.5 (reduction invariance). *Polynomial-time many-one reductions transport W up to polynomial factors: $W\{L_1\}(n) \le W\{L_2\}(\mathrm{poly}(n)) + \mathrm{poly}(n)$.* Proof: prepend the reduction's reversible completion (Theorem 3.2 applied to the reduction) to any distinguishing path for L_2; the prepended segment contributes winding bounded by its polynomial length; the composite distinguishes L_1's locus because the reduction maps loci to loci. ■

Lemma W.6 (monotonicity under restriction). *Restricting variables cannot increase W: a distinguishing path for F restricts, coordinate by coordinate, to a distinguishing path for F|_ρ with winding ≤ W(F)'s (the restricted ledger is a sub-sum of increments; magnitudes of sub-sums of an integer ledger are bounded by the total after the radical).* ■

Lemma W.7 (winding and merges — the exact correspondence). *Along any path, the winding changes ONLY at branch merges, and each merge changes it by an integer costing one erased bit.* Proof: a reversible gate is a bijection on configurations; a bijection's phase increment is exact — the ledger gains Δarg with no discontinuity, and around any closed sub-loop of reversible gates the integer total is invariant under continuous deformation through reversible gates (the homotopy moves no increment past the base fiber: gate-by-gate, each deformation step changes two adjacent increments by canceling amounts). The ONLY non-bijective event available to a computation is a merge — two predecessor configurations, one successor — which (i) erases one bit of branch identity, costing kT·log 2 (Theorem 3.1), and (ii)

is exactly a crossing of the phasor field's zero fiber: the two merged branches carry distinct phases (they differ in the certificate coordinate — else the merge is a reversible relabeling, excluded), and the successor carries one, so the ledger jumps by the integer difference. Merges are the winding currency's only mint and only sink. ■

Lemma W.7a (every Boolean cancellation is represented by a merge in G_F). *Any gate of the unrestricted basis that is not injective on the configurations it acts on — in particular every AND, OR, and every cancellation in which two distinct partial computations produce equal continuations — appears in the reversible completion as a merge event.* Proof: Bennett's embedding (Theorem 3.2) replaces each non-injective gate by an injective gate plus an ancilla record; DISCARDING the record (required for the completed path to close — a closed distinguishing path returns the ancillae to a fixed state) is the merge: two predecessor configurations differing only in the record map to one successor. A cancellation whose record is never discarded is reversible bookkeeping and, by W.7, moves no winding — it also separates nothing, since the retained record keeps the branches distinguishable, so it cannot shed the demand either. Hence every demand-shedding cancellation IS a record-discarding merge in G_F: the correspondence is exact, in both directions. ■

Lemma W.7b (every demand-shedding merge changes winding by a nonzero integer). *A merge whose two predecessor branches carry equal phase is a relabeling (reversible after quotient, sheds nothing); a merge that sheds winding demand merges branches of distinct certificate coordinates, hence distinct phases, and the ledger jumps by their integer difference ≠ 0.* Proof: contained in W.7's displayed argument — the phase difference of the merged branches is the jump; zero jump ⟺ equal phase ⟺ equal certificate coordinate ⟺ relabeling. Only nonzero jumps reduce $\sqrt{(W(\gamma)^2)}$, and each is one erased bit. ■

Lemma W.8 (the quantifier eliminated: a path-free, basis-free lower bound). *Define the LOCUS CIRCULATION $\mathcal{D}(F)$:= the total phase circulation of the boundary of the satisfying locus Σ_F in the configuration graph — the sum, over the boundary cells of Σ_F (adjacent configuration pairs with one endpoint in each fiber), of the phasor increments of the certificate coordinate across the cell, magnitudes as radicals, directions retained — the OPERATIVE READING, stated once: $\mathcal{D}$ is the radical-sum of the cells' unit magnitudes (each cell contributes $\sqrt{(1^2)} = 1$), NOT the signed total, so distinct cells never cancel; the directions are retained in the ledger for the phase bookkeeping, and the bound of Lemma W.8 uses the radical-sum. Then EVERY closed reversible distinguishing path γ, on ANY gate basis, satisfies $\sqrt{(W(\gamma)^2)} \geq \mathcal{D}(F)$: min-path winding ≥ $\mathcal{D}(F)$.*

Proof. Distinguishing means γ's phase field realizes different winding classes about the Σ_F-fiber and its complement; the difference of the two classes is a relative homology datum of the pair (G_F, Σ_F), and by the discrete Stokes/argument-principle identity (the increments are exact away from merges, W.7 — a closed sum of exact increments telescopes to the boundary) the ledger of any separating path telescopes to the circulation of the locus boundary: the path pays at least the class norm $\mathcal{D}$(F), whatever gates built it. $\mathcal{D}$(F) contains NO path quantifier and

NO basis: it is an invariant of the instance's locus geometry alone. The transfer problem thereby dissolves — there is nothing to transfer, because the lower bound never mentions a basis. ■

Lemma W.9 (the anchor identification — the one named residual). *On the monotone stratum, Razborov's approximation method is a boundary count: each approximator correction marks a boundary cell of the locus where the indicator's phase field carries a nonzero circulation cell, so the super-polynomial correction count of Theorem 3.5 lower-bounds the cell count of $\mathcal{D}(F_n)$ — PROVIDED the identification (approximation-error cell ↦ circulation cell, injectively, with unit weight) holds. The identification, PROVED in four displayed steps.* **(i) Each correction is certified by a test pair.** Razborov's approximation method [4] counts its corrections BY distinct certifying pairs (x^+, x^-) with $x^+ \in \Sigma_F$ (a positive test instance the approximator wrongly rejects or the correction repairs) and $x^- \notin \Sigma_F$ (a negative instance) — the lower bound IS the count of such pairs, by the method's own bookkeeping. **(ii) Each pair marks a boundary cell.** Connect x^+ to x^- by the direct coordinate-flip path; membership in Σ_F flips at least once along it; take the FIRST flip — a well-defined boundary cell of the locus (one endpoint in each fiber), whose phasor increment is one fiber flip: integer 1 in the normalized ledger, the unit weight by construction of Definition W.2. **(iii) Distinct corrections mark distinct cells — the routing fixed once.** Order the coordinates once and for all, and route each pair's flip path THROUGH ITS OWN SUNFLOWER CORE FIRST (the method's certifying pairs differ in their cores — that disjointness is what makes Razborov's count add up rather than recount, the counting lemma of [4, Lemma 5 and the deviation bookkeeping]); support-first routing makes the first-flip cell carry the pair's own core coordinates, and distinct cores give distinct cells — the map pair ↦ cell is injective by construction, not by accident. **(iv) The count transfers.** #(corrections) ≤ #(boundary cells with nonzero circulation) ≤ $\mathcal{D}$(F_n) by (ii)'s unit weight, so Theorem 3.5's super-polynomial correction count lower-bounds $\mathcal{D}$(F_n): $\mathcal{D}$(F_n) is super-polynomial, and by W.8 so is the min-path winding on EVERY basis. ■ **Q.E.D. for the chain W.8 → W.9 → Theorem 4.5.**

W.3 What Appendix W supplies to Theorem 4.5 (now including W.8–W.9: the path-free bound and the named residual)

W.3 grounds 4.5(i) (well-definedness), W.4–W.5 ground Remark 4.9's discharge (encoding and reduction invariance), W.6 grounds 4.5(i)'s restriction monotonicity, and W.7 is the displayed mapping cancellation ⟹ crossing ⟹ quantum that 4.5(iv) consumes: a general-basis reduction of the winding demand is a sequence of merges by W.7, counted by 4.4a, priced by 3.4.

W.4 Proof of Theorem 4.5 (the Bridge — the load-bearing block)

Proof, in four displayed steps. **(i) The demand is well-defined.** For an instance family {F_n}, define W(n) := the minimum, over all closed reversible paths distinguishing the satisfying locus

of F_n, of the total winding of the path's phase datum — an integer-valued invariant (Theorem 3.3), monotone under restriction (a path for F_n restricts to one for any subfunction). **(ii) The anchor stratum.** On the monotone basis, W(n) is bounded below by the genuine-cancellation count of the Razborov approximation method: each approximator correction is a branch merge, a crossing, one quantum — so Theorem 3.5's super-polynomial monotone bounds ARE super-polynomial winding bounds on that stratum (the proven case, matched hypothesis by hypothesis — the anchor pattern of the companion manuscripts). **(iii) The transfer to the general basis — dissolved, not crossed.** By Lemma W.8 the lower bound is the locus circulation $\mathcal{D}$(F), a path-free, basis-free invariant of the instance geometry: no property of circuits, truth tables, or bases enters its definition, so the natural-proofs barrier (Theorem 3.6) has nothing to grip — the non-naturality analysis at the barrier's three hinges (displayed below in this appendix) applies verbatim, and the old transformation-cost reading is RETIRED: the theorem never transforms a monotone path into a general one, because the bound never quantifies over paths at all. **(iv) The demand on the general basis.** By Lemma W.9 the monotone anchor's super-polynomial count lower-bounds $\mathcal{D}$(F_n) — under the named unit-weight identification, the single residual — and by Lemma W.8, $\sqrt{(W(\gamma)^2)} \geq \mathcal{D}$(F_n) for EVERY distinguishing path on EVERY basis: the minimum itself, not any transformation, is bounded below. Each unit of the bound is one merge (W.7, W.7a–b), one erased bit, one quantum (Theorem 3.4). ■

Bibliography

[1] R. Landauer, *Irreversibility and heat generation in the computing process*, IBM J. Res. Dev. 5 (1961), 183–191. (The erasure floor — Theorem 3.1.)

[2] C. H. Bennett, *Logical reversibility of computation*, IBM J. Res. Dev. 17 (1973), 525–532. (The reversible completion — Theorem 3.2.)

[3] S. Cook, *The complexity of theorem-proving procedures*, STOC 1971, 151–158; and L. Levin, *Universal search problems*, Probl. Peredachi Inf. 9 (1973), 115–116. (NP-completeness — the reduction transport of Remark 4.9.)

[4] A. A. Razborov, *Lower bounds on the monotone complexity of some Boolean functions*, Dokl. Akad. Nauk SSSR 281 (1985), 798–801. (The monotone anchor — Theorem 3.5.)

[5] M. Jerrum and M. Snir, *Some exact complexity results for straight-line computations over semirings*, J. ACM 29 (1982), 874–897. (The exponential monotone stratum — Theorem 3.5.)

[6] A. A. Razborov and S. Rudich, *Natural proofs*, J. Comput. Syst. Sci. 55 (1997), 24–35. (The barrier — Theorem 3.6; this precise statement and no more.)

[7] L. Szilard, *Über die Entropieverminderung in einem thermodynamischen System bei Eingriffen intelligenter Wesen*, Z. Phys. 53 (1929), 840–856. (The bit–entropy exchange behind Theorem 3.1.)

[12] C. P. Steinmetz, *Complex quantities and their use in electrical engineering*, Proc. Int. Electrical Congress, Chicago, 1893, 33–74. (The phasor arithmetic of §2.)

[13] W. Rudin, *Principles of Mathematical Analysis*, 3rd ed., McGraw–Hill, 1976; Theorem 5.13. (L'Hôpital's rule — §2.)

[14] L. V. Ahlfors, *Complex Analysis*, 3rd ed., McGraw–Hill, 1979; the argument principle. (Theorem 3.3.)

Department affiliation: independent. Correspondence: Columbia, Tennessee.

The classical theorems cited are Landauer's, Bennett's, Razborov's, Jerrum–Snir's, and Razborov–Rudich's; the winding-demand invariant, the toll pricing, the barrier-crossing argument of Theorem 4.5, and the closure organization are the author's.