A Finite-Dimensional Bounded-Perturbation Certificate

Gaussian Reference Integrals, Scalar Proof Checks, and a Domain-Validation Agenda

A uniform cubic remainder bound yields a Gaussian approximation certificate

Dr. Tamás Nagy

tnagyphd@gmail.com

Working Paper • 2026-04-09

Build-std v2.0+739bece5f9 | 2026-07-20 21:45 | 37ccba85

A shared quadratic part gives a shared Gaussian reference term; a uniform remainder bound controls how far each integral can depart from it.

Abstract

Saddle-point approximations replace an action by its quadratic part, but transferring a quantitative error statement between applications requires an explicit remainder estimate. Here (C, ρ) is a supplied certificate pair, not a canonical invariant.

Let $(\Gamma, \mathcal{B}, \lambda)$ be a measure space, let measurable $\mathcal{A}_2, R : \Gamma \to \mathbb{R}$, let $K_0 > 0$, and suppose $0 < F_{\text{Gauss}} = K_0 \int_\Gamma e^{-\mathcal{A}_2}\, d\lambda < \infty$. If $|R| \le C/\rho^3 \le 1$ almost everywhere, then

$$F = F_{\text{Gauss}} \cdot (1 + \delta), \qquad |\delta| \le \frac{2C}{\rho^3}$$

where $F = K_0 \int_\Gamma e^{-\mathcal{A}_2 - R}\, d\lambda$ is positive and finite. The proof is an elementary change to the normalized reference probability measure followed by $|e^x - 1| \le 2|x|$ for $|x| \le 1$. The exponent three is part of the remainder parameterization, not a derived universal physical exponent.

The paper records a domain-validation agenda across twelve areas. None currently supplies the measurable normalized integral and full-domain remainder estimate needed to instantiate the theorem. Numerical calculations illustrate quartic, cubic, product, magnetization-Laplace, and Poisson/Stirling asymptotics, but do not validate those hypotheses. The companion checks 25 scalar consequences of supplied inequalities; it does not formalize integration, normalization, or the manuscript theorem. The established contribution is therefore the conventional finite-dimensional certificate and a falsifiable agenda for future domain work.

Keywords: saddle-point approximation, Gaussian approximation, uniform remainder bound, bounded perturbation, reproducibility

MSC 2020: 41A60, 81S40, 60H30, 11M06, 91A10

---

1. Introduction

1.1 A Recurring Approximation Pattern

The saddle-point (or steepest descent) approximation is widely used in mathematics and physics [6]. In its classical form, an integral dominated by a non-degenerate critical point is approximated by a Gaussian integral.

This paper asks whether error bounds from different applications can be compared through one explicitly defined remainder parameter.

Consider three practitioners who have never spoken to each other:

- **In a COS-inspired conditional mapping**, doubling an assigned strip-width parameter tightens the certificate by a factor of eight; this is not a general COS truncation-error theorem.
- **In a low-Reynolds-number conditional mapping**, assigning $\rho = K/\mathrm{Re}$ makes the certificate scale as Re^3; this is not a general law for the Stokes-to-Navier–Stokes solution gap.
- **For the Euler-product analogy**, truncating the local logarithmic series after grade 2 leaves a tail beginning at p^{-3s}; this is not a saddle-point integral.

The pattern extends further:

- In **machine learning**, a proposed assignment $\rho = N/P$ would produce an $(P/N)^3$ certificate if a matching uniform remainder bound were proved.
- In **game theory**, a proposed curvature assignment would produce a τ^{-3} certificate in a specified local neighborhood.
- In **Bayesian inference**, the template gives an $n^{-3/2}$ bound only after that finite-sample remainder bound is independently established; it is not the generic Bernstein–von Mises rate.
- For **variational quantum circuits**, the algebra defines a scale where a quadratic certificate becomes uncontrolled, not a barren-plateau theorem.
- A **magnetization Laplace toy model** illustrates an $O(1/N)$ correction; no universal finite-size theorem for particle systems is claimed.
- For **nonlinear control**, the template is conditional on a uniform local remainder estimate and does not predict universal policy switching.

Twelve mappings are discussed in Section 5, but only the abstract conditional certificate is established.

Some are candidate integral constructions; others, notably Euler products and rough volatility, are analogies only. None is currently a theorem instance.

1.2 A Finite-Dimensional Bounded-Perturbation Certificate

We initially attempted to unify these domains by matching their cumulant generating functions directly. This failed because the domains live in different function spaces — a characteristic function in finance cannot be mapped linearly to a Navier-Stokes velocity field. We spent considerable effort trying to build a dictionary of functional transformations before realizing that the algebraic structure of the Taylor expansion itself sidesteps the problem entirely.

The common template starts from a quadratic expansion at a non-degenerate critical point. A local Taylor expansion alone is not enough for the global integral bound below: the higher-order remainder must also satisfy a uniform estimate on the full integration domain, or a separately controlled localization error must be added.

Specifically, each $\mathcal{A}$ decomposes as:

$$\mathcal{A} = a_* + \mathcal{A}_2 + \mathcal{A}_{\geq 3}$$

Where $\mathcal{A}_2$ is the quadratic part and $\mathcal{A}_{\geq 3}$ is the remainder.

The paper uses an **admissible certificate pair** (C, ρ): both numbers are supplied, $C > 0$, $\rho > 0$, and the remainder must satisfy $|\mathcal{A}_{\geq 3}| \leq C/\rho^3$ almost everywhere. This avoids pretending that ρ is a canonical invariant: rescaling C rescales ρ, the zero remainder naturally permits $\rho = \infty$, and an unbounded remainder admits no such global pair. A local variant requires an additional tail estimate and is not proved here.

The central result is:

Theorem 1 (Finite-Dimensional Remainder Certificate; conventional manuscript theorem).

Let $(\Gamma, \mathcal{B}, \lambda)$ be a measure space. In the intended Euclidean case, Γ is a Lebesgue-measurable subset of $\mathbb{R}^n$ and λ is restricted Lebesgue measure.

Let $\mathcal{A}_2$ and R be measurable real-valued functions on Γ. Let $K_0 > 0$, and assume

$$0 < F_{\text{Gauss}} := K_0 \int_\Gamma e^{-\mathcal{A}_2}\, d\lambda < \infty.$$

Let $C, \rho > 0$ satisfy $B := C/\rho^3 \leq 1$ and $|R| \leq B$ λ-almost everywhere. Then $e^{-\mathcal{A}_2 - R}$ is measurable and integrable,

$$F = F_{\text{Gauss}} \cdot (1 + \delta), \qquad |\delta| \leq \frac{2C}{\rho^3}$$

where $F := K_0 \int_\Gamma e^{-\mathcal{A}_2 - R}\, d\lambda$ is positive and finite. For $\Gamma = \mathbb{R}^n$ and a positive-definite Hessian, F_{Gauss} has the standard determinant form. On a restricted domain it is a truncated Gaussian integral and is not determined by the Hessian alone.

Three consequences follow:

- **Conditional agreement.** Two systems with the same domain, measure, normalization, quadratic action, and compatible remainder bounds share the same Gaussian reference term and are close to it.
- **Cubic decay.** The certified bound $2C/\rho^3$ on $|\delta|$ scales as ρ^{-3}. Doubling ρ tightens that bound by a factor of eight — exactly for the certificate, not only in an asymptotic limit.
- **Product factorization.** If the measure and action genuinely factor, then $F_{\text{total}} = F_1 F_2$ and the algebraic cross-correction is $O(\rho^{-6})$ when both component bounds have the same scaling.

1.3 Domain-Validation Agenda

Earlier drafts classified twelve proposed mappings by how a domain parameter would relate to ρ. No mapping has yet passed the theorem's application gate. In every domain, the required deliverable is the same:

1. specify $(\Gamma, \mathcal{B}, \lambda)$ and measurable $\mathcal{A}_2, R$;
2. prove $0 < K_0 \int e^{-\mathcal{A}_2}\, d\lambda < \infty$;
3. prove an almost-everywhere full-domain bound $|R| \leq C/\rho^3 \leq 1$, or a localized analogue with an explicit tail term;
4. identify the domain observable with the resulting *integral ratio*, without changing object type to a pointwise payoff, density, PDE solution, or value function.

Section 5 records candidate domains and these falsifiable obligations. It derives no Reynolds-number, sample-size, curvature, circuit-depth, or finite-size law.

1.4 Companion Verification

The companion bundle verifies 25 elementary real-algebra consequences after the analytic hypotheses are supplied. It does not formalize integration, Taylor remainders, probability measures, PDE estimates, or any domain mapping. The canonical scalar source is platonic.py; its explicit 25-item contract list prevents an empty or partial replay from passing silently. The integration theorem is a conventional manuscript theorem, not a formally verified result.

An earlier development inventory mixed 30 core-algebra attempts with 80 domain-labeled or threshold templates. Review found one false implication and 21 automation failures. Four honest scalar supports were repaired with direct proofs, five optional duplicates were retired, and all 80 domain or threshold templates were removed from the formal inventory because their labels require model-specific analysis. The current replay passes all 25 declared contracts with no rejected target. No Lean artifact for Theorem 1 is claimed here unless the separately reported minimal export compiles without sorry.

1.5 What Is and What Is Not Proved

The companion checks establish only algebra downstream of explicitly assumed inequalities. The analytic derivation of Theorem 1 is given in Section 2.3 at the prose level and is not encoded there. The scripts also do not establish the existence of a perturbation expansion.

The identification of ρ in each domain (Section 5) is a modeling proposal. The cited literature motivates the domains but does not establish the displayed Latent bounds, and the toy computations in Section 7 do not validate those domain mappings.

The defensible result is therefore the conditional finite-dimensional certificate plus illustrative examples; the domain bridges are a research program.

1.6 Related Work

The saddle-point method (method of steepest descent) dates to Riemann and has been refined over a century. Bleistein and Handelsman [6] give the classical uniform asymptotic theory. Hörmander [15] provides the modern functional-analytic foundation.

The main contribution of the present paper is not a new asymptotic technique or a sharper Laplace theorem. The displayed estimate is an elementary bounded-perturbation comparison, packaged as a certificate with an explicit assumption ledger and proposed cross-domain mappings. Classical asymptotic theory [6], Bayesian Laplace approximations [20], and modern non-asymptotic Gaussian-approximation bounds [21–23] are substantially deeper: they derive useful error controls from local smoothness, concentration, and dimension-dependent conditions. Here the full-domain remainder bound is supplied as a hypothesis. Transfer to a domain still requires a domain-specific proof of that bound.

For individual domains, extensive saddle-point literature exists:

- in finance, Lewis [16] and Friz–Gerhold–Gulisashvili–Sturm [17]
- in statistical physics, Varadhan [18] and Ellis [19]
- in Bayesian inference, Tierney and Kadane [20]

The Latent framework does not supersede these; it proposes a common language for comparing remainder estimates when the cubic parameterization can be justified.

1.7 Organization

Section 2 proves the finite-dimensional remainder certificate. Section 3 records its direct consequences, and Section 4 gives conditional perturbative and product algebra. Section 5 presents twelve proposed or analogical domain mappings and distinguishes them from the established core result. Section 6 discusses limitations and open problems. Section 7 reports five numerical scaling illustrations. Section 8 describes the scope of the kernel checks.

---

2. The Framework

2.1 Grade Decomposition

For the theorem proved here, fix a measure space $(\Gamma, \mathcal{B}, \lambda)$ and measurable functions $\mathcal{A}_2, R : \Gamma \to \mathbb{R}$. The intended Euclidean case uses a Lebesgue-measurable $\Gamma \subseteq \mathbb{R}^n$ with restricted Lebesgue measure. Write

$$\mathcal{A}(\gamma) = a_* + \mathcal{A}_2(\gamma) + R(\gamma).$$

In the motivating saddle-point case, coordinates are centered at a critical point, $\mathcal{A}_2$ is a positive-definite quadratic form, and R is the higher-order remainder. Those Taylor origins are motivation, not assumptions needed by the bounded-perturbation theorem.

Working in coordinates centered at the critical path γ^* (where $\nabla\mathcal{A}(\gamma^*) = 0$), the quadratic part is simply $\mathcal{A}_2(\gamma) = \frac{1}{2}\gamma^T H \gamma$, with $H = \operatorname{Hess}\mathcal{A}|_{\gamma^*}$ acting as the Hessian.

Definition 1 (Admissible certificate pair). A pair $(C, \rho) \in (0, \infty)^2$ is admissible when $|R| \leq C/\rho^3$ λ-almost everywhere. It is useful for Theorem 1 when additionally $C/\rho^3 \leq 1$. The pair is a parameterization of a supplied essential-supremum bound, not a canonical invariant. If $R = 0$, arbitrarily large ρ are admissible; if R is essentially unbounded, no global pair exists.

Why the exponent 3? The exponent 3 in ρ^{-3} is a uniform worst-case certificate tied to the chosen parameterization. Grade-1 terms vanish at a critical point and grade-2 terms are absorbed into F_{Gauss}. Grade 3 can therefore be the first remaining Taylor grade, but the theorem itself uses only the supplied bound.

2.2 The Gaussian Approximation

The *Gaussian approximation* is the integral with only the grade-2 action:

$$F_{\text{Gauss}} = K_0 \int_\Gamma e^{-\mathcal{A}_2(\gamma)}\, d\lambda(\gamma),$$

which Theorem 1 explicitly assumes to be positive and finite.

For $\Gamma = \mathbb{R}^n$ and $\mathcal{A}_2(\gamma) = \frac{1}{2}\gamma^T H \gamma$ with H positive definite, this equals

$$F_{\text{Gauss}} = K_0 (2\pi)^{n/2} (\det H)^{-1/2}.$$

Where:

- $K_0 = e^{-a_*}$ absorbs the constant action value at the saddle point
- n is the ambient finite dimension
- The determinant $\det H$ encodes the curvature at the critical point.

The kernel target T10 checks the elementary positivity implication from assumed positive factors; it does not formalize the Gaussian integral.

2.3 The Correction Factor

The exact integral is $F = K_0 \int_\Gamma e^{-\mathcal{A}_2(\gamma) - R(\gamma)}\, d\lambda(\gamma)$. Measurability of $\mathcal{A}_2$ and R makes the integrand measurable. Since $|R| \le B$ almost everywhere, $e^{-B} e^{-\mathcal{A}_2} \le e^{-\mathcal{A}_2 - R} \le e^{B} e^{-\mathcal{A}_2}$ almost everywhere, so $0 < F < \infty$. Define the normalized probability measure

$$d\mu_G(\gamma) = \frac{K_0 e^{-\mathcal{A}_2(\gamma)}}{F_{\text{Gauss}}}\, d\lambda(\gamma),$$

we have

$$F = F_{\text{Gauss}} \cdot (1 + \delta)$$

where $\delta = \int_\Gamma (e^{-R} - 1)\, d\mu_G$. If $|R| \le C/\rho^3 \le 1$ almost everywhere, then:

$$|\delta| \le \frac{2C}{\rho^3}$$

Why the factor of 2? The elementary inequality $|e^x - 1| \le (e-1)|x| < 2|x|$ holds for $|x| \le 1$ (since $e - 1 \approx 1.718 < 2$).

The factor 2 is a convenient upper bound, not a sharp constant. Integrating the pointwise inequality against the probability measure μ_G proves the displayed bound. This argument is finite-dimensional and requires the full-domain remainder estimate.

3. Core Theorems

3.1 The Saddle-Point Sandwich

The first structural result confines F to a symmetric interval around F_{Gauss}.

Theorem 2 (Saddle-Point Sandwich). *Under Theorem 1's hypotheses:*

$$F_{\text{Gauss}} \cdot \left(1 - \frac{2C}{\rho^3}\right) \leq F \leq F_{\text{Gauss}} \cdot \left(1 + \frac{2C}{\rho^3}\right)$$

The kernel targets saddle_sandwich *(T16),* saddle_upper *(T6), and* saddle_lower *(T7) check this consequence from an assumed correction bound.*

The sandwich immediately gives the absolute error:

Corollary (Excess and Deficit Bounds).

$$|F - F_{\text{Gauss}}| \leq F_{\text{Gauss}} \cdot \frac{2C}{\rho^3}$$

Represented by kernel targets excess_bounded *(T8) and* deficit_bounded *(T9).*

3.2 The Precision Theorem

How large must ρ be to achieve a desired relative precision ε?

Theorem 3 (Conditional Precision). *Under Theorem 1's hypotheses, if $\rho^3\varepsilon \geq 2C$, then:*

$$\left|\frac{F - F_{\text{Gauss}}}{F_{\text{Gauss}}}\right| \leq \varepsilon$$

Represented by kernel target universal_rho_bound *(T20).*

For fixed C, this gives the sufficient certificate condition $\rho \geq (2C/\varepsilon)^{1/3}$.

Algebraically, tightening the certificate by a factor of 10 requires multiplying ρ by $10^{1/3} \approx 2.15$, provided the same parameterization and constant remain valid.

3.3 Conditional Two-System Comparison

Two systems can be compared through one Gaussian reference term only when their domain, measure, normalization, and quadratic action agree.

Theorem 4 (Two-System Error Bound). *Suppose two systems share the same domain, measure, normalization, and quadratic action, and each satisfies Theorem 1 with the same C, ρ. Then their integrals differ by at most $4C/\rho^3$ relative to their common F_{Gauss}.*

This is the ordinary triangle-inequality consequence of the two assumed bounds. It is not assigned a separate formal declaration because the cross-domain interpretation remains conditional. The scalar equality of two stipulated Gaussian inputs does not prove that a Hessian alone determines a restricted-domain integral.

This is a conditional comparison statement. Given a shared reference integral and the stated bounds, the two integrals agree to the specified precision.

The non-trivial content is establishing that specific domains satisfy the hypothesis. Section 5 proposes twelve mappings but does not prove those analytic estimates.

3.4 Monotonicity

The framework has clean monotonicity properties:

- **Higher ρ tightens the expression.** The bound C/ρ^3 is strictly decreasing in ρ (bound_improves, T4).
- **Larger C loosens the expression.** For fixed ρ, the bound grows with C (larger_C_wider, T19).
- **Higher ρ tightens the sandwich.** The width $2F_{\text{Gauss}} \cdot 2C/\rho^3$ shrinks strictly with ρ (tighter_rho_tighter, T18).

---

4. Perturbation Theory and Product Structure

4.1 The Perturbation Series

If an application independently supplies a convergent expansion for the correction, write:

$$\delta = \delta_1 + \delta_2 + \delta_3 + \cdots$$

and suppose $|\delta_k|$ is controlled by powers of $B = C/\rho^3 < 1$. Theorem 1 alone does not establish this expansion.

The following elementary inequalities then compare the assumed powers:

Theorem 5 (Order Domination). *For $0 \leq B \leq 1$: $B^2 \leq B$ and $B^3 \leq B^2$.*

Represented by kernel targets square_le_linear *(T21) and* cube_le_square *(T22).*

The truncated perturbation series at order k has a computable error:

Theorem 6 (Two-Term Bound). *If $|\delta_1| \leq B$ and $|\delta_2| \leq B^2$ with $0 \leq B \leq 1$, then $|\delta_1 + \delta_2| \leq B + B^2 \leq 2B$.*

Represented by kernel targets two_orders_upper *(T23) and* two_orders_lower *(T24).*

For $0 < B < 1$, the upper bound $B + B^2$ is strictly smaller than $2B$.

Theorem 7 (Refined vs. Crude Bound). *For $0 < B < 1$: $B + B^2 < 2B$. The refined bound is strictly tighter than the crude $2B$ estimate.*

Represented by kernel target refined_strictly_less *(T25).*

This comparison is useful only after the existence of the expansion and the separate term bounds have been established.

4.2 Product of Independent Saddle-Points

When both the action and measure factor into independent subsystems, the total integral is a product:

$$F_{\text{total}} = F_1 \cdot F_2 = F_{\text{Gauss},1}(1+\delta_1) \cdot F_{\text{Gauss},2}(1+\delta_2)$$

Expanding:

$$(1+\delta_1)(1+\delta_2) = 1 + \delta_1 + \delta_2 + \delta_1\delta_2$$

Represented by ring-identity target product_expansion *(T26).*

The cross-term $\delta_1\delta_2$ is the interaction term in the product. Its magnitude:

Theorem 8 (Cross-Term Bound). *If $|\delta_1| \le B_1$ and $|\delta_2| \le B_2$, then $|\delta_1\delta_2| \le B_1 B_2$.*

Represented by kernel targets cross_term_upper *(T27) and* cross_term_lower *(T28).*

If both subsystem corrections are $O(C/\rho^3)$, their cross-term is $O(C^2/\rho^6)$. Its ratio to a nonzero $O(\rho^{-3})$ primary scale tends to zero as $\rho \to \infty$; $\rho > 1$ alone does not make it negligible for arbitrary C or cancellations.

Summing the primary corrections and the cross-term yields the total error of the factorized approximation, which is bounded by the square of the individual bounds.

Theorem 9 (Product-Correction Upper Bound). *If $|\delta_1|, |\delta_2| \le B$ with $B \ge 0$, the total product correction satisfies:*

$$\delta_1 + \delta_2 + \delta_1\delta_2 \le 2B + B^2 = (1+B)^2 - 1$$

These optional envelope restatements are not part of the 25-item formal inventory; the checked product identity and cross-term bounds contain the relevant scalar content.

---

5. Research Agenda: Domain Instantiation Obligations

None of the twelve candidates below is a theorem instance, corollary, or machine-checked bridge. Each is retained only as a research question. Promotion requires all four obligations in Section 1.3, including a measurable normalized integral and a full-domain bound. A parameter substitution into $2C/\rho^3$ is not evidence that those obligations hold.

5.1 COS Option Pricing (Direct Mapping: $\rho = \eta$)

Practitioners of the COS method [1] use analyticity to control spectral convergence. The assignment below illustrates an eightfold change in the *postulated cubic certificate* when η doubles; it is not a derivation of the standard COS truncation error.

The COS method prices European options by expanding the transition density in a Fourier-cosine series. Convergence is controlled by the analyticity strip width η of the characteristic function $\varphi(u)$.

The log-characteristic function decomposes as:

$$\log \varphi(u) = \underbrace{i\mu u - \tfrac{1}{2}\sigma^2 u^2}_{\text{grade 2 (Gaussian)}} + \underbrace{\sum_{k\geq 3} \kappa_k (iu)^k / k!}_{\text{grade}\geq 3}$$

Where κ_k are the cumulants. The grade-2 part gives Black-Scholes pricing. The higher cumulants (skewness, kurtosis, …) are the corrections.

Conditional bridge. If the COS pricing error can be represented by Theorem 1 with $\rho = \eta$ and a full-domain remainder constant M, then:

$$P_{\text{COS}} = P_{\text{BS}} \cdot (1+\delta), \qquad |\delta| \leq \frac{2M}{\eta^3}$$

Where M bounds the higher cumulants and P_{BS} is the Black-Scholes price.

Conditional arithmetic targets:

- **Wider strip, tighter bound:** $\eta_2 > \eta_1 \Rightarrow M/\eta_2^3 < M/\eta_1^3$.
- **Precision target:** $\eta^3 \varepsilon \geq 2M \Rightarrow |P_{\text{COS}} - P_{\text{BS}}| \leq \varepsilon \cdot P_{\text{BS}}$.
- **Two independent assets:** the proposed joint pricing correction would be bounded by $(1+B)^2$ if the factorization assumptions held.
- **Gaussian limit:** when all higher cumulants vanish ($M = 0$), the stipulated correction vanishes.

This implication is conditional; Fang and Oosterlee's spectral-error analysis is not replaced by the displayed bound.

5.2 Navier-Stokes Regularity (Inverse Mapping: $\rho = K/\text{Re}$)

This subsection asks whether a cubic certificate can be obtained for a specified low-Reynolds-number perturbative regime. Merely observing that convection is nonlinear does not prove a cubic norm estimate for the solution gap.

The incompressible Navier–Stokes energy pairing does not provide the needed remainder: under standard divergence-free periodic, no-flux, or decaying conditions, $\langle u \cdot \nabla u, u\rangle = 0$. Nor is the Stokes solution generally zero. A viable finite-dimensional bridge would therefore need a different explicitly defined action or likelihood, its reference measure and normalization, and a nontrivial full-domain remainder estimate.

Conditional bridge. Suppose a specified finite-dimensional discretization and norm admit Theorem 1's remainder estimate with $\rho = K/\text{Re}$, where K is geometry-dependent. Then the certificate becomes:

$$|\delta_{\text{NS}}| \leq \frac{2C \cdot \text{Re}^3}{K^3}.$$

This algebra does not prove a continuum Navier–Stokes solution-error estimate, a turbulent transition law, or a universal cubic Reynolds-number law.

Conditional arithmetic targets:

- **Reynolds cubed:** the conditional cubic expression.
- **Lower Reynolds is better:** under the proposed assignment, $\mathrm{Re}_2 < \mathrm{Re}_1$ gives a tighter expression.
- **Halving Reynolds:** under the same assignment, $\mathrm{Re} \to \mathrm{Re}/2$ reduces the expression by a factor of 8.
- **Coupled flows:** a product model would have cross-corrections of order $O(\mathrm{Re}^6)$; no flow factorization is proved.

5.3 Euler Product Structure (Direct Mapping: $\rho_p = p^s$)

The Euler product $\zeta(s) = \prod_p (1-p^{-s})^{-1}$ is already a product of "saddle-points" — each prime contributes an independent factor. Expanding each factor reveals the grade structure immediately:

$$-\log(1-p^{-s}) = p^{-s} + \frac{p^{-2s}}{2} + \frac{p^{-3s}}{3} + \cdots$$

Unlike the other domains, the Euler product is not an integral with a saddle point — it is a formal product over primes. The "Grade Decomposition" here is a power series analogy, not a genuine instance of the path-integral framework.

For this analogy, "grade" simply records the power index in the local logarithmic series:

$$-\log(1-p^{-s}) = \sum_{k\geq 1} \frac{p^{-ks}}{k}.$$

Assigning grade k to the kth summand gives:

- The leading term p^{-s} is grade 1
- The term $p^{-2s}/2$ is grade 2
- The grade-≥ 3 correction is $O(p^{-3s})$

The cubic decay $O(p^{-3s})$ matches the chosen certificate exponent, but the mechanism is distinct from saddle-point integration. The Euler bridge is a structural analogy rather than an application of Theorem 1.

Analogy only. Assigning $\rho_p = p^s$ rewrites the third and later local series terms at a cubic scale. This does not produce a Gaussian reference integral, a probability normalization, or independent saddle points, so Theorem 1 is not applicable.

Conditional arithmetic targets:

- **Larger prime, smaller assigned expression:** $p_2 > p_1 \Rightarrow C/p_2^{3s} < C/p_1^{3s}$.
- **Deeper in the strip:** $s_2 > s_1$ tightens each assigned per-prime expression.
- **Two prime factors:** the algebraic cross-term is bounded by $B_p B_q$ after separate bounds are supplied.
- **Convergent-region analogy:** for $s > 1$, the smallest prime ($p = 2$) has the assigned $\rho \geq 2$, giving $C/\rho^3 \leq C/8$.

Connection to the Riemann Hypothesis. The Euler product discussion provides only a structural analogy: for $\mathrm{Re}(s) > 1$ the per-prime correction bounds are uniformly small (in the sense of the displayed per-prime p^{-3s} estimate), whereas for $0 < \mathrm{Re}(s) < 1$ such bounds need not be small for the smallest primes. This heuristic comparison yields no information about the location of zeta zeros; it only motivates a way to *phrase* a potential "transition" picture.

5.4 ML Scaling Laws (Inverse Mapping: $\rho = N/P$)

For machine learning, N/P is explored as a candidate dimensionless parameter. No universal relation between this ratio and a loss landscape's Taylor remainder is known.

A model with P parameters trained on N data points has a loss landscape $\mathcal{L}(\theta)$ minimized at θ^*. The loss expands as:

$$\mathcal{L}(\theta) = \mathcal{L}(\theta^*) + \underbrace{\tfrac{1}{2}(\theta - \theta^*)^T H(\theta - \theta^*)}_{\text{grade 2 (quadratic loss)}} + \underbrace{O(\|\theta - \theta^*\|^3)}_{\text{grade}\geq 3}$$

Where $H = \nabla^2 \mathcal{L}|_{\theta^*}$ is the Hessian of the loss.

Conditional bridge. If a specified model, dataset, neighborhood, and norm satisfy Theorem 1 with $\rho = N/P$, then:

$$|\delta_{\text{ML}}| \leq 2C \cdot \left(\frac{P}{N}\right)^3 .$$

The displayed expression is a hypothesis to test, not a general cubic overfit law.

This proposal is distinct from empirical neural scaling laws $L \propto N^{-\alpha}$ [3], which concern loss values. It could be tested by comparing Hessian-based local predictions with the true loss over a specified neighborhood and checking whether the required uniform remainder estimate holds.

Conditional arithmetic targets:

- **Double data rule:** doubling N with P fixed reduces the proposed expression by a factor of 8.
- **Larger assigned N/P tightens the expression.**
- **Ensemble proposal:** bounded cross-corrections would additionally require a justified factorization.

5.5 Game Theory (Direct Mapping: $\rho = \tau$)

Nash [4] proved equilibrium existence. Local stability and approximation error require additional game-specific assumptions.

An n-player game with payoff functions $U_i(x_1, \ldots, x_n)$ has a Nash equilibrium x^* where $\nabla_{x_i} U_i(x^*) = 0$ for all i. The payoff expands as:

$$U_i(x) = U_i(x^*) + \underbrace{\tfrac{1}{2}(x - x^*)^T H_i(x - x^*)}_{\text{grade 2 (quadratic payoff)}} + \underbrace{O(\|x - x^*\|^3)}_{\text{grade}\geq 3}$$

Where H_i is the Hessian of U_i at equilibrium.

Research obligation. A game-theoretic instance must first construct a positive finite normalization integral over strategy space. Theorem 1 does not bound a pointwise payoff ratio, and $U_i(x^*)$ may be zero or negative. Curvature alone does not supply the required global remainder bound.

Conditional arithmetic targets:

- **More curvature under the proposal:** $\tau_2 > \tau_1$ tightens the assigned expression.
- **Conditional precision:** $\tau^3\varepsilon \geq 2C$ makes the stipulated certificate at most ε.
- **Multi-player factorization proposal:** separately bounded sub-game corrections would have a scalar product envelope $2B + B^2$.
- **Certificate below one:** $\tau^3 > 2C$ makes the stipulated relative-error bound less than one; it does not prove equilibrium existence or stability.

5.6 Rough Volatility (Parametric Mapping: $\alpha = H - 1/2$)

The five bridges above all share a cubic correction: the leading error is $O(\rho^{-3})$ because grade-1 vanishes at the saddle point and grade-2 is absorbed into the Gaussian. Rough volatility breaks this pattern.

The rough Bergomi model [5] drives spot variance with a fractional Brownian motion of Hurst index $H \in (0, 1/2)$. The ATM skew has the short-maturity power-law behavior $\psi(T) \sim CT^{H-1/2}$ under model-specific assumptions. The following comparison uses $\alpha = H - 1/2$ only as an exponent parameter:

- **Reference exponent** ($\alpha = 0$, i.e. $H = 1/2$): the non-rough boundary; this is not an exact Black–Scholes limit for general Heston parameters
- **Grade 1** ($\alpha \neq 0$): the leading rough correction, proportional to α

Here the leading correction is *first-order* in α, not third. The roughness parameter $\alpha = H - 1/2 < 0$ directly controls the skew exponent:

$$\psi(T) \sim C \cdot T^{\alpha}, \qquad \mathrm{Var}(\text{hedge error}) \sim C' \cdot T^{2\alpha}.$$

The hedge error exponent 2α is exactly twice the skew exponent — the exponent-space analog of the cross-term δ^2 from the product structure (Section 4.2).

The exponent hierarchy. All rough vol observables are ordered:

$$\underbrace{2\alpha - 1}_{\text{forward slope}} < \underbrace{2\alpha}_{\text{hedge error}} < \underbrace{\alpha}_{\text{skew}} < 0.$$

This is an algebraic hierarchy of the displayed exponents, not a theorem about a rough-volatility model.

Arithmetic targets:

- *Skew exponent:* $H < 1/2 \Rightarrow \alpha < 0$.
- *Hedge exponent:* $2(H - 1/2) = 2H - 1$.
- *Rougher exponent:* $H_1 < H_2 \Rightarrow \alpha_1 < \alpha_2$.
- *Boundary exponents:* $H = 1/2 \Rightarrow \alpha = 0$ and $2\alpha = 0$; this algebraic identity does not imply that all model corrections vanish.

Rough volatility is not an instance of Theorem 1 in this manuscript. It is retained as an exponent-space analogy: $\alpha = H - 1/2$ shifts the short-maturity power law, while simple ordering identities can be checked algebraically.

5.7 Certificate Thresholds

The abstract bound has two arithmetic thresholds. The proof of Theorem 1 uses $B = C/\rho^3 \leq 1$; the looser displayed relative-error certificate falls below one when $2C/\rho^3 < 1$. Neither threshold is, by itself, a physical phase transition.

Definition 2 (Certificate threshold). For fixed C, define

$$\rho^* = (2C)^{1/3},$$

the value at which the *upper bound* $2C/(\rho^*)^3$ equals one. This does not imply that the actual error equals 100%.

This is only the defining scalar equality for the certificate threshold.

Proposition 1 (Bound Dichotomy).

- If $\rho > \rho^*$, the certified relative-error upper bound is less than one.
- If $\rho < \rho^*$, that upper bound exceeds one and is not informative enough to certify relative accuracy. It does not follow that the actual correction exceeds the signal or that any perturbation series diverges.

Conditional perturbation convergence. If a perturbation series has terms dominated by a geometric sequence B^k with $B = C/\rho^3 < 1$, the standard geometric-series argument gives convergence and the corresponding tail estimate. This condition is not a proof that any such series exists in the proposed applications.

Deep subcritical amplification. For $\rho < 1$, the expression C/ρ^3 exceeds C. This statement concerns the certificate expression, not the system's physical nonlinearity.

Exact substitution audit. The table applies the same relative-error certificate $|\delta| \leq 2C/\rho^3$ mechanically to every candidate. A row marked "not applicable" does not supply a ρ-parameterization and therefore has no certificate constant to audit.

Candidate	Substituted certificate	Unit-bound threshold
COS	$2M/\eta^3$	$\eta^3 = 2M$
Navier–Stokes	$2C\,\mathrm{Re}^3/K^3$	$2C\,\mathrm{Re}^3 = K^3$
Euler-product analogy	$2C/p^{3s}$	$p^{3s} = 2C$
ML scaling	$2C(P/N)^3$	$(N/P)^3 = 2C$
Game theory	$2C/\tau^3$	$\tau^3 = 2C$
Rough volatility	not applicable: exponent analogy only	not applicable
Bayesian inference	$2C/n^{3/2}$	$n^{3/2} = 2C$
Quantum circuits	$2CL^3/K^3$	$2CL^3 = K^3$
Statistical mechanics	$2C/N$	$N = 2C$
Stochastic control	$2C/\rho^3$	$\rho^3 = 2C$

Candidate	Substituted certificate	Unit-bound threshold
Large deviations	$2C/k$	$k = 2C$
Renormalization group	$2C/\rho^3$	$\rho^3 = 2C$

These are substitutions into the abstract inequality; they do not prove turbulence onset, an interpolation transition, or a cooperation threshold.

The certificate gap. The ratio ρ/ρ^* measures distance from the certificate threshold: if $\rho = k\rho^*$ with $k > 1$, then $2C/\rho^3 = 1/k^3$. Being twice above the threshold gives an eightfold certificate margin.

Possible relationships to known domain transitions remain conjectural:

1. *Turbulence.* Connecting the certificate threshold to laminar–turbulent onset would require a continuum estimate and calibration of C, K; it is not established here.
2. *Double descent.* The equality $N = P$ gives the proposed value $\rho = 1$, but no derivation connects the bound to test-error double descent.
3. *The critical strip.* The Euler-product series is absolutely controlled for $\mathrm{Re}(s) > 1$. The artificial local threshold above does not characterize the critical strip and yields no conclusion about zeta zeros.

The six mappings above span finance, fluid dynamics, number theory, machine learning, game theory, and stochastic volatility. The following six broaden the comparison but remain conditional or analogical.

5.8 Bayesian Inference (Direct Mapping: $\rho = \sqrt{n}$)

Under regularity conditions, Bernstein–von Mises theory gives asymptotic Gaussian posterior behavior, typically at an $O(n^{-1/2})$ first correction scale in standard formulations [8, 14]. A local cubic Taylor remainder evaluated on $O(n^{-1/2})$ fluctuations has pointwise size $O(n^{-1/2})$, not automatically $O(n^{-3/2})$. The assignment $\rho = \sqrt{n}$ below is therefore only a conditional template and requires an independently proved finite-sample remainder estimate.

Research obligation. A Bayesian instance may compare the exact marginal likelihood with its Gaussian-reference integral after proving the theorem's hypotheses. Theorem 1 does not give a pointwise posterior density-ratio bound. The assignment $\rho = \sqrt{n}$ remains unproved.

Under the stipulated assignment $\rho = \sqrt{n}$, doubling n multiplies ρ by $\sqrt{2}$ and divides the certificate by $2\sqrt{2} \approx 2.83$. This is an algebraic consequence of the assignment, not a general statistical rate.

This conditional $O(n^{-3/2})$ expression must not be identified with the Bernstein–von Mises theorem; the manuscript provides no statistical proof of it.

If two normalized approximation factors genuinely multiply and each correction has already been bounded, the algebraic product identity from Section 4.2 applies. This does not by itself give a posterior-error theorem for combined experiments.

5.9 Quantum Variational Circuits (Inverse Mapping: $\rho = K/L$)

McClean et al. [9] identified barren plateaus through concentration and exponentially small gradients in classes of random circuits. A breakdown of a local quadratic approximation is not equivalent to exponentially flat gradients.

The assignment $\rho = K/L$ is proposed as a testable depth parameter; it is not derived from the concentration mechanism in [9].

Conditional circuit-depth statement. Under the proposed assignment and an independently established remainder bound, the certificate exceeds one when

$$\frac{2CL^3}{K^3} > 1,$$

equivalently $2CL^3 > K^3$. This says only that the stipulated quadratic-error bound becomes non-informative; it neither proves exponential flatness nor identifies a barren-plateau threshold.

Under the stipulated inverse mapping, halving depth divides the certificate expression by eight. No physical analogy to turbulence is needed for this arithmetic statement.

When a circuit model and its approximation factorize exactly into independent blocks, the product identity applies. Weakly coupled or entangled blocks require a separate interaction estimate.

5.10 Statistical Mechanics (Direct Mapping: $\rho = N^{1/3}$)

Partition functions motivate Gaussian and saddle-point approximations in statistical mechanics [10]. This does not make every finite-size model an instance of Theorem 1.

Finite-size corrections depend on the Hamiltonian, dimension, boundary conditions, and distance from criticality. The manuscript's numerical example is an explicitly defined one-dimensional magnetization Laplace integral with an $O(1/N)$ saddle correction; it does not prove an $O(1/N)$ law for arbitrary N-particle systems.

Toy-model parameterization. If a specified finite-size model independently satisfies Theorem 1 with $\rho = N^{1/3}$, then:

$$\left| \frac{F_N}{F_{\mathrm{MF}}} - 1 \right| \le \frac{2C}{N} = \frac{2C}{\rho^3},$$

where F_N and F_{MF} must be defined for that model. This is only an arithmetic substitution.

A subtlety of the proposed 3D volume scaling is that doubling N multiplies ρ by $2^{1/3} \approx 1.26$, halving the certificate expression. The full eightfold change requires octupling N.

No identification of a thermodynamic critical temperature with ρ^* is proved. At most, a calibrated certificate reaching one could signal loss of control of a Gaussian approximation; it cannot establish the existence or location of a phase transition.

Exactly independent lattice factors obey the elementary product identity. Weak coupling and cluster expansions require additional estimates.

5.11 Stochastic Optimal Control (Direct Mapping: ρ = control authority)

LQR is exact for linear dynamics and quadratic costs. Quantifying error under nonlinear dynamics or costs requires model-specific HJB stability and remainder estimates.

The Hamilton-Jacobi-Bellman equation for stochastic optimal control [11] has the LQR as its grade-2 component — the exact solution when the dynamics and cost are both quadratic. When either departs from quadratic, the grade-≥ 3 terms generate an LQR error.

Research obligation. A control instance must construct a positive finite normalization integral whose action represents the specified control problem. Theorem 1 does not bound a pointwise value-function ratio, and $V_{\mathrm{LQR}}(x)$ can vanish. HJB stability and nonlinear remainder estimates are separate obligations.

If a model justifies interpreting stronger control authority as larger ρ, doubling that assigned parameter divides the certificate expression by eight.

When ρ drops below ρ^*, the certificate becomes at least one and no longer certifies relative accuracy. This does not imply a universal qualitative change in the optimal policy; saturation and switching are model-dependent possibilities outside the present bound.

Only exactly factorized controller models inherit the product identity directly. Weak coupling requires a separate interaction estimate.

5.12 Large Deviation Theory (Direct Mapping: $\rho = k^{1/3}$)

This section records a classical Poisson saddlepoint example [7, 12]. It does not establish a new saddlepoint theorem.

For n iid random variables with cumulant generating function $\kappa(t)$, the probability $P(\bar{X}_n \geq a)$ involves an exponential integral whose grade-2 component is the Gaussian tail (the CLT) and whose higher grades give the precise tail correction.

Poisson example. Stirling's expansion gives the saddlepoint approximation to $P(X = k)$ for a Poisson(λ) random variable with correction:

$$\delta = \frac{P_{\mathrm{exact}}(k)}{\widehat{p}_{\mathrm{saddle}}(k)} - 1 = O(1/k) = O(\rho^{-3}),$$

where $\widehat{p}_{\mathrm{saddle}}(k) = \exp(k - \lambda - k\ln(k/\lambda))/\sqrt{2\pi k}$.

As $k \to \infty$, Stirling's correction vanishes. This is distinct from asserting a central-limit theorem for fixed λ. Cramér regularity supports classical saddlepoint methods for many light-tailed laws, but failure of a global moment-generating function does not imply that every saddlepoint or asymptotic method is unreliable at all scales. The Poisson calculation cannot be promoted to a universal non-exponential-family theorem.

5.13 Renormalization Group (Direct Mapping: ρ = distance from RG fixed point)

The renormalization-group language offers a possible analogy for scale-dependent suppression [13], but it does not derive the certificate.

Under coarse-graining, physical systems flow toward RG fixed points. Near a fixed point, couplings decompose into relevant (growing), irrelevant (decaying), and marginal directions. The irrelevant directions decay as powers of the distance to the fixed point — and this *is* the Grade Decomposition. The grade-2 component is the linearized RG flow. The higher-grade corrections are the nonlinear terms that die away as the fixed point is approached.

Conditional RG statement. If a model-specific RG flow admits Theorem 1's cubic remainder parameterization, then the same arithmetic bound applies.

RG irrelevant exponents are generally model- and universality-class-dependent; no general reason forces the leading correction exponent to be three. Relating an RG eigenvalue to C or ρ remains an open modeling problem.

The certificate threshold ρ^* does not prove an RG crossover or a change of universality class.

5.14 The Prime Transition Ordering

Under the artificial local threshold assigned in Section 5.7, primes cross that threshold one at a time as $\mathrm{Re}(s)$ decreases. This arithmetic ordering is not a phase transition of the zeta function.

Under the purely algebraic assignment $\rho_p = p^s$, smaller primes have smaller assigned parameters. As s decreases toward 1/2:

1. $p = 2$ crosses the assigned threshold first because it has the smallest ρ_p.
2. Then $p = 3$, then $p = 5$, and so on; this is only magnitude ordering.
3. Under the artificial local bound, a prime can cross the chosen certificate threshold; this has no established implication for the analytically continued zeta function.

The product-correction identity applies to actual factor corrections with controlled signs and magnitudes, not merely to upper bounds B_p, B_q. Therefore the statement that subcritical upper bounds force explosive growth is not justified.

Connection to zeta zeros (speculative). The ordering of the assigned parameters suggests a possible question about zeta zeros, but the present work gives no constraint on zero locations.

---

6. Discussion

6.1 What the Template Does—and Does Not—Unify

Why compare the same formula across twelve domains?

A smooth finite-dimensional functional has a local Taylor expansion near a non-degenerate critical point. The full-domain uniform remainder bound used by Theorem 1 is an additional hypothesis and is not automatic.

The supplied ratio C/ρ^3 records the asserted remainder size. The integral certificate is the consequence once all analytic hypotheses hold.

The downstream algebra does not change, but the difficult analytic work is domain-specific: defining the integral, normalization, norm, controlled domain, C, and ρ, then proving the remainder estimate.

The RG comparison in Section 5.13 is heuristic and does not force a cubic correction exponent.

6.2 The Direct/Inverse Contrast

The direct and inverse mapping types have distinct practical signatures (the parametric type, represented by rough volatility, is discussed separately in Section 6.5):

Property	Direct (ρ = param)	Inverse ($\rho = K$/param)
Improving ρ means…	increase the parameter	*decrease* the parameter
Error scales as…	C/param^3	$C \cdot \text{param}^3/K^3$
Practitioners want to…	maximize the parameter	maximize the parameter
But the error…	decreases cubically	*increases* cubically

Under the proposed inverse mappings, increasing the domain parameter loosens the certificate.

No conclusion about turbulence or overparameterized learning follows until the corresponding mappings are proved. Algebraically, doubling the proposed inverse parameter multiplies the *upper-bound expression* by eight.

6.3 The Eightfold Rule

Doubling ρ divides the certificate expression by eight:

$$\frac{C/(2\rho)^3}{C/\rho^3} = \frac{1}{8}$$

The factor $2^3 = 8$ is exact for the expression C/ρ^3. It does not imply that an actual error ratio is exactly eight, and it applies to a domain only after its error has been bounded in this form.

- Widen an assigned strip parameter from 2 to 4: the conditional bound is 8× smaller.
- Halve Reynolds number under the proposed inverse mapping: the conditional bound is 8× smaller.
- Move to a prime $2^{1/s}$ times larger: 8× smaller per-prime correction.
- Double the dataset when $\rho = N/P$: the conditional bound is 8× smaller.
- Double the assigned payoff curvature: the conditional bound is 8× smaller.

This is a scaling identity for the certificate, not a domain-independent cost-benefit law.

6.4 The Second-Order Signal

When genuinely independent subsystems combine, the cross-term $\delta_1\delta_2 = O(\rho^{-6})$ is asymptotically smaller than a nonzero $O(\rho^{-3})$ primary scale as $\rho \to \infty$. The condition $\rho > 1$ alone does not guarantee numerical negligibility.

In the Euler product, this is the two-prime correction. In portfolio pricing, this is the correlation effect between assets. In coupled NS, this is the interaction between flow regions.

The scalar product identity gives the exact decomposition:

$$(1+\delta_1)(1+\delta_2)-1=\underbrace{\delta_1+\delta_2}_{O(\rho^{-3})}+\underbrace{\delta_1\delta_2}_{O(\rho^{-6})}$$

The first-order sum $\delta_1+\delta_2$ is the superposition. The second-order product $\delta_1\delta_2$ is the interaction.

The scale separation requires genuine factorization, controlled constants, and absence of cancellation in the primary corrections.

6.5 The Three Bridge Types

The twelve domain bridges organize into three distinct types, each with characteristic scaling:

Type	Mapping	Scaling when parameter doubles	Domains
Direct	$\rho = x$	$\delta \to \delta/8$	COS, Euler, Game, Bayesian, Stat Mech, Control, Large Dev, RG
Inverse	$\rho = K/x$	$\delta \to 8\delta$	NS, ML, Quantum
Parametric	$\rho \sim f(\alpha)$	depends on f	Rough vol

The Bayesian bridge has an additional subtlety: because $\rho = \sqrt{n}$, doubling the *data* (not ρ) gives only a $2\sqrt{2} \approx 2.83\times$ improvement instead of $8\times$. This is the statistical estimation bottleneck — the square root slows the convergence compared to domains with linear ρ mapping.

6.6 Limitations

1. **Bounded-remainder regime only.** The proof requires $B = C/\rho^3 \leq 1$; $\rho > 1$ alone is neither necessary nor sufficient. Outside this regime the displayed exponential inequality does not provide the stated certificate.
2. **Conditional exponent and non-universal prefactor.** The exponent -3 is built into the definition of ρ and the assumed remainder bound; it is not derived as a universal physical exponent. Each application must justify both the parameterization and C.
3. **Smoothness required.** Not every system has a smooth action functional with a non-degenerate critical point. Bifurcations, degenerate saddles, and topological obstructions are outside the theory.
4. **Statistical scope.** Standard regularity does not by itself give the displayed finite-sample $O(n^{-3/2})$ density-ratio bound. That bridge remains unproved.

6.7 Open Problems

1. **Sharp constants.** Compute C exactly for the COS and Euler product bridges. The current bounds are not sharp.
2. **Higher-order product structure.** Extend the two-subsystem product (Section 4.2) to n independent subsystems. The product

$$\prod_{j=1}^{n}(1+\delta_j)$$

generates pair, triple, and higher cross-terms. Under a common $O(\rho^{-3})$ component bound, these begin at orders ρ^{-6} and ρ^{-9}. When does the full tower converge absolutely?

3. **Infinite-dimensional path spaces.** The current formulation assumes finite-dimensional Γ. Extending to infinite-dimensional path spaces (genuine path integrals in the sense of Feynman) requires regularization. Does the grade structure survive regularization?

4. **Euler-product analogy.** Determine whether any genuine positive finite-dimensional integral representation yields a controlled remainder related to a local Euler factor. The current power-series analogy has no implication for zeta zeros.

5. **Circuit-depth remainder estimates.** Determine whether any explicit ansatz family admits a uniform quadratic-remainder estimate with a depth-dependent parameter. Only after such an estimate is proved can it be compared with known barren-plateau concentration results.

6. **Bayesian model selection.** When comparing two models with different C_1, C_2, the model with smaller C (simpler grade structure) has better Laplace approximation at any sample size. Is C computable from the Fisher information matrix? If so, the framework gives a finite-sample correction to BIC/AIC.

7. **Quantitative RG comparison.** Determine whether a model-specific change of variables can relate the certificate to an RG irrelevant exponent. No universal relation between that exponent, C, and ρ^* is currently established.

6.8 Conclusion

The manuscript proves a finite-dimensional conditional certificate: a supplied full-domain bound $|\mathcal{A}_{\geq 3}| \leq C/\rho^3 \leq 1$ controls the relative error of the Gaussian reference integral. The companion kernel checks elementary algebra downstream of assumed inequalities. The twelve domain mappings remain proposals or analogies and require separate analytic validation; the toy computations in Section 7 illustrate scaling behavior but do not validate universal physical, statistical, or machine-learning laws.

---

7. Numerical Illustrations

We illustrate several asymptotic rates using five models with exact or high-precision reference values, plus a doubling-ratio check on the half-line cubic integral.

Scope and limitations. These computations do not validate Theorem 1 because the polynomial remainders in the full-line and half-line integrals are unbounded on their integration domains, whereas the theorem assumes a full-domain uniform remainder bound. The displayed constants are selected numerical comparison curves, not fitted or analytically derived certificate constants. A rigorous application would need localization plus an explicit tail estimate. The examples also do not validate the domain mappings in Section 5.

7.1 Model 1: Quartic Path Integral

The simplest non-Gaussian integral:

$$F(\rho) = \int_{-\infty}^{\infty} \exp\left(-\frac{\rho^2 x^2}{2} - \frac{x^4}{4}\right) dx, \qquad F_{\text{Gauss}} = \frac{\sqrt{2\pi}}{\rho}.$$

This model has no cubic term ($\mathcal{A}_3 = 0$); its observed correction is $O(\rho^{-4})$, faster than the cubic comparison curve.

ρ	δ (exact)	$2C/\rho^3$ ($C = 0.5$, selected comparison)	Below curve
1	-0.2279	1.000	✓
2	-0.0380	0.125	✓
5	-0.00119	0.008	✓
10	-7.5×10^{-5}	0.001	✓
20	-4.7×10^{-6}	1.25×10^{-4}	✓

Empirical rate: $|\delta| \sim \rho^{-3.98}$, confirming the $O(\rho^{-4})$ rate from the absent grade-3 term.

7.2 Model 2: Half-Line Cubic Integral

Restricting the Gaussian integral to a half-line removes the parity cancellation of the cubic perturbation, so this example displays the expected cubic asymptotic scale.

$$F(\rho) = \int_{0}^{\infty} \exp\left(-\frac{\rho^2 x^2}{2} - \frac{x^3}{3}\right) dx, \qquad F_{\text{Gauss}} = \frac{\sqrt{\pi/2}}{\rho}.$$

The half-line domain $[0, \infty)$ breaks the symmetry that suppresses odd-grade contributions. The cubic term $x^3/3$ now contributes at first order, giving a genuine $O(\rho^{-3})$ correction.

ρ	δ (exact)	$2C/\rho^3$ ($C = 0.6$, selected comparison)	Below curve
1	-0.2434	1.200	✓
2	-0.0562	0.150	✓
5	-0.00420	0.0096	✓
10	-5.3×10^{-4}	0.0012	✓
20	-6.6×10^{-5}	1.5×10^{-4}	✓

Empirical rate: $|\delta| \sim \rho^{-2.98}$, saturating the universal $O(\rho^{-3})$ bound.

7.3 The Eightfold Rule in Practice

Using the half-line cubic model (which has a clean $O(\rho^{-3})$ rate), we verify the eightfold rule by doubling ρ at each step:

ρ	$\lvert\delta\rvert$	Ratio to previous	Expected
2	5.623×10^{-2}	—	—
4	8.115×10^{-3}	6.93	8
8	1.036×10^{-3}	7.83	8
16	1.298×10^{-4}	7.98	8
32	1.623×10^{-5}	8.00	8

The observed ratio approaches 8, consistent with a numerically estimated asymptotic $|\delta| \sim c\rho^{-3}$. The ratio is not exactly 8 at finite ρ; only the comparison expression has an exact doubling ratio.

7.4 Product Cross-Terms

For two independent quartic integrals $F_1(\rho) \cdot F_2(\rho)$, the product correction $\delta_{\text{prod}} = \delta_1 + \delta_2 + \delta_1\delta_2$. The cross-term $\delta_1\delta_2$ should decay at twice the individual rate:

ρ	$\delta_1\delta_2$	$\lvert\delta_1\delta_2\rvert/\lvert\delta_1 + \delta_2\rvert$
2	1.4×10^{-3}	1.9%
5	1.4×10^{-6}	0.06%
10	5.6×10^{-9}	0.004%
20	2.2×10^{-11}	0.0002%

Empirical rate: $|\delta_1\delta_2| \sim \rho^{-7.96}$ (expected -8 for quartic $\times$ quartic). At the sampled values with $\rho \geq 3$, the cross-term is less than 1% of the leading correction.

7.5 Magnetization Laplace toy model: Finite-Size Corrections

For an explicitly defined magnetization Laplace toy model (often used as a mean-field-style saddle-point testbed), we compute the partition function via the magnetization integral $Z = \int_{-1}^{1} e^{N \cdot s(m)}\, dm$ where $s(m)$ encodes entropy and energy. (This is not presented as the standard finite-volume nearest-neighbor 1D Ising partition function.) The Gaussian approximation is the saddle-point around $m = 0$.

N	$\rho = N^{1/3}$	δ (exact)	$2C/\rho^3$ ($C = 0.5$, selected comparison)	Below curve
8	2.0	-0.1102	0.125	✓
64	4.0	-0.0152	0.0156	✓
216	6.0	-0.00459	0.00463	✓
512	8.0	-0.00195	0.00195	✓
1000	10.0	-0.000998	0.00100	✓

Empirical rate: $|\delta| \sim \rho^{-2.97}$, confirming the $O(1/N) = O(\rho^{-3})$ finite-size correction.

The comparison constant $C = 0.5$ was selected for this table, not derived as a theorem constant. For this particular magnetization integral, the observed large-N coefficient is close to one in $\delta_N \sim -1/N$, which explains the near-saturation of the curve $2C/N = 1/N$. This statement is not about the standard nearest-neighbor 1D Ising partition function.

7.6 Poisson/Stirling Reparameterization

The Poisson saddlepoint approximation is a classical Stirling calculation, reparameterized here by $\rho = k^{1/3}$. It is not an application of Theorem 1. For $X \sim \text{Poisson}(\lambda)$:

$$\hat{p}(k) = \frac{\exp(k - \lambda - k\ln(k/\lambda))}{\sqrt{2\pi k}}$$

k	$\rho = k^{1/3}$	δ (exact)	$2C/\rho^3$ ($C = 0.1$, selected comparison)	Below curve
8	2.0	-0.01036	0.025	✓
30	3.1	-0.00277	0.0067	✓
125	5.0	-0.000666	0.0016	✓
500	7.9	-0.000167	0.0004	✓
1000	10.0	-0.0000833	0.0002	✓

Empirical rate: $|\delta| \sim \rho^{-3.00}$ (exactly).

This is the cleanest illustration: Stirling's expansion gives a Poisson saddlepoint correction of order $1/k$, and the chosen reparameterization $\rho = k^{1/3}$ rewrites it as $O(\rho^{-3})$. It is not evidence for a universal non-exponential-family theorem.

7.7 Summary of Numerical Results

Model	Cubic contribution?	Empirical rate	Selected comparison C	Numerical comparison
Quartic integral	No	$\rho^{-3.98}$	0.50	below curve at sampled $\rho \geq 1$
Half-line cubic	Yes	$\rho^{-2.98}$	0.60	below curve at sampled $\rho \geq 1$
Product cross-term	—	$\rho^{-7.96}$	—	identity evaluated at sampled $\rho \geq 1$
Magnetization toy model	Effective asymmetry	$\rho^{-2.97}$	0.50	below curve at sampled $N \geq 8$
Poisson saddlepoint	Stirling correction	$\rho^{-3.00}$	0.10	below curve at sampled $k \geq 8$

When the cubic term is absent in the quartic model, the observed correction decays faster than ρ^{-3}.

The half-line cubic and magnetization toy examples display numerically estimated $O(\rho^{-3})$ rates. The product cross-term in the quartic example displays twice the individual asymptotic exponent. These are illustrations, not proofs of the abstract certificate or domain mappings.

---

8. Formal Verification Summary

Verified surface	Declared contracts	Passed	Rejected
Scalar core algebra (platonic.py)	25	25	0
Domain proposals	0	0	0
Threshold analogies	0	0	0
Total	**25**	**25**	**0**

The historical 110-target development inventory is retained only in the machine-readable claim map for auditability. Eighty domain or threshold templates and five optional scalar duplicates are explicitly retired; they cannot contribute to a verified count.

The previous false green arose at the harness boundary, not from a proof of the counterexample: a failed automatic attempt was discarded, and an empty contract list then passed vacuously. The current scalar source uses an explicit 25-name inventory, raises if any declaration is missing, and passes all 25 replay contracts. A regression test pins the old counterexample and the fail-closed behavior.

Declaration of Generative AI and AI-assisted technologies in the writing process

During the preparation of this work the author used Cursor and its integrated AI models (Claude and GPT families) in order to assist with manuscript drafting, structural editing, and language polishing. After using this tool/service, the author reviewed and edited the content as needed and takes full responsibility for the content of the publication.

References

1. F. Fang and C. W. Oosterlee (2009). A novel pricing method for European options based on Fourier-cosine series expansions. *SIAM J. Sci. Comput.*, 31(2):826–848.
2. C. L. Fefferman (2000). Existence and smoothness of the Navier-Stokes equation. *Clay Mathematics Institute Millennium Prize Problems.*
3. J. Kaplan et al. (2020). Scaling laws for neural language models. *arXiv:2001.08361.*
4. J. F. Nash (1951). Non-cooperative games. *Annals of Mathematics*, 54(2):286–295.

5. C. Bayer, P. K. Friz, and J. Gatheral (2016). Pricing under rough volatility. *Quantitative Finance*, 16(6):887–904.

6. N. Bleistein and R. A. Handelsman (1986). *Asymptotic Expansions of Integrals.* Dover. Corrected republication of the 1975 Holt, Rinehart and Winston edition. ISBN 978-0-486-65082-1.

7. H. E. Daniels (1954). Saddlepoint approximations in statistics. *Annals of Mathematical Statistics*, 25(4):631–650.

8. L. Le Cam (1986). *Asymptotic Methods in Statistical Decision Theory.* Springer.

9. J. R. McClean, S. Boixo, V. N. Smelyanskiy, R. Babbush, and H. Neven (2018). Barren plateaus in quantum neural network training landscapes. *Nature Communications*, 9:4812.

10. R. K. Pathria and P. D. Beale (2011). *Statistical Mechanics.* Academic Press, 3rd edition.

11. W. H. Fleming and R. W. Rishel (1975). *Deterministic and Stochastic Optimal Control.* Springer.

12. A. Dembo and O. Zeitouni (1998). *Large Deviations Techniques and Applications.* Springer, 2nd edition.

13. K. G. Wilson and J. Kogut (1974). The renormalization group and the expansion. *Physics Reports*, 12(2):75–199.

14. A. W. van der Vaart (1998). *Asymptotic Statistics.* Cambridge University Press.

15. L. Hörmander (1990). *The Analysis of Linear Partial Differential Operators I.* Springer-Verlag, 2nd ed. ISBN 978-3-540-52345-1. Digital reprint doi:10.1007/978-3-642-61497-2.

16. A. L. Lewis (2000). *Option Valuation under Stochastic Volatility: With Mathematica Code.* Finance Press.

17. P. K. Friz, S. Gerhold, A. Gulisashvili, and S. Sturm (2011). On refined volatility smile expansion in the Heston model. *Quantitative Finance*, 11(8):1151–1164.

18. S. R. S. Varadhan (1984). *Large Deviations and Applications.* SIAM.

19. R. S. Ellis (1985). *Entropy, Large Deviations, and Statistical Mechanics.* Springer.

20. L. Tierney and J. B. Kadane (1986). Accurate approximations for posterior moments and marginal densities. *Journal of the American Statistical Association*, 81(393):82–86.

21. V. Spokoiny (2023). Dimension free nonasymptotic bounds on the accuracy of high-dimensional Laplace approximation. *SIAM/ASA Journal on Uncertainty Quantification*, 11(3):1044–1068. doi:10.1137/22M1495688. Earlier version: arXiv:2204.11038.

22. A. Katsevich (2025). Improved dimension dependence in the Bernstein–von Mises theorem via a new Laplace approximation bound. *Information and Inference: A Journal of the IMA*, 14:iaaf020. doi:10.1093/imaiai/iaaf020.

23. M. Kasprzak, R. Giordano, and T. Broderick (2025). How good is your Laplace approximation of the Bayesian posterior? Finite-sample computable error bounds for a variety of useful divergences. *Journal of Machine Learning Research*, 26(87):1–81.

Appendix A: Numerical Code

The numerical illustrations in Section 7 are reproducible from the companion numerical_examples.py script using NumPy and SciPy. The script requests an absolute quadrature tolerance of 10^{-12}; this is a solver setting, not a guarantee that every reported value has 10^{-12} absolute error.

Appendix B: Theorem Index

ID	Name	Statement (abbreviated)	File
T1	grade_tail_upper	$A - A_2 \le C/\rho^3$	platonic.py
T2	grade_tail_lower	$A - A_2 \ge -C/\rho^3$	platonic.py
T3	grade2_dominance	$A_2 > C/\rho^3 \Rightarrow A > 0$	platonic.py
T4	bound_improves	$\rho_2 > \rho_1 \Rightarrow C/\rho_2^3 < C/\rho_1^3$	platonic.py
T5	precision_from_rho	$\rho^3\varepsilon \ge C$ implies the grade-tail bound $C/\rho^3 \le \varepsilon$	platonic.py
T6	saddle_upper	$F \le F_g(1 + 2C/\rho^3)$	platonic.py
T7	saddle_lower	$F \ge F_g(1 - 2C/\rho^3)$	platonic.py
T8	excess_bounded	$F - F_g \le F_g \cdot 2C/\rho^3$	platonic.py
T9	deficit_bounded	$F_g - F \le F_g \cdot 2C/\rho^3$	platonic.py
T10	gaussian_positive	$K > 0, \sqrt{\det H} > 0 \Rightarrow F_g > 0$	platonic.py
T11	hessian_determines_gaussian	Equal stipulated scalar inputs give equal Gaussian reference values	platonic.py
T12	integral_near_gaussian_U	**Upper** multiplicative consequence of an assumed correction bound	platonic.py
T16	saddle_sandwich	Full sandwich bound	platonic.py
T17	precision_sandwich	Precision-parameterized sandwich	platonic.py
T18–20	Bound algebra	Monotonicity, C-monotonicity, conditional precision	platonic.py
T21–22	Conditional power algebra	$B^2 \le B$ and $B^3 \le B^2$ for $0 \le B \le 1$	platonic.py
T23–25	Conditional perturbation algebra	Power inequalities and two-term bounds; no series existence proof	platonic.py

ID	Name	Statement (abbreviated)	File
T26–28	Products	Ring identity and two cross-term bounds	platonic.py

The domain mappings and threshold analogies have no theorem identifiers and are absent from the formal inventory.