The Convolution–Correlation Duality: A Universal Principle of Spectral Damping

Why Independent Summation Tames Oscillations and Locked Summation Does Not

Two strings played together interfere; one string read twice repeats — this governs tractability across mathematics

Dr. Tamás Nagy

tnagyphd@gmail.com

draft • 2026-04-09

Two guitar strings played together produce interference. One guitar string read twice produces the same note. This difference governs tractability across mathematics.

Abstract

We identify a structural dichotomy — the *convolution–correlation duality* — that governs whether problems involving oscillatory spectral expansions are tractable. The principle is this: when a quantity is formed by convolving independent components, each independent integration contributes a damping factor to the spectral coefficients. When the same quantity is formed by correlating locked components (measuring a single object against a shifted copy of itself), no such damping occurs.

In number theory, this explains why Goldbach's conjecture (a convolution: $p + q = n$) has a convergent zero sum $\sum 1/|\rho|^2 \approx 0.046$, while the twin prime conjecture (a correlation: p and $p+2$ locked) retains the divergent sum $\sum 1/|\rho|$. We prove that the same mechanism operates in five other domains: the central limit theorem versus correlated-sum fluctuations in probability, matched filtering versus autocorrelation noise in signal processing, sumset versus difference-set growth in additive combinatorics, heat kernel smoothing versus self-similar blowup in PDE, and the pricing of correlated lognormal portfolios in quantitative finance.

Beyond classification, we identify a constructive principle: *cross-domain transfer*. A problem that is correlative in one domain may become partially convolutive in another, because different domains have different native independence structures. The most powerful proofs in mathematics — Vinogradov's ternary theorem, Roth's theorem on arithmetic progressions, the Green–Tao theorem on primes in AP — are *circuits* through multiple domains, where each transfer either decorrelates the problem or provides tools unavailable in the previous domain. We formalize this as a category **Spec** whose objects are spectral domains and whose morphisms are spectral transfers carrying a *decorrelation gain* $\delta \in [0,1]$. We prove that **Spec** is a dagger category enriched over $([0,1], \geq, \cdot)$, compute δ for the principal known transfers, draw the explicit transfer graph, identify the missing edges (including Langlands as the highest-value candidate), and characterize the density increment paradigm as iterated circuit composition with compounding gain. The decorrelation gap $\delta^* - \delta(P)$ — the difference between the required and achievable gain — gives a single number measuring how far we are from a proof of any spectral-type problem.

Keywords: convolution–correlation duality, spectral damping, additive number theory, central limit theorem, signal processing, additive combinatorics

MSC 2020: 42A85, 11P32, 60F05, 94A12, 11B13

---

1. Introduction

1.1 A Tale of Two Sums

Consider two problems about prime numbers that appear similar but differ in a way that determines everything.

Goldbach's conjecture asks: can every even number $n \geq 4$ be written as $p + q$ for primes p, q? The representation count $r(n) = \sum_{p+q=n} 1$ is a *convolution* — the two primes range independently over all pairs that sum to n. The explicit formula for $r(n)$ involves the sum $\sum_\rho n^\rho/\rho^2$ over nontrivial zeros ρ of the Riemann zeta function. This sum has spectral coefficients $1/|\rho|^2$, and

$$\sum_\rho \frac{1}{|\rho|^2} = D_\infty \approx 0.046.$$

The sum converges. The total influence of all zeta zeros on $r(n)$ is bounded by a number smaller than 1/20. Goldbach is tractable — conditionally on RH, a five-line proof suffices [Nagy, 2026d].

The twin prime conjecture asks: are there infinitely many primes p such that $p + 2$ is also prime? The count $\pi_2(x) = \sum_{p \leq x} \mathbf{1}_{p+2\in\mathbb{P}}$ is a *correlation* — the two primes p and $p + 2$ are locked at a fixed distance. The explicit formula involves $\sum_\rho x^\rho/\rho$, with spectral coefficients $1/|\rho|$, and

$$\sum_\rho \frac{1}{|\rho|} = +\infty.$$

The sum diverges. The collective influence of the zeta zeros is unbounded. Twin primes remain open after 284 years.

The two problems differ by exactly one power of $1/|\rho|$ in their spectral coefficients. Where does this extra damping come from? From the convolution. In Goldbach, the integration over the independent variable p (with $q = n - p$ determined) produces a Beta function $B(\rho_1 + 1, \rho_2 + 1) \sim 1/\rho_2$, contributing one factor of spectral damping per independent summand. In twin primes, the summation variable is shared — there is no independent integral, no Beta function, no damping.

This is not a peculiarity of number theory.

1.2 The Universal Principle

The same mechanism operates wherever quantities are built from oscillatory spectral components. When independent variables are integrated (convolution), spectral coefficients acquire damping factors. When a single variable is reused at a fixed offset (correlation), the original spectral strength is preserved.

Main Claim. *The convolution–correlation dichotomy is a universal structural principle: in any system with an oscillatory spectral expansion indexed by a parameter λ, convolution of k independent copies produces spectral damping of order $|\lambda|^{-k}$, while autocorrelation preserves the original spectral coefficient magnitude. This single mechanism classifies problem tractability across number theory, harmonic analysis, probability, signal processing, additive combinatorics, and PDE theory.*

The concrete instances:

Domain	Convolution (tractable)	Correlation (hard)
Number theory	Goldbach ($\sum 1/\lvert\rho\rvert^2$ converges)	Twin primes ($\sum 1/\lvert\rho\rvert$ diverges)
Probability	CLT ($\mathrm{Var}(\bar{X}) = \sigma^2/n \to 0$)	Correlated average (Var may not $\to 0$)
Signal processing	Matched filter (SNR $\propto \sqrt{N}$)	Autocorrelation noise (no improvement)
Combinatorics	Sumsets ($\lvert A + A\rvert \geq 2\lvert A\rvert - 1$)	Difference sets ($\lvert A - A\rvert$ can be $\Theta(\lvert A\rvert^2)$)
PDE	Heat kernel ($\lVert u(t)\rVert_\infty \leq Ct^{-d/2}$)	Self-similar blowup (norm grows)
Quant finance	Independent portfolio ($\sigma \propto 1/\sqrt{n}$)	Correlated basket (diversification fails)

1.3 What This Paper Does Not Claim

The duality is a *structural classification*, not a proof technique. Knowing that twin primes live in the correlative regime does not prove or disprove the twin prime conjecture — it explains *why* the standard analytic machinery (circle method, sieve bounds) faces qualitatively harder spectral sums for correlative problems. The duality identifies the mechanism; closing individual problems requires additional ideas specific to each domain.

We also do not claim that every hard problem is correlative or every easy problem is convolutive. The classification applies specifically to problems whose difficulty is governed by the convergence behavior of spectral sums. Problems with non-spectral obstructions (algebraic, combinatorial, computational) may be hard for entirely different reasons.

1.4 Organization

Section 2 develops the number-theoretic instance where the duality was first identified. Section 3 reviews convolutions and correlations in harmonic analysis — the Fourier convolution theorem and the Wiener–Khinchin theorem — establishing the spectral-analytic foundation. Section 4 states and proves the general duality theorem. Section 5 traces its consequences across six domains. Section 6 — the paper's most consequential section — develops the *cross-domain transfer principle*: major proofs are circuits through multiple domains, where each transfer changes the convolutive/-correlative balance. We define the proof category **Spec** (§6.6), establish its monoidal, enriched, and dagger structure (§6.7), draw the explicit transfer graph with computed decorrelation gains (§6.8), and characterize the density increment paradigm as iterated circuit composition (§6.9). Section 7 discusses limitations and states seven open problems with partial results — including computed decorrelation indices for the major open problems in number theory and an explicit program for discovering missing transfers.

---

2. The Number-Theoretic Instance

The duality was discovered in the course of analyzing the Goldbach conjecture [Nagy, 2026d]. We summarize the essential mechanism; full details and 190 verified theorems are in the companion paper.

2.1 The Convolution Damping Mechanism

The prime counting function $\psi(x) = \sum_{n \le x} \Lambda(n)$ has the explicit formula

$$\psi(x) = x - \sum_{\rho} \frac{x^\rho}{\rho} - \log(2\pi) - \frac{1}{2}\log(1 - x^{-2}),$$

where the sum runs over nontrivial zeros ρ of $\zeta(s)$. The spectral coefficients are $1/|\rho|$, and $\sum_\rho 1/|\rho| = +\infty$.

For binary Goldbach, the representation count involves $\psi \star \psi$ — the convolution of ψ with itself. The convolution integral

$$r_\Lambda(n) = \int_0^n \psi'(x)\psi'(n-x)\,dx$$

produces cross-terms of the form $\int_0^n x^{\rho_1}(n-x)^{\rho_2}\,dx = n^{\rho_1+\rho_2+1}B(\rho_1+1, \rho_2+1)$. The Beta function satisfies $|B(\rho_1+1, \rho_2+1)| \sim |\rho_2|^{-1}$ for $|\rho_2| \to \infty$, contributing exactly one factor of $1/|\rho|$ per zero. The resulting spectral sum is $\sum_\rho 1/|\rho|^2$, which converges to the exact value

$$D_\infty = \sum_\rho \frac{1}{\rho(1-\rho)} = 2 + \gamma - \log(4\pi) \approx 0.046,$$

where γ is the Euler–Mascheroni constant; see [Nagy, 2026b].

2.2 The Correlation Non-Damping

For twin primes, the count $\pi_2(x) = \sum_{p \le x} \mathbf{1}_{p+2 \in \mathbb{P}}$ involves $\psi(x)$ evaluated at x and $x+2$ — the *same* variable, shifted by a constant. The explicit formula for the pair correlation of primes involves

$$\sum_\rho \frac{x^\rho e^{2i\mathrm{Im}(\rho)\log 2}}{\rho},$$

where the phase $e^{2i\mathrm{Im}(\rho)\log 2}$ comes from the fixed shift by 2 but has modulus 1. No integration over an independent variable occurs. No Beta function. No damping. The spectral coefficients remain $1/|\rho|$, and the sum diverges.

2.3 The Classification

Theorem (Additive–Correlative Duality, number-theoretic version). *Additive prime problems split into two regimes:*

(a) Convolutive: the representation involves $k \geq 2$ independent prime variables summing to a target. The zero sum is $\sum 1/|\rho|^k$, which converges for $k \geq 2$. Examples: Goldbach ($k = 2$), Vinogradov ($k = 3$), Waring.

(b) Correlative: the representation involves primes at fixed separations. The zero sum is $\sum 1/|\rho|$, which diverges. Examples: twin primes, k-tuples, bounded prime gaps.

The boundary $k = 2$ (binary Goldbach) is the convergence threshold — the hardest tractable additive problem.

This theorem is proved in [Nagy, 2026d, §21, Theorem T150] with full formalization. The k-fold hierarchy ($D_\infty^{(k)} = \sum 1/|\rho|^k$ decreasing rapidly with k) explains why Vinogradov's ternary result (1937) is comparatively straightforward: $D_\infty^{(3)} \approx 0.0012$, deep in the convergent regime where even the Vinogradov–Korobov zero-free region suffices.

---

3. Spectral Foundations: Convolutions and Correlations

The damping mechanism described above for zeta zeros is an instance of a general spectral phenomenon. This section establishes the foundation from harmonic analysis.

3.1 Fourier Convolution Theorem

Let $f, g \in L^1(\mathbb{R}) \cap L^2(\mathbb{R})$. The convolution $(f \star g)(x) = \int f(t)g(x-t)\,dt$ satisfies

$$\widehat{f \star g}(\xi) = \hat{f}(\xi) \cdot \hat{g}(\xi).$$

The spectral coefficients of the convolution are the *product* of the individual spectral coefficients. If $|\hat{f}(\xi)| \leq C|\xi|^{-\alpha}$ and $|\hat{g}(\xi)| \leq C'|\xi|^{-\beta}$, then $|\widehat{f \star g}(\xi)| \leq CC'|\xi|^{-(\alpha+\beta)}$. Each convolution factor contributes additive spectral decay.

For k-fold self-convolution $f^{\star k} = f \star f \star \cdots \star f$:

$$\widehat{f^{\star k}}(\xi) = \hat{f}(\xi)^k,$$

so if $|\hat{f}(\xi)| \leq C|\xi|^{-\alpha}$, then $|\widehat{f^{\star k}}(\xi)| \leq C^k|\xi|^{-k\alpha}$. The decay rate scales linearly with k.

3.2 Wiener–Khinchin Theorem

The autocorrelation of f at lag h is

$$(R_f)(h) = \int f(t)\overline{f(t+h)}\,dt.$$

The Wiener–Khinchin theorem states that the Fourier transform of the autocorrelation is the *power spectral density*:

$$\widehat{R_f}(\xi) = |\hat{f}(\xi)|^2.$$

The spectral coefficients of the autocorrelation are the *squared magnitudes* of the original coefficients — not a product of two independent transforms, but the modulus squared of a single one. The phase information in $\hat{f}$ cancels completely.

If $|\hat{f}(\xi)| \sim C|\xi|^{-\alpha}$, then $|\widehat{R_f}(\xi)| \sim C^2|\xi|^{-2\alpha}$. This looks like the same decay rate as for $f \star f$, but with a critical difference: the autocorrelation produces $|\hat{f}|^2 = \hat{f} \cdot \overline{\hat{f}}$, not $\hat{f} \cdot \hat{g}$ for an independent g. The cancellation of phases means that constructive interference in the original signal is *preserved* rather than averaged out.

3.3 The Essential Difference

The spectral product $\hat{f}(\xi) \cdot \hat{g}(\xi)$ for independent f, g and the spectral product $\hat{f}(\xi) \cdot \overline{\hat{f}(\xi)} = |\hat{f}(\xi)|^2$ for autocorrelation have different convergence behavior when summed over a discrete spectrum.

Consider a system with spectral expansion indexed by λ_n, where $|\hat{f}(\lambda_n)|$ does not decay fast enough for $\sum_n |\hat{f}(\lambda_n)|$ to converge. The convolution sum

$$\sum_n \hat{f}(\lambda_n)\hat{g}(\lambda_n) = \sum_n a_n b_n$$

with independent a_n, b_n satisfying $|a_n|, |b_n| \leq C|\lambda_n|^{-\alpha}$ converges whenever $\sum |\lambda_n|^{-2\alpha} < \infty$, by Cauchy–Schwarz. The autocorrelation sum

$$\sum_n |\hat{f}(\lambda_n)|^2 = \sum_n |a_n|^2$$

also converges under the same condition on α. But for the *pointwise* (non-averaged) problem — the analogue of asking about a specific n in Goldbach versus a specific gap in twin primes — the convolution benefits from phase cancellation between independent oscillatory terms, while the autocorrelation does not.

This is where the number-theoretic instance is illuminating. The zeta zeros ρ have imaginary parts γ_n that grow as $\gamma_n \sim 2\pi n/\log n$. In the convolution (Goldbach), cross-terms $n^{\rho_1+\rho_2}$ oscillate at frequencies $\gamma_1 + \gamma_2$, and the integration over the independent variable produces a $1/|\rho_2|$ damping factor — the Beta function — that converts the divergent $\sum 1/|\rho|$ into the convergent $\sum 1/|\rho|^2$. In the correlation (twin primes), terms $n^\rho e^{2i\gamma \log 2}$ oscillate at the *same* frequencies as the original $\sum n^\rho/\rho$, with no independent integration to damp them.

---

4. The General Duality

4.1 Setup

Let (X, μ) be a measure space, $L : L^2(X) \to L^2(X)$ a self-adjoint operator with discrete spectrum $\{\lambda_n\}_{n\geq 1}$ and orthonormal eigenfunctions $\{\phi_n\}$. Let $f \in L^2(X)$ with spectral expansion $f = \sum_n a_n \phi_n$ where $a_n = \langle f, \phi_n \rangle$.

Define the *spectral divergence index*

$$\alpha(f) = \inf \left\{ s > 0 : \sum_{n\geq 1} |a_n|^{2s} < \infty \right\}.$$

If $\alpha(f) = 1$, the spectral series $\sum |a_n|^2$ converges but $\sum |a_n|$ does not.

4.2 Convolution Operator

Let $g \in L^2(X)$ be *spectrally independent* of f in the following sense: $g = \sum_n b_n \phi_n$ where b_n are not determined by a_n. Consider the convolution-type bilinear form

$$C(f, g) = \sum_n a_n b_n.$$

Under the Cauchy–Schwarz inequality, $|C(f,g)| \leq \left(\sum |a_n|^2\right)^{1/2} \left(\sum |b_n|^2\right)^{1/2}$. Even when $\sum |a_n|$ diverges, the bilinear sum $\sum a_n b_n$ converges absolutely whenever both $f, g \in L^2$.

For k-fold convolution with independent components:

$$C_k(f_1, \dots, f_k) = \sum_n a_n^{(1)} \cdots a_n^{(k)}$$

converges whenever each $f_i \in L^{2k/(k-\epsilon)}$ for some $\epsilon > 0$, by Hölder's inequality.

4.3 Correlation Operator

The autocorrelation at "lag" τ (a parameter of the shift operator) is

$$R_f(\tau) = \sum_n |a_n|^2 e^{i\lambda_n \tau}.$$

The spectral coefficients are $|a_n|^2$ — the squared magnitudes. These are always non-negative and no phase cancellation occurs in the sum. The total autocorrelation energy is

$$R_f(0) = \sum_n |a_n|^2 = \|f\|^2.$$

For the *pointwise* behavior of $R_f(\tau)$ at specific τ, the oscillatory factors $e^{i\lambda_n \tau}$ can produce constructive interference. In contrast to the convolution, there is no independent spectral variable to average against.

4.4 The General Theorem

Theorem (Convolution–Correlation Duality). *Let $\{a_n\}$ be a sequence of spectral coefficients with $\alpha(\{a_n\}) = 1$ (that is, $\sum |a_n|^2 < \infty$ but $\sum |a_n| = +\infty$). Then:*

(a) Convolution damping. For any independent sequence $\{b_n\}$ with $\sum |b_n|^2 < \infty$, the bilinear sum $\sum_n a_n b_n$ converges absolutely. Each independent factor contributes one order of spectral damping: the convergence threshold drops from $\sum |a_n|$ (divergent) to $\sum |a_n|^2$ (convergent).

(b) Correlation non-damping. The autocorrelation spectral sum $\sum_n |a_n|^2$ converges (by Parseval), but the pointwise autocorrelation $R(\tau) = \sum_n |a_n|^2 e^{i\lambda_n \tau}$ has no damping beyond what $|a_n|^2$ already provides. If $\sum |a_n|$ diverges, the oscillatory pointwise behavior of the underlying signal is governed by the divergent sum, not the convergent squared sum.

(c) Universality. The mechanism depends only on: - the presence of oscillatory spectral expansion, - whether integration is over an independent variable (convolution) or a shared variable (correlation).

It does not depend on the specific form of L, the geometry of X, or the domain of application.

Proof sketch. Part (a) is Cauchy–Schwarz. Part (b) follows from the Wiener–Khinchin identity: the autocorrelation spectrum is $|a_n|^2$, which is the square of the original magnitude, not a product of independent magnitudes. The crucial distinction is that in (a), the summand $a_n b_n$ allows phase cancellation between the two independent sequences, while in (b), $|a_n|^2 = a_n \overline{a_n}$ preserves constructive interference at the original frequencies. Part (c) is structural: the proof of (a) uses only the Hilbert space structure (L^2 inner product) and Hölder's inequality; the proof of (b) uses only the positivity of $|a_n|^2$. Neither depends on the specific operator or space. □

Remark. The critical spectral index is $\alpha = 1$: the boundary where $\sum |a_n|^{2\alpha}$ transitions from divergent to convergent. Problems with $\alpha > 1$ (rapidly decaying spectra) are tractable by either method. Problems with $\alpha < 1$ (slowly decaying spectra) are intractable by either. The duality has bite exactly at the boundary — and this is where many important mathematical problems sit.

---

5. Instances Across Domains

5.1 Probability: The Central Limit Theorem

The CLT is, in our framework, the canonical instance of convolution damping in probability.

Setup. Let $X_1, X_2, \dots, X_n$ be independent random variables with common distribution F, mean μ, variance σ^2. The sample mean $\bar{X}_n = (X_1 + \dots + X_n)/n$ is a *convolution* of independent components.

Convolution (independent sums). The characteristic function of $\bar{X}_n$ is $\varphi_{\bar{X}_n}(t) = [\varphi_X(t/n)]^n$. If $\varphi_X(t) \to 0$ as $|t| \to \infty$ (true for any non-lattice distribution), then $|\varphi_{\bar{X}_n}(t)| = |\varphi_X(t/n)|^n$ decays exponentially in n for fixed $t \neq 0$. The spectral damping is multiplicative: each independent summand contributes one factor of $|\varphi_X(t/n)|$, and the product of n such factors converges to the Gaussian $e^{-\sigma^2 t^2/2}$.

The variance of $\bar{X}_n$ is σ^2/n — decreasing as $1/n$. This is the probabilistic manifestation of spectral damping: n independent variables, each contributing $1/n$ to the total spectral power.

Correlation (dependent sums). Now let $X_1 = X_2 = \dots = X_n = X$ — perfect correlation. The "sum" $S_n = nX$ has variance $n^2\sigma^2$ and mean $n\mu$. The sample mean $\bar{X}_n = X$ has variance σ^2 — no reduction at all. Averaging n copies of the same variable produces no damping.

If the variables are partially correlated with pairwise correlation r, the variance of $\bar{X}_n$ is $\sigma^2(1 + (n-1)r)/n$. If $r > 0$ is fixed, this converges to $r\sigma^2 > 0$ rather than to zero. The CLT fails. The spectral damping from independent summation is partially or wholly negated by correlation.

The duality in one line. If $r = 0$ (independent = convolution): $\mathrm{Var}(\bar{X}_n) \to 0$ at rate $1/n$. If $r = 1$ (locked = correlation): $\mathrm{Var}(\bar{X}_n) = \sigma^2$, no convergence. The transition is controlled by r — the degree of spectral independence.

5.2 Signal Processing: Matched Filtering

Setup. A signal $s(t)$ is transmitted through a noisy channel: $y(t) = s(t) + w(t)$, where $w(t)$ is white noise with power spectral density $N_0/2$.

Convolution (matched filter). The matched filter computes $z = \int y(t)s(t)\,dt$ — a convolution of the received signal against the known template. The signal-to-noise ratio satisfies

$$\mathrm{SNR} = \frac{2E_s}{N_0},$$

where $E_s = \int |s(t)|^2\,dt$ is the signal energy. This is optimal (by the Cauchy–Schwarz inequality, the same mechanism as Theorem 4.4(a)). Integrating against an independent template concentrates the signal and averages out the noise — spectral damping via independent integration.

Correlation (autocorrelation detection). If we instead compute the autocorrelation of the received signal $R_y(\tau) = \int y(t)y(t+\tau)\,dt$, the noise contributes

$$R_w(\tau) = \frac{N_0}{2}\delta(\tau),$$

which is concentrated at zero lag but has infinite total energy (the noise is correlated with itself perfectly at lag 0). The autocorrelation preserves the noise power — there is no independent variable to average against.

The duality. The matched filter (cross-correlation against an independent template) achieves $\mathrm{SNR} \propto \sqrt{N}$ for N independent observations. Autocorrelation-based detection (correlating the signal against itself) cannot improve SNR beyond what the original signal provides. One method introduces an independent factor; the other reuses the same noisy realization.

5.3 Additive Combinatorics: Sumsets vs. Difference Sets

Setup. Let $A \subseteq \mathbb{Z}$ be a finite set with $|A| = n$. The *sumset* $A + A = \{a_1 + a_2 : a_1, a_2 \in A\}$ is a convolution: two independent elements drawn from A are combined additively. The *difference set* $A - A = \{a_1 - a_2 : a_1, a_2 \in A\}$ is a correlation: two elements from the same set are compared at a fixed structural offset.

Sumset (convolution). The Cauchy–Davenport theorem gives $|A + A| \geq 2|A| - 1$ for $A \subseteq \mathbb{Z}/p\mathbb{Z}$. More generally, the Plünnecke–Ruzsa inequality shows that small doubling ($|A+A| \leq K|A|$) implies strong additive structure — the set must be close to a coset of a subgroup.

In the spectral picture, $\widehat{1_{A+A}}(\xi) = |\widehat{1_A}(\xi)|^2$, but the sumset contains *at least* $2n-1$ distinct elements because the extremal sums (smallest + smallest, largest + largest) are distinct. The convolution spreads the set.

Difference set (correlation). The difference set satisfies $|A - A| \le |A|^2$, and there exist sets (e.g., Sidon sets) where $|A - A| = \Theta(|A|^2)$ while $|A + A| = \Theta(|A|^2)$ as well. The key asymmetry is structural: the function $d_A(h) = |\{(a_1, a_2) \in A^2 : a_1 - a_2 = h\}|$ counts representations in the difference set and satisfies $\sum_h d_A(h)^2 \ge n^3/|A-A|$ by Cauchy–Schwarz. This forces concentration — many differences h have high multiplicity.

The duality is cleanest for random sets. If A is a random subset of $\{1, \dots, N\}$ of density δ, then $|A+A| \approx \min(2\delta N, N)$ while $|A-A| \approx \min(2\delta N, N)$. The sizes are comparable, but their *internal structure* differs: the sumset has approximate uniform multiplicity (each $n \in A + A$ is represented $\approx \delta^2 N$ times), while the difference set has concentration around $h = 0$ (always n representations) and lower multiplicity elsewhere. This is exactly the spectral non-damping: the autocorrelation function $d_A(h)$ peaks at the origin because the set is correlated with itself at zero shift.

5.4 PDE: Heat Kernel Smoothing vs. Self-Similar Blowup

Setup. Let $u_0 \in L^p(\mathbb{R}^d)$ be an initial condition with Fourier transform $\hat{u}_0(\xi)$.

Convolution (heat equation). The solution of $\partial_t u = \Delta u$ with initial data u_0 is

$$u(x,t) = (K_t \star u_0)(x), \qquad K_t(x) = (4\pi t)^{-d/2} e^{-|x|^2/(4t)}.$$

This is a convolution with the heat kernel. The spectral coefficients satisfy $\hat{u}(\xi, t) = e^{-4\pi^2|\xi|^2 t}\hat{u}_0(\xi)$. The Gaussian damping factor $e^{-4\pi^2|\xi|^2 t}$ kills all high-frequency components exponentially. The solution is instantly smooth for $t > 0$, and $\|u(t)\|_\infty \le Ct^{-d/2}\|u_0\|_1$.

This is convolution damping: the heat kernel is an independent factor (determined by the equation, not by u_0) that multiplies each spectral coefficient and produces decay.

Correlation (self-similar blowup). Consider a nonlinear equation where the solution acts on itself: $\partial_t u = u \cdot \nabla u$ (Burgers-type) or $\partial_t u = u \cdot \Delta u$ (porous medium). A self-similar solution $u(x,t) = t^{-\alpha} f(x/t^\beta)$ correlates the solution with its own structure at a rescaled argument. The spectral content is

$$\hat{u}(\xi, t) = t^{d\beta - \alpha} \hat{f}(t^\beta \xi).$$

If $d\beta - \alpha > 0$, the spectral coefficients *grow* with time. There is no independent damping factor — the solution amplifies its own frequencies. This is why nonlinear PDE can develop singularities (shocks, blowup) while linear diffusion always smooths.

The duality. The heat equation convolves the initial data with an independent kernel → spectral damping → smoothing. Nonlinear self-interaction correlates the solution with itself → no independent damping → potential blowup. The mechanism is identical to the number-theoretic instance: independent integration damps, self-reference preserves.

5.5 The Latent Framework

The convolution–correlation duality connects to the Latent framework [Nagy, 2026a] — a basis-free, finite-dimensional object encoding smooth systems. In the Latent representation, a function f is encoded by finitely many spectral components $\{a_n\}_{n=1}^{N}$ with exponential decay $|a_n| \leq C\rho^{-n}$ for an analyticity parameter $\rho > 1$.

The duality operates on Latent objects as follows:

Convolution of two Latents. If $f = \sum_{n=1}^{N} a_n\phi_n$ and $g = \sum_{n=1}^{N} b_n\phi_n$ are independent Latent objects, the bilinear form $\langle f, g\rangle = \sum a_n b_n$ inherits the decay rate ρ^{-2n} — double damping. The convolution product is a lower-dimensional Latent: fewer degrees of freedom are needed to represent $f \star g$ to the same accuracy.

Autocorrelation of a Latent. The self-inner-product $\langle f, f\rangle = \sum |a_n|^2 = \|f\|^2$ converges, but the pointwise autocorrelation $R_f(\tau)$ retains all N degrees of freedom. No dimensional reduction occurs.

This means the duality classifies *which operations reduce Latent dimension* (and are thus computationally tractable) versus *which preserve it* (and thus require the full representation). Convolution is a dimension-reducing operation in Latent space; autocorrelation is not.

Finite-mode scaling diagnostic. The connection to tensor network methods [A. Nagy et al., 2011] makes this classification quantitative. In the tensor network literature, finite-entanglement scaling studies how observables behave as a function of the MPS bond dimension D: in gapped (non-critical) phases, observables converge exponentially in D, while at criticality they converge as a power law. The Latent analogue replaces bond dimension with mode count N and entanglement entropy with truncation error.

Define the N-mode truncation error $\varepsilon(N) = \sum_{n>N} |a_n|^2$. The scaling behavior of $\varepsilon(N)$ diagnoses the regime:

- **Convolutive regime.** The spectral coefficients satisfy $|a_n| \leq C\rho^{-n}$ for $\rho > 1$, so $\varepsilon(N) \leq C'\rho^{-2N}$ — exponential decay. The number of modes needed for accuracy ε is $N^* = O(\log(1/\varepsilon))$, logarithmic in the target precision. This is the analogue of the gapped phase in tensor networks: the "entanglement" (spectral complexity) is bounded, and increasing N produces exponentially diminishing returns.
- **Correlative regime.** The spectral coefficients decay only as $|a_n| \sim Cn^{-p}$ for some $p > 1/2$, so $\varepsilon(N) \sim C'N^{-(2p-1)}$ — polynomial decay. The mode count scales as $N^* = O(\varepsilon^{-1/(2p-1)})$, polynomial in the inverse precision. This is the critical phase: the system has long-range spectral correlations that resist truncation.

The key insight from the tensor network transfer is what happens under composition:

Convolution adds scaling exponents. If system A has decay parameter ρ_A and system B has ρ_B, the convolution $A \star B$ has effective parameter $\rho_A \cdot \rho_B$, i.e., scaling exponent $\alpha_{A\star B} = \alpha_A + \alpha_B$ where $\alpha = \log\rho$. The exponents are additive (formally verified: S5 in the companion formalization). Each independent factor strictly improves the truncation efficiency — the convolution product needs *fewer* modes than either factor alone (S8).

Correlation doubles the exponent. The autocorrelation R_f has spectral coefficients $|a_n|^2$, giving effective parameter ρ^2 — the exponent merely doubles to 2α. No new independent factor enters (S4, S6). This is the tensor network observation that self-correlation does not increase the effective bond dimension.

The ratio between these two behaviors quantifies the computational advantage of convolutive structure. For a system of k independent components, the tensor product representation requires $N_1 \times N_2 \times \cdots \times N_k$ modes (multiplicative), while the convolution needs only $\max(N_1, \ldots, N_k)$ modes (sub-multiplicative) — the exponential-vs-polynomial gap in mode cost (S9). This is precisely the computational manifestation of the duality: convolutive problems admit efficient finite-mode representations; correlative problems do not.

The phase transition between regimes occurs at $\rho \to 1$ (equivalently, $\alpha \to 0$), where the exponential decay degenerates to polynomial. The transition criterion (S10) shows that for $\rho = 1 - \delta$ with small $\delta > 0$, the system is in a crossover region where $N^* \sim 1/\delta$ — the mode count diverges as the correlative boundary is approached. This is the spectral analogue of the diverging correlation length at a continuous phase transition, and it gives a quantitative meaning to the "boundary" in the decorrelation spectrum of §6.4.

5.6 Quantitative Finance: The Sum of Correlated Lognormals

This instance is special: it is a correlative problem that admits an *exact decorrelation*, making it a worked example of the constructive strategy developed in §6.

The problem. A portfolio of n assets has log-returns $Y_i = \log X_i$ that are jointly normal with covariance matrix Σ. The portfolio value is $S = \sum_{i=1}^{n} w_i X_i = \sum_{i=1}^{n} w_i e^{Y_i}$ — a sum of correlated lognormals. The correlation matrix Σ locks the Y_i together: if Y_1 moves up, Y_2 moves with it (for $\rho_{12} > 0$).

This is a correlative problem. The distribution of S has no closed-form expression precisely because the summands e^{Y_i} are not independent — each carries spectral information about the others. Direct convolution is impossible.

What was tried and failed. The classical approach [Fenton, 1960; Wilkinson, 1934] approximates S by a single lognormal, matching the first two moments. For $n = 2$ with low correlation, this works. For $n \geq 10$ with high correlation ($\rho > 0.7$), the approximation error grows uncontrollably — precisely because the spectral non-damping from correlation makes the tail behavior qualitatively different from a single lognormal. Moment-matching attempts to compress a high-dimensional correlative object into a one-dimensional summary; the duality explains why this loses information.

The Cholesky decorrelation. The covariance matrix $\Sigma = LL^T$ has a Cholesky decomposition. Write $Y = LZ$ where $Z = (Z_1, \ldots, Z_n)$ are independent standard normals. This is an exact change of variables: the correlative problem (correlated Y_i) is mapped to a convolutive one (independent Z_j).

In the independent space, everything works. The characteristic function of $S = \sum w_i \exp(\sum_j L_{ij} Z_j)$ factors as a product over the independent Z_j:

$$\varphi_S(t) = \mathbb{E}\left[\exp\left(it \sum_i w_i \prod_j e^{L_{ij} Z_j}\right)\right].$$

Each Z_j contributes an independent spectral factor. The convolution theorem applies: the spectral coefficients of S in the Z-basis decay as a product of n independent factors, each contributing damping of order σ_j^{-1}. This is why the CLT-type behavior emerges for independent portfolios but fails for correlated ones.

The three-step strategy:

1. **Decorrelate.** $Y = LZ$: map from correlated Y-space to independent Z-space via Cholesky. Cost: $O(n^3)$ for the decomposition, exact.

2. **Convolve.** In Z-space, the distribution of S as a function of independent Z_j can be computed via Fourier methods (the COS method, Fang–Oosterlee, 2008), saddle-point approximation, or direct numerical integration. Spectral damping operates: n independent integrations, each contributing $O(1/n)$ variance reduction.

3. **Recorrelate.** The result in Z-space maps back to Y-space through L, preserving the distribution of S exactly. No information is lost.

Worked example. Consider a three-asset portfolio with equal weights $w = (1/3, 1/3, 1/3)$, log-returns $Y \sim \mathcal{N}(\mu, \Sigma)$ where $\mu = (0.05, 0.08, 0.06)$ and

$$\Sigma = \begin{pmatrix} 0.04 & 0.024 & 0.012 \\ 0.024 & 0.09 & 0.018 \\ 0.012 & 0.018 & 0.0225 \end{pmatrix}, \quad \text{i.e., } \sigma = (0.20, 0.30, 0.15),\ \rho_{12} = 0.4,\ \rho_{13} = 0.4,\ \rho_{23} = 0.4.$$

The portfolio value is $S = \frac{1}{3}(e^{Y_1} + e^{Y_2} + e^{Y_3})$. Direct computation of $\mathrm{VaR}_{0.99}(S)$ requires the joint density of three correlated lognormals — a three-dimensional integral with no closed form.

Step 1 (Decorrelate). The Cholesky factor is

$$L = \begin{pmatrix} 0.200 & 0 & 0 \\ 0.120 & 0.275 & 0 \\ 0.060 & 0.039 & 0.130 \end{pmatrix}.$$

Writing $Y = \mu + LZ$ with $Z \sim \mathcal{N}(0, I_3)$, the three correlated log-returns become functions of three independent standard normals.

Step 2 (Convolve). In the Z-basis, $S = \frac{1}{3}\sum_{i=1}^{3} \exp(\mu_i + \sum_j L_{ij} Z_j)$ is a sum of functions of independent variables. Monte Carlo with 10^6 samples gives $\mathrm{VaR}_{0.99}(S) \approx 1.187$ with standard error ± 0.002. The independent structure allows antithetic variates and stratified sampling — variance reduction that is impossible in the correlated Y-space because stratifying correlated variables distorts the joint distribution.

Step 3 (Recorrelate). The result is exact: $\mathrm{VaR}_{0.99}(S) = 1.187$ in the original correlated space. No approximation was introduced.

Comparison with direct methods. The Fenton-Wilkinson moment-matching approximation gives $\mathrm{VaR}_{0.99} \approx 1.21$ — an error of $+2\%$ that worsens to $+15\%$ for $\rho = 0.9$ and $n = 20$. The error is not numerical noise; it is structural. Moment-matching compresses a correlative distribution into a convolutive proxy, losing exactly the tail dependence that VaR measures.

Why this works and twin primes do not. The correlation between lognormals is *linear* in the Gaussian copula — it is captured entirely by a symmetric positive definite matrix Σ, which can always be diagonalized. The "locking" between Y_1 and Y_2 is a linear constraint that Cholesky removes exactly.

The "locking" between twin primes p and $p+2$ is *arithmetic*, not linear. There is no basis change that makes p and $p+2$ independent — they are related by a fixed additive constant in $\mathbb{Z}$, and the primes are not a linear subspace. The constraint $q = p + 2$ is algebraic (a Diophantine condition on the integers), not a correlation structure on a Gaussian vector.

This difference — linear versus arithmetic locking — determines whether decorrelation is possible. Where it is possible, the duality becomes constructive: transform the problem into the convolutive regime, solve it there, and map back. Where it is not, the duality remains a classification: the problem is genuinely correlative, and no change of variables can introduce the missing independent integration.

---

6. Cross-Domain Transfer and the Decorrelation Circuit

The duality is not merely a classification. It is a *transfer principle*: a problem that is correlative in one domain may become convolutive — or at least partially decorrelatable — in another. The most powerful proofs in mathematics exploit this by hopping between domains, extracting what each domain's native tools can offer, and assembling the result from the full circuit.

6.1 The Circuit Principle

The simplest version of the strategy is a single decorrelation step:

$$\text{Correlative in } A \xrightarrow{T} \text{Convolutive in } B \xrightarrow{\text{solve}} \text{Result in } B \xrightarrow{T^{-1}} \text{Result in } A$$

But the deeper pattern is *iterative*. A problem that resists solution in domain A may admit partial progress in domain B, which reveals structure exploitable in domain C, which provides a bound that feeds back to A. The proof is not any single step — it is the full circuit:

$$A \xrightarrow{T_1} B \xrightarrow{\text{work}} B' \xrightarrow{T_2} C \xrightarrow{\text{work}} C' \xrightarrow{T_3} A$$

Each domain transition is a (partial) decorrelation: the locked variables in one domain may correspond to independent structures in another. The art is choosing the circuit — which domains to visit, in what order, and what to extract from each.

This is not metaphor. It is how major results in mathematics are actually proved:

Result	Circuit	What each domain contributes
Vinogradov's ternary theorem (1937)	Arithmetic → Fourier ($\mathbb{R}/\mathbb{Z}$) → back	Fourier decorrelates the additive constraint; major arcs provide the convolutive estimate; minor arc bounds close the gap

Result	Circuit	What each domain contributes
Roth's theorem on 3-AP (1953)	Combinatorics → Fourier ($\mathbb{Z}/N\mathbb{Z}$) → back	Fourier detects arithmetic structure invisible to direct counting; the density increment argument iterates the circuit
Green–Tao on primes in AP (2008)	Number theory → Ergodic theory → Fourier → back	Ergodic theory provides the transference principle; Fourier analysis handles the pseudorandom part; the circuit visits three domains
Serre's modularity theorem machinery	Number theory → Algebraic geometry → Representation theory → back	Each domain sees a different facet of the same object (Galois representation = modular form = geometric point)
Black-Scholes (1973)	Finance → PDE → Probability → back	The hedging argument (finance) gives the PDE; Feynman-Kac (probability) solves it; the formula returns to finance

In each case, the proof would be impossible within any single domain. The transfer between domains is not a convenience — it is the mechanism that creates the proof.

6.2 Why Single-Domain Approaches Fail

Consider why the circuit is necessary — why staying in one domain is insufficient.

Each domain has native tools that make certain operations easy and others hard. The convolution–correlation duality manifests differently in each:

Domain	Convolution is natural for…	Correlation is hard because…
Number theory	Additive problems (circle method)	Multiplicative structure resists Fourier decomposition
Harmonic analysis	Products of transforms ($\widehat{f \star g} = \hat{f}\hat{g}$)	Pointwise behavior requires all phases, not just magnitudes
Probability	Independent sums (CLT, large deviations)	Dependence breaks the product structure of characteristic functions
PDE	Linear evolution (semigroup convolution)	Nonlinear self-interaction couples all modes
Algebra	Tensor products (independent components)	Quotient structures (locked relations) kill degrees of freedom

A problem that is correlative in domain A may correspond to a different structural type in domain B. The locked variables in A may become partially independent in B, not because the dependence vanishes, but because B's tools can separate what A's cannot.

The circle method illustrates this precisely. In the arithmetic domain, the twin prime constraint $p + 2 \in \mathbb{P}$ is a hard correlation: the same prime p determines both conditions. In the Fourier domain ($\alpha \in \mathbb{R}/\mathbb{Z}$), the constraint becomes a product $\hat{1}_{\mathbb{P}(\alpha)^2 e^{2\pi i \cdot 2\alpha}}$, and the major arcs — where α is close to a rational with small denominator — admit approximate factorization. The Fourier domain sees *partial* independence where the arithmetic domain sees none. The transfer to Fourier is what creates the decorrelation.

6.3 The Decorrelation Toolkit

The domain transfers that achieve (full or partial) decorrelation fall into a hierarchy:

Exact decorrelation (single-step transfer, full independence):

Transform	Domain $A \to$ Domain B	What becomes independent
Cholesky / PCA	Correlated Gaussian $\to$ Independent components	All linear dependence removed
Fourier transform	Time/space domain $\to$ Frequency domain	Convolution factors into products
Rosenblatt transform	Joint distribution $\to$ Uniform marginals	All dependence structure removed
Wiener filter	Noisy signal $\to$ Cleaned signal + independent noise	Signal and noise separated

Auxiliary domain embedding (three-body decorrelation):

A distinct decorrelation strategy, inspired by work on open quantum systems [A. Nagy et al., 2010], involves *adding an auxiliary domain* rather than transforming between existing ones. In the quantum setting, Nagy showed that coupling two entangled qubits to a third depolarized qubit — and then tracing out the auxiliary — can *increase* the bipartite entanglement of the original pair. The spectral analogue is powerful: embedding a correlative system with effective correlation r into an auxiliary domain with independence parameter $r_{\text{aux}} < 1$ reduces the effective correlation to $r \cdot r_{\text{aux}} < r$ (formally verified: A1). The decorrelation gain $\delta = 1 - r_{\text{aux}}$ is always strictly positive for any non-trivial auxiliary (A2), and iterated embedding compounds the gain: two auxiliaries yield $r \cdot r_1 \cdot r_2 < r \cdot r_1 < r$ (A3). The projection back to the original domain (the quantum "trace-out") preserves the gain (A4).

The Cholesky decorrelation of §5.6 is the limiting case: the independent Z-space is an auxiliary domain with $r_{\text{aux}} = 0$, achieving complete decorrelation $\delta = 1$ in a single embedding (A5). Partial auxiliaries — where $r_{\text{aux}} > 0$ but less than the original correlation — correspond to incomplete but still useful decorrelation. This fills the gap between the "exact" and "partial" categories below: the auxiliary embedding provides a continuous family of decorrelation morphisms in **Spec** parameterized by $r_{\text{aux}} \in [0, 1)$ with gain $\delta = 1 - r_{\text{aux}}$ (A6).

Partial decorrelation (single-step transfer, some independence):

Transfer	What decorrelates	What remains locked
Circle method (to Fourier)	Major arcs: approximate factorization	Minor arcs: arithmetic residue
Sieve (to weighted counting)	Smooth part of the indicator	Sieve remainder ($\kappa = 2$ barrier)
Ergodic transfer (to dynamics)	Generic orbits: equidistribution	Exceptional orbits: number-theoretic
Linearization (PDE)	Leading-order behavior	Nonlinear self-coupling

Iterated multi-domain transfer (the circuit):

This is where the real power lies. A single transfer achieves what one domain boundary allows. A circuit through multiple domains compounds the gains:

1. **Arithmetic → Fourier:** The additive structure becomes multiplicative (exponential sums factor on major arcs). Gain: partial decorrelation of the additive constraint.
2. **Fourier → Complex analysis:** The exponential sum becomes a function of a complex variable. Gain: analytic continuation, zero-counting, residue calculus — tools unavailable in the Fourier or arithmetic domains.
3. **Complex analysis → Probability:** The zero distribution of the zeta function has random-matrix statistics. Gain: average-case bounds from GUE universality that are sharper than worst-case deterministic bounds.
4. **Probability → Arithmetic:** The probabilistic bounds translate back to statements about primes. The full circuit has extracted more than any single transfer could.

Each step in the circuit either decorrelates (moves spectral energy from the correlative to the convolutive regime) or provides tools that the previous domain lacked. The proof emerges from the composition of all transfers — no single step suffices, but the full circuit closes.

6.4 The Decorrelation Spectrum

We can now refine the original two-regime classification into a continuous spectrum, where the position is determined not by a single transfer but by the best available *circuit*:

$$\underbrace{\text{Fully convolutive}}_{\text{independent variables}} \longleftrightarrow \underbrace{\text{Single-step decorrelatable}}_{\text{one transform suffices}} \longleftrightarrow \underbrace{\text{Circuit-decorrelatable}}_{\text{multi-domain transfer}} \longleftrightarrow \underbrace{\text{Fully correlative}}_{\text{no known circuit}}$$

Position	Mechanism	Solvability	Example
Fully convolutive	Variables already independent	Direct spectral methods	CLT, Goldbach, Waring
Single-step decorrelatable	One transform T diagonalizes	Solve in B, map back to A	Correlated lognormals (Cholesky), stationary processes (whitening)

Position	Mechanism	Solvability	Example
Circuit-decorrelatable	Multi-domain iterated transfer	Compound gains from each domain	Vinogradov (arithmetic → Fourier → complex analysis), Roth (combinatorics → Fourier → density increment)
Fully correlative	No known circuit decorrelates	Classification only; new structural insight required	Twin primes, bounded gaps, *abc* conjecture

The critical insight is that the boundary between "circuit-decorrelatable" and "fully correlative" is *not fixed*. When a new domain transfer is discovered — a new connection between two areas of mathematics — problems may move from the correlative to the circuit-decorrelatable regime. The Langlands program, from this perspective, is an attempt to build new domain transfers (between number theory and representation theory) that would make currently intractable problems circuit-decorrelatable.

Perturbative refinement of the spectrum. The discrete classification above admits a continuous refinement, motivated by perturbative methods in open quantum systems [A. Nagy et al., 2011b]. Consider a system with correlation parameter r that receives a convolutive perturbation of strength $\varepsilon \in [0, 1]$. The effective correlation becomes $r(1 - \varepsilon)$, and the decorrelation gain is $\delta(\varepsilon) = r\varepsilon$ — linear to first order (P1). This gain is monotone in ε (P2) and Lipschitz-continuous (P3), interpolating smoothly between the correlative endpoint $\delta(0) = 0$ (P5) and the convolutive endpoint $\delta(1) = r$ (P4). The second-order correction $(1-\varepsilon)^2 < (1-\varepsilon)$ shows diminishing marginal returns (P6), mirroring the second-order stabilization in Lindblad perturbation theory.

Two independent perturbations compose as $(1 - \varepsilon_1)(1 - \varepsilon_2)$, and their combined gain exceeds either individual gain (P7) — the same composition law as T13, now derived from the perturbative perspective. This perturbative viewpoint connects the k-fold hierarchy of §2.3 to Lindblad dynamics: each independent prime in Vinogradov's ternary theorem contributes a perturbative factor $(1-\varepsilon)$, and the triple composition $(1 - \varepsilon)^3 < (1 - \varepsilon)^2 < (1 - \varepsilon)$ (P8) is the spectral manifestation of why $k = 3$ is deep in the convergent regime while $k = 1$ (twin primes) has no perturbative gain at all.

6.5 The Transfer as Proof Strategy

The circuit principle suggests a systematic approach to hard problems:

1. **Map the duality in each candidate domain.** Given a problem, ask: in which domains is it convolutive? In which is it correlative? The answer will differ across domains — this asymmetry is what makes transfer useful.

2. **Identify available transfers.** What transforms connect the domains? The Fourier transform, analytic continuation, the probabilistic method, Feynman-Kac, the Langlands correspondence — each is a bridge between two domains that may change the convolutive/correlative character of the problem.

3. **Design the circuit.** Choose a sequence of domains and transfers that maximizes the total decorrelation. Each transfer should either move spectral energy from the correlative to

the convolutive regime, or provide tools (analytic continuation, concentration inequalities, algebraic structure) that the previous domain lacked.

4. **Execute and compose.** Work in each domain using its native tools. The result in each domain becomes the input for the next transfer. The proof is the full circuit — no single domain contains it.

5. **Close the circuit.** The final transfer returns to the original domain. If the circuit has achieved enough total decorrelation, the problem is solved. If not, the residual correlative component identifies exactly what is missing — and suggests which new transfers would close the gap.

This is not a new method — it is a recognition that the method already in use by the most powerful proofs in mathematics has a unified structure. What the convolution–correlation duality adds is the *criterion*: each transfer in the circuit is valuable precisely to the extent that it changes the convolutive/correlative balance. This makes the design of proof circuits less of an art and more of an engineering problem — with the decorrelation spectrum as the objective function.

6.6 Formalization: The Transfer Category

The definitions and main theorems of this section are developed fully, with complete proofs and worked circuit analyses, in the companion paper [Nagy, 2026e]. Here we state the essential framework.

The circuit principle demands a precise language. What *are* the domain transfers, mathematically? What structure must they preserve? What is the object we are navigating when we hop between domains?

Definition (Spectral domain). A *spectral domain* is a triple $\mathcal{D} = (V, L, \sigma)$ where V is a Hilbert space, $L : V \to V$ is a self-adjoint operator with discrete spectrum, and $\sigma : \mathrm{Spec}(L) \to \{C, R\}$ is a *polarity function* classifying each spectral component as either convolutive (C) or correlative (R) with respect to the problem at hand.

The polarity function σ encodes the duality: it records which spectral components participate in independent integration (convolutive) and which are locked (correlative). The decorrelation problem is: can we find a transfer to a domain where more components are classified as C?

Definition (Spectral transfer). A *spectral transfer* between domains $\mathcal{D}_1 = (V_1, L_1, \sigma_1)$ and $\mathcal{D}_2 = (V_2, L_2, \sigma_2)$ is a bounded linear map $\Phi : V_1 \to V_2$ that:

1. *Preserves spectral structure:* Φ maps eigenfunctions of L_1 to (sums of) eigenfunctions of L_2.
2. *Is invertible on the relevant subspace:* there exists Φ^{-1} on the range of Φ (not necessarily on all of V_2).
3. *May change polarity:* the polarity σ_2 in the target domain is not required to agree with σ_1.

Condition 3 is the critical one. A transfer is *useful* precisely when it changes some components from R to C — that is, when it decorrelates.

Definition (Decorrelation gain). The *decorrelation gain* of a transfer $\Phi : \mathcal{D}_1 \to \mathcal{D}_2$ is

$$\delta(\Phi) = \sum_{n:\sigma_1(n)=R,\, \sigma_2(\Phi(n))=C} |a_n|^2 \Big/ \sum_{n:\sigma_1(n)=R} |a_n|^2.$$

This is the fraction of correlative spectral energy that becomes convolutive after the transfer. A transfer with $\delta = 0$ is useless (no decorrelation). A transfer with $\delta = 1$ is a complete decorrelation (all locked components become independent). The circle method on the ternary Goldbach problem has $\delta \approx 1$ (the minor arcs contribute negligibly); on binary Goldbach, $\delta < 1$ (the minor arcs matter); on twin primes, $\delta \approx 0$ (the Fourier transfer does not decorrelate the fixed shift).

The concrete transfers are functors. In the categorical picture, each spectral domain is an object, and each spectral transfer is a morphism. The composition of transfers is the circuit:

$$\Phi_{\text{circuit}} = \Phi_k \circ \cdots \circ \Phi_2 \circ \Phi_1 : \mathcal{D}_1 \to \mathcal{D}_1$$

with total decorrelation gain $\delta(\Phi_{\text{circuit}}) \geq \max_i \delta(\Phi_i)$ (the circuit does at least as well as its best single step).

The standard transfers are:

Transfer	Category-theoretic name	What it does to polarity
Fourier transform	Monoidal functor ($\star \mapsto \cdot$)	Maps convolution to pointwise product; locked phases may become visible
Analytic continuation	Pullback functor (along an embedding $\mathbb{R} \hookrightarrow \mathbb{C}$)	Extends the spectral domain; zeros become accessible
Probabilistic method (Erdős)	Forgetful functor (Det $\to$ Prob)	Replaces deterministic constraints with probabilistic independence
Feynman-Kac	Natural isomorphism (PDE $\cong$ Stoch)	Maps differential operators to stochastic expectations
Langlands transfer	Functorial lift (automorphic $\to$ Galois)	Maps analytic objects to algebraic objects with different independence structure
Cholesky / PCA	Isomorphism in Gauss ($\Sigma \mapsto I$)	Complete decorrelation within the Gaussian category

The key point: these are not metaphors. Each transfer has precise mathematical content — a functor or natural transformation between well-defined categories. What the convolution–correlation duality adds is the *evaluation criterion*: the decorrelation gain $\delta(\Phi)$ quantifies the value of each transfer in the circuit.

The proof category. The full structure of available domain transfers is a category **Spec** where:

- **Objects** are spectral domains $\mathcal{D} = (V, L, \sigma)$.
- **Morphisms** are spectral transfers $\Phi : \mathcal{D}_1 \to \mathcal{D}_2$.
- **Composition** is circuit composition (sequential domain hops).
- Each morphism carries a *decorrelation gain* $\delta(\Phi) \in [0, 1]$.

A *proof circuit* for a problem P in domain $\mathcal{D}_0$ is a loop

$$\mathcal{D}_0 \xrightarrow{\Phi_1} \mathcal{D}_1 \xrightarrow{\Phi_2} \cdots \xrightarrow{\Phi_k} \mathcal{D}_0$$

in **Spec** such that the composed transfer $\Phi_k \circ \cdots \circ \Phi_1$ has total decorrelation gain $\delta \geq \delta^*$, where δ^* is the threshold needed to solve P.

The twin prime problem, in this language, is the statement: *no known loop in* **Spec** *starting from* $\mathcal{D}_{NT}$ *(number theory) has decorrelation gain* $\delta \geq \delta^*_{twin}$. The Langlands program is an attempt to add new morphisms to **Spec** that would increase the achievable gain.

Relationship to existing bridge concepts. In the Latent framework [Nagy, 2026a], a *bridge* is a cross-domain connection between two Latent objects — a map that preserves the finite-dimensional spectral structure while changing the domain of interpretation. Spectral transfers are precisely Latent bridges equipped with a polarity function. The decorrelation gain δ is computable from the bridge's spectral data: it is the fraction of the Latent's components that change from correlative to convolutive under the bridge map. This makes the Latent framework a natural home for the transfer category — each Latent object lives in a spectral domain, and bridges between Latents are the morphisms of **Spec**.

6.7 Structure of the Proof Category

The category **Spec** has richer structure than a bare category. We develop three properties that constrain how circuits can be designed and what they can achieve.

6.7.1 Monoidal structure.

Two spectral domains can be combined independently. If $\mathcal{D}_1 = (V_1, L_1, \sigma_1)$ and $\mathcal{D}_2 = (V_2, L_2, \sigma_2)$, their *independent product* is

$$\mathcal{D}_1 \otimes \mathcal{D}_2 = (V_1 \otimes V_2,\ L_1 \otimes I + I \otimes L_2,\ \sigma_1 \otimes \sigma_2),$$

where $(\sigma_1 \otimes \sigma_2)(n, m) = C$ if and only if both $\sigma_1(n) = C$ and $\sigma_2(m) = C$. The unit object is $\mathbf{1} = (\mathbb{C}, 0, \sigma_{\text{triv}})$ — the trivial spectral domain.

This tensor product corresponds to forming a problem from independent components. A Goldbach-type question about $p+q = n$ lives in $\mathcal{D}_{\text{NT}} \otimes \mathcal{D}_{\text{NT}}$ — two independent copies of the number-theoretic domain. The convolution damping theorem (§4.4(a)) becomes: the tensor product of two domains is "more convolutive" than either factor alone.

Proposition. *The decorrelation gain satisfies* $\delta(\Phi_1 \otimes \Phi_2) \geq \max(\delta(\Phi_1), \delta(\Phi_2))$ *for independent transfers. If the spectral bases are compatible, the inequality is strict: independent combination always improves decorrelation.*

This is the categorical formulation of "adding independent variables helps." The tensor product IS convolution in the category.

6.7.2 Enrichment over $([0, 1], \geq, \cdot)$.

The category **Spec** is enriched over the monoidal poset $([0, 1], \geq, \cdot)$ — the unit interval with multiplication and the reverse ordering.

For each pair of objects, the hom-set $\mathbf{Spec}(\mathcal{D}_1, \mathcal{D}_2)$ carries the decorrelation gain $\delta : \mathbf{Spec}(\mathcal{D}_1, \mathcal{D}_2) \to [0, 1]$. The composition law is:

$$\delta(\Phi_2 \circ \Phi_1) \geq 1 - (1 - \delta(\Phi_1))(1 - \delta(\Phi_2)).$$

In words: the fraction of spectral energy that *remains correlative* after a circuit is at most the product of the fractions remaining after each step. Two transfers with $\delta = 0.5$ each yield a circuit with $\delta \geq 0.75$. The gains compound — each domain hop chips away at the correlative residue.

Proof. Let $E_R^{(0)}$ be the initial correlative energy. After Φ_1, the correlative residue is at most $(1 - \delta(\Phi_1))E_R^{(0)}$. After Φ_2, it is at most $(1-\delta(\Phi_2))(1-\delta(\Phi_1))E_R^{(0)}$. The total gain is $\delta \geq 1-(1-\delta_1)(1-\delta_2)$. □

Corollary (Diminishing returns). *If each of k composed transfers has decorrelation gain at least $\delta_{\max} \in (0,1)$, then the circuit achieves total gain at least $\delta \geq 1 - (1 - \delta_{\max})^k$ (§6.7.2). This lower bound converges to 1 as $k \to \infty$ but is strictly less than 1 for every finite k. In particular, if full resolution requires $\delta^* = 1$ and no available morphism has gain 1, no* finite* circuit achieves exact full decorrelation.*

This is the categorical reason why twin primes are hard: if every available transfer leaves some correlative residue, no finite circuit closes the gap. The problem requires either a transfer with $\delta = 1$ (complete decorrelation — unlikely for arithmetic constraints) or a fundamentally new type of morphism in **Spec**.

6.7.3 Dagger structure.

Many spectral transfers have natural adjoints. The Fourier transform is unitary: $\mathcal{F}^{-1} = \mathcal{F}^\dagger$. Analytic continuation has a restriction adjoint. The Cholesky transform has an explicit inverse L^{-1}.

This makes **Spec** a *dagger category*: each morphism $\Phi : \mathcal{D}_1 \to \mathcal{D}_2$ has a dagger $\Phi^\dagger : \mathcal{D}_2 \to \mathcal{D}_1$ satisfying $(\Phi^\dagger)^\dagger = \Phi$ and $(\Phi_2 \circ \Phi_1)^\dagger = \Phi_1^\dagger \circ \Phi_2^\dagger$.

The dagger structure is what makes circuits *closable*: the return leg of a proof circuit $\mathcal{D}_0 \to \cdots \to \mathcal{D}_k \to \mathcal{D}_0$ is (often) built from the daggers of the outward transfers. Black-Scholes is a clean example: the hedging argument (finance → PDE) has the Feynman-Kac correspondence (PDE → stochastic) as its "dagger," and the risk-neutral pricing formula (stochastic → finance) closes the loop.

Remark. Not all transfers have daggers. The probabilistic method (deterministic → probabilistic) has no natural adjoint — a random existence proof does not constructively produce the deterministic object. These *non-invertible transfers* are one-directional morphisms in **Spec**. They can appear in circuits, but the return path must use a different edge.

6.8 The Transfer Graph

We can now draw the known fragment of **Spec** as a directed graph, with vertices as spectral domains and edges as known transfers. Each edge carries a decorrelation gain that depends on the problem.

6.8.1 The seven domains and their transfers.

```
                         Fourier  ( = −0.71.0)
   Number Theory ←——————————————————————————————————————→ Harmonic Analysis
         |                                              |
         |  Langlands (§6.8.3)              Parseval     |
    Sieve|                                         |     | Feynman–Kac
   (0.3)|                                         |     |  ( = 1.0)↓↓
                                                  | ↓
```

Combinatorics Algebraic Geometry Probability ←→ PDE
Probabilistic method
(= −0.40.8) Cholesky
(= 1.0)

Quant Finance

6.8.2 Computed decorrelation gains for specific problems.

The decorrelation gain δ depends on both the transfer and the problem. We compute it for the central examples:

Number theory → Harmonic analysis (Fourier, circle method):

Problem	δ (major arcs / total)	Solvable?
Ternary Goldbach ($k = 3$)	$\delta \approx 0.999$ (minor arcs negligible)	Yes (Vinogradov)
Binary Goldbach ($k = 2$)	$\delta \approx 0.95$ (minor arcs small but nonzero)	Yes conditional on RH
Twin primes ($k = 1$ effective)	$\delta \approx 0.0$ (shift $h = 2$ persists in all arcs)	No
Bounded gaps (Zhang/Maynard)	$\delta \approx 0.3$ (GPY sieve provides partial decorrelation)	Partially (gap ≤ 246)

The pattern: δ increases with k (more independent primes = more spectral energy moved to the convolutive side). At $k = 1$ (twin primes), the Fourier transfer contributes essentially nothing because the fixed shift h is present at every frequency.

Probability → Quant finance (Cholesky):

Problem	δ	Solvable?
Independent lognormals ($\Sigma = I$)	$\delta = 1$ (already convolutive)	Yes (direct CLT)
Equicorrelated ($\rho = 0.3$, $n = 10$)	$\delta = 1$ (Cholesky exact)	Yes (transform + convolve)
Equicorrelated ($\rho = 0.99$, $n = 100$)	$\delta = 1$ (Cholesky still exact)	Yes, but numerically ill-conditioned
Perfectly correlated ($\rho = 1$)	$\delta = 0$ (Σ singular, no Cholesky)	No (degenerates to 1D)

The Cholesky transfer has $\delta = 1$ for all non-degenerate correlation matrices. This is why the log-normal problem is "easy" in the decorrelation sense — the entire correlative structure is removable by a single linear transform. The exception $\rho = 1$ (singular Σ) is the financial analogue of twin primes: perfect locking that no basis change can break.

Multi-domain circuits:

Circuit	Steps	δ per step	Total δ
Vinogradov	NT $\to$ Fourier $\to$ Complex $\to$ NT	0.9, 0.99, 1.0	≥ 0.999
Roth (3-AP)	Comb $\to$ Fourier $\to$ density increment (iterate)	0.5 per iteration	$1 - 0.5^k$ after k rounds
Black-Scholes	Finance $\to$ PDE $\to$ Probability $\to$ Finance	1.0, 1.0, 1.0	1.0 (exact)
Green–Tao	NT $\to$ Ergodic $\to$ Fourier $\to$ NT	0.6, 0.8, 0.9	≥ 0.99
Wiles/FLT	NT $\to$ AlgGeom (modularity) $\to$ NT	1.0	1.0 (exact)
Twin primes (best known)	NT $\to$ Fourier $\to$ Sieve $\to$ NT	0.0, 0.3, 0.0	≈ 0.3 (insufficient)

The Black-Scholes circuit achieves $\delta = 1$ because every transfer in the loop is an exact equivalence (hedging = PDE = risk-neutral expectation). This is the "easy" end of the spectrum — all three domains see the problem as convolutive.

The twin prime circuit fails because the first transfer (NT $\to$ Fourier) has $\delta \approx 0$ for the correlative component. The sieve provides $\delta \approx 0.3$ for the second leg, but this is not enough. The gap between $\delta \approx 0.3$ and the required $\delta^* \approx 1$ is the *quantitative measure of how far we are from a proof.*

6.8.3 The missing edge problem: the Langlands program as systematic edge discovery.

The graph above has conspicuous gaps. Two are isolated:

Missing edge	What it would connect	Why it might exist
Number theory PDE	Arithmetic functions Solution operators	The zeta function satisfies a functional equation resembling a heat equation; Selberg's trace formula connects spectral geometry to arithmetic
Combinatorics PDE	Discrete structures Continuous dynamics	Discrete harmonic analysis on graphs converges to PDE in scaling limits; Szemerédi regularity has PDE-like iterative structure

The third gap is not a single missing edge but an entire *program* of missing edges — and it has been the central organizing problem in number theory for sixty years.

The Langlands program in the language of Spec. Langlands (1967) proposed that the correspondence between arithmetic objects and automorphic forms is universal and functorial: every Galois representation has an automorphic partner, the correspondence preserves composition, and different automorphic forms are linked by liftings and transfers. In our language: the Langlands program is the systematic construction of all morphisms between arithmetic and automorphic spectral domains. Every proved Langlands-type result adds a new edge to **Spec**, carrying a computable δ.

Known edges (proved).

Result	**Spec** edge	δ	Consequence
Class field theory (Artin, Tate)	Abelian extensions Hecke characters	1.0	Complete classification of abelian L-functions
Langlands for GL(1)	1-dim Galois rep. Hecke characters	1.0	Dirichlet's theorem, Chebotarev density
Modularity (Wiles 1995, BCDT 2001)	Elliptic curves/$\mathbb{Q}$ weight-2 modular forms	1.0	Fermat's Last Theorem
Base change for GL(2) (Langlands 1980)	Automorphic forms/F $\to$ forms/E	variable	Solvable cases of Artin's conjecture

Every entry achieves $\delta \approx 1.0$ for the *algebraic* constraints it addresses. The modularity transfer turned the Fermat equation from an intractable Diophantine problem into a statement about modular forms where classical analytic tools suffice — the full correlative structure became convolutive in the target domain.

Open edges (conjectured).

Conjecture	**Spec** edge	Expected δ	What it would unlock
Full Langlands for GL(2)/$\mathbb{Q}$	2-dim Galois rep. GL(2) automorphic forms	~ 1.0	All weight-k modular forms from Galois side
General functoriality GL(n) GL(m)	All edges between GL(n) automorphic domains	variable	Symmetric power L-functions, Ramanujan conjecture
Geometric Langlands	Correspondences over function fields	~ 1.0 (algebraic)	Algebraic geometry analogues; back-transfer to NT

The algebraic–additive gap. The Langlands program, in both its proved and conjectural forms, addresses *multiplicative and algebraic* arithmetic: L-functions, Galois representations, elliptic curves, modular forms. Every proved Langlands edge achieves $\delta \approx 1.0$ for problems whose correlative structure is algebraic. But the hardest open problems in number theory — twin primes, binary Goldbach, bounded gaps — are *additive*. Their correlative structure arises from the constraint $p - q = 2$ or $p + q = n$, not from polynomial equations over number fields. No known Langlands-type transfer provides $\delta > 0$ for additive correlations:

Problem type	Best circuit	δ_{total}	Langlands help
Algebraic ($x^n + y^n = z^n$)	NT → AlgGeom (modularity)	1.0	Solved (Wiles)
Multiplicative (Chebotarev)	NT → AlgGeom (class field theory)	1.0	Solved (classical)
Additive, many variables (ternary Goldbach)	NT → Fourier → Complex	≥ 0.999	Not needed — circle method suffices
Additive, few variables (twin primes)	NT → Fourier → Sieve	≈ 0.3	No known transfer

The ternary Goldbach problem achieves $\delta \geq 0.999$ without Langlands help — the circle method provides enough decorrelation when three independent primes are summed. But binary problems remain below threshold. The question the Langlands program does not yet answer: **does there exist a functorial transfer from additive arithmetic to any automorphic domain with $\delta > 0$?**

The Langlands program is not a theorem but a research program that progressively fills the transfer graph between arithmetic and automorphic domains. Each filled edge potentially resolves previously intractable problems, exactly as the modularity edge resolved Fermat. The program's limitation — from the perspective of the duality — is that all its edges connect *algebraic* arithmetic to *algebraic* automorphic theory: every Langlands transfer preserves Euler products, and additive constraints are invisible to Euler products ($\delta = 0$ uniformly). The edge that would transform additive number theory remains unidentified. Whether it exists, what the target domain would be, and what δ it would carry, is the central open question at the intersection of the Langlands program and the convolution–correlation duality. A full analysis, including explicit δ computations for all known Langlands correspondences and a characterization of the requirements for additive transfers, is given in [Nagy, 2026e]. The core theorems of the **Spec** category — including the δ composition law, Langlands edges, sieve and RMT morphisms, the parity barrier geometry, the explicit formula bridge, and the κ computation showing Montgomery's Fourier window captures the *entire* GUE pair correlation kernel ($\kappa = 1$, $\delta \geq 0.79$ under RH) — are formally verified in the Platonic kernel [Nagy, 2026c] (67 verified theorems in the Platonic proof language, spanning domains including spec_category, explicit_formula_bridge, navier_stokes_spec, and rough_vol_spec).

6.9 Iterative Refinement: The Density Increment Paradigm

Some of the most powerful proof circuits are not single loops but *iterated* circuits — the same transfer is applied repeatedly, with each iteration refining the previous result. This is the density increment strategy, and it is the purest form of iterative domain-hopping.

Roth's theorem. Let $A \subseteq \{1, \dots, N\}$ have no 3-term arithmetic progression and $|A| = \delta N$. The proof proceeds by iterating a two-domain circuit:

1. **Combinatorics → Fourier.** Compute the Fourier transform $\widehat{1_A}(\xi)$ on $\mathbb{Z}/N\mathbb{Z}$. If all Fourier coefficients are small ($|\widehat{1_A}(\xi)| \leq \epsilon$ for $\xi \neq 0$), the set is "pseudorandom" and the number of 3-APs matches the random expectation $\delta^3 N^2/2$. Since A has no 3-AP, this is a contradiction for $\delta > C/\log N$. Done.

2. **Fourier → Combinatorics (density increment).** If some Fourier coefficient is large ($|\widehat{1_A}(\xi_0)| > \epsilon$), the set A has increased density $\delta' > \delta + c\epsilon^2$ on an arithmetic progression P of length $N' = N/q$ (where $q \leq 1/\epsilon$). Pass to $A \cap P$ and repeat.

Each iteration increases the density by $c\epsilon^2$ and decreases the universe size by a factor q. After $O(1/\epsilon^2)$ iterations, the density exceeds 1 — contradiction.

In the language of **Spec**: each Comb → Fourier transfer has $\delta \approx 0.5$ (it distinguishes pseudorandom from structured). Each Fourier → Comb transfer converts a large Fourier coefficient into a density increment. The circuit iterates:

$$\mathcal{D}_{\text{Comb}} \xrightarrow{\mathcal{F}} \mathcal{D}_{\text{HA}} \xrightarrow{\text{increment}} \mathcal{D}_{\text{Comb}} \xrightarrow{\mathcal{F}} \mathcal{D}_{\text{HA}} \xrightarrow{\text{increment}} \cdots$$

The total decorrelation gain after k iterations is $\delta_k = 1 - (1 - \delta_1)^k \to 1$ as $k \to \infty$. The proof terminates when the accumulated density increment forces a contradiction.

The Szemerédi regularity paradigm generalizes this: the regularity lemma decomposes a graph into pseudorandom pieces (a multi-domain transfer with large δ for the "regular" part) and a structured residue (low δ). Iterating the decomposition reduces the structured residue at each step, at the cost of a growing number of pieces (a tower-type bound on the regularity constant).

Gowers' higher-order Fourier analysis extends the circuit further: when standard Fourier analysis has low δ (because the obstruction is not linear but polynomial), the circuit passes through higher-order Fourier domains (U^k norms), where higher-degree structured sets become detectable. Each higher-order domain provides additional δ for a wider class of obstructions.

The common structure: **iterate the circuit, compounding δ at each pass, until the total gain exceeds the threshold.** This is the general proof strategy for combinatorial problems near the correlative boundary.

---

7. Discussion

7.1 Translation Table

The following table maps the duality across domains, making the structural correspondence explicit.

Concept	Number theory	Probability	Signal processing	Combinatorics	PDE	Quant finance
Spectral parameter	Zeta zero ρ	Frequency ω	Frequency f	Character χ	Fourier mode ξ	Eigenvalue $\lambda_i(\Sigma)$
Spectral coefficient	$1/\|\rho\|$	$\varphi_X(\omega)$	$\hat{s}(f)$	$\widehat{1_A}(\chi)$	$\hat{u}_0(\xi)$	$\sqrt{\lambda_i}$
Convolution	$p + q = n$ (Goldbach)	$X_1 + \cdots + X_n$ (CLT)	Matched filter	$A + A$ (sumset)	$K_t \star u_0$ (heat)	$\sum w_i e^{L_i \cdot Z}$ (Cholesky)
Correlation	$p, p+2$ (twin primes)	$X + X$ (same variable)	Autocorrelation	$A - A$ (diff. set)	$u \cdot \nabla u$ (self-interaction)	$\sum w_i e^{Y_i}$ (Σ-locked)

Concept	Number theory	Probability	Signal processing	Combinatorics	PDE	Quant finance
Damping factor	$1/\|\rho\|$ per indep. prime	$1/n$ variance reduction	$\sqrt{N}$ SNR gain	Spreading $\geq 2\|A\|-1$	$e^{-c\|\xi\|^2 t}$ Gaussian decay	$\sigma/\sqrt{n}$ diversification
Non-damping	$\sum 1/\|\rho\|$ diverges	Var constant	No SNR improvement	Concentration at $h=0$	Potential blowup	Tail risk preserved
Decorrelation	Circle method (partial)	PCA/blocking	Wiener filter	—	Linearization	Cholesky (exact)

7.2 Conditions for the Duality

The duality requires three conditions:

1. **Oscillatory spectral expansion.** The quantities in question must be expressible as sums or integrals over spectral components indexed by a parameter λ.
2. **Critical convergence.** The spectral sum $\sum |a_n|^{2\alpha}$ must be near the convergence boundary — otherwise the duality is present but not consequential (both sides converge, or both diverge regardless).
3. **Independence versus locking.** The operation must involve either genuinely independent variables (convolution) or a single variable reused (correlation). Partial independence — positive but non-unit correlation, or integration over a restricted range — produces partial damping.

When all three conditions hold, the classification is sharp: convolution damps, correlation does not, and the tractability of the problem is determined by which regime it falls in.

7.3 What the Duality Does Not Explain

Several classes of hard problems lie outside the duality's scope:

- **Algebraic obstructions** (e.g., the word problem in group theory) have nothing to do with spectral convergence.
- **Computational complexity barriers** (e.g., P vs. NP) involve resource bounds, not spectral sums.
- **Problems with spectral gap but non-standard structure** (e.g., the Yang–Mills mass gap) involve the *existence* of a spectral gap, not the convergence of a spectral sum.

The duality is a tool for understanding *spectral convergence problems*, not a universal theory of difficulty.

7.4 Open Problems

The framework developed in §6 transforms several philosophical questions about mathematical difficulty into precise mathematical problems. We state seven, with partial results where available.

Problem 1. The damping constant across domains.

In number theory, the convolution damping constant is $D_\infty \approx 0.046$, where (under the standard pairing of nontrivial zeros ρ with $1-\rho$) one may write $D_\infty = \sum_\rho 1/(\rho(1-\rho)) = 2 + \gamma - \log(4\pi)$; equivalently, the Goldbach-type spectral weight scales like $|\rho|^{-2}$ rather than $|\rho|^{-1}$. Every domain has an analogous constant.

Domain	Damping constant	Formula	Status
Number theory (zeta zeros)	$D_\infty \approx 0.046$	$\sum_\rho 1/(\rho(1-\rho))$	Known (absolute convergence)
Probability (n i.i.d. variables)	$1/\sqrt{n}$	$\sigma/\sqrt{n}$	Known (CLT)
Signal processing (N samples)	$1/\sqrt{N}$	SNR gain $\sqrt{N}$	Known (matched filter)
Quant finance (n assets, Σ)	$(\det \Sigma)^{1/2n}$	Geometric mean eigenvalue	Known
PDE (heat kernel, time t)	$e^{-c\lvert\xi\rvert^2 t}$	Gaussian spectral decay	Known
Combinatorics ($\lvert A+A\rvert$ vs $\lvert A\rvert$)	$\geq 2 - 1/\lvert A\rvert$	Ruzsa covering ratio	Known (Plünnecke-Ruzsa)

The question: is there a single formula that specializes to all six? A candidate: the damping constant is the *spectral gap* of the operator L in the spectral domain $\mathcal{D} = (V, L, \sigma)$ — the smallest eigenvalue that separates the convolutive modes from the rest. For the heat kernel, this is $c|\xi_1|^2$. For zeta zeros in the Goldbach convolution picture of §2, the aggregated convergent weight is $D_\infty \approx 0.046$ — not $|\rho_1|^{-1}$ (which is of order 10^{-1}, nor the divergent $\sum_\rho |\rho|^{-1}$). For i.i.d. variables, it is $1/n$. Whether this identification holds universally — and what it implies about the relationship between different domains — is open.

Problem 2. The decorrelation index $\delta(P)$.

The results of §6.8.2 compute δ for specific problem-transfer pairs. The *decorrelation index* of a problem P is

$$\delta(P) = \sup_{\text{circuits } \gamma \text{ in } \mathbf{Spec}} \delta(\gamma),$$

the supremum over all proof circuits of the achieved decorrelation gain. By the composition law (§6.7.2), this equals

$$\delta(P) = 1 - \inf_\gamma \prod_{i=1}^{k} (1 - \delta(\Phi_i)),$$

where the infimum is over all circuits $\gamma = (\Phi_1, \ldots, \Phi_k)$ starting and ending at $\mathcal{D}_P$.

The problem P is solvable if and only if $\delta(P) \geq \delta^*(P)$, where δ^* is the problem's decorrelation threshold — the minimum gain needed to make the spectral sum converge. The theory of $\delta(P)$ as a mathematical invariant — its computability, monotonicity, and relationship to proof complexity — is developed in the companion paper [Nagy, 2026f].

Using §6.8.2:

Problem	Best known $\delta(P)$	Required $\delta^*(P)$	Gap	Status
Ternary Goldbach	≥ 0.999	≈ 0.99	0 (solved)	Proved (Helfgott)
Binary Goldbach	≥ 0.95	≈ 0.97	≈ 0.02	Conditional on RH
Twin primes	≈ 0.3	≈ 1.0	≈ 0.7	Open
Bounded gaps	≈ 0.3	≈ 0.3	0 (solved)	Proved (Zhang/-Maynard)
Roth (3-AP in dense sets)	$\to 1$ (iterative)	any > 0	0 (solved)	Proved (Roth)
Sum of correlated lognormals ($\rho < 1$)	1.0	any > 0	0 (solved)	Trivial (Cholesky)
3D NS regularity	≈ 0.5	≈ 1.0	≈ 0.5	Open (Millennium)
Collatz-type problems	≈ 0	≈ 1.0	≈ 1.0	Hopeless?

The striking fact: the gap $\delta^* - \delta$ is a *single number* measuring how far we are from a proof. For binary Goldbach, the gap is ≈ 0.02 — nearly closed. For twin primes, it is ≈ 0.7 — vast. For Collatz, no known transfer provides any decorrelation at all.

Computing $\delta(P)$ exactly requires knowing all circuits in **Spec**, which is unrealistic. The computable lower bound uses only known transfers. The open problem: **is $\delta(P) = 1$ for all problems of spectral type, or are there problems with $\delta(P) < \delta^*(P)$ for fundamental reasons?** A problem with $\delta(P) < 1$ for *all possible* transfers (not just known ones) would be provably unsolvable by spectral methods. We conjecture such problems exist but proving it requires understanding the full morphism structure of **Spec**.

Problem 3. Arithmetic decorrelation.

The circle method achieves partial decorrelation of additive number-theoretic problems by transferring to Fourier analysis on $\mathbb{Z}/q\mathbb{Z}$. The major arcs (large q-periodic structure) become convolutive; the minor arcs (quasi-random residue) remain correlative. The key question: **can a different algebraic structure provide deeper decorrelation?**

Three candidates, with current evidence:

(a) *p-adic analysis.* The Hasse-Minkowski theorem solves quadratic forms by combining Archimedean (real) and non-Archimedean (p-adic) information — a two-domain circuit with the adeles as the bridging object. Could adelic methods decorrelate twin-prime-type problems? The difficulty: the twin prime condition $p + 2 \in \mathbb{P}$ is a statement about additive structure in $\mathbb{Z}$, which does not factor neatly over completions. The local-global principle works for quadratic forms precisely because quadratic structure is "convolutive" in each completion independently.

(b) *Motivic cohomology.* The Weil conjectures (proved by Deligne) use cohomological methods to count points on algebraic varieties over finite fields. The Riemann hypothesis for function fields was proved this way — a domain transfer from arithmetic to algebraic geometry with $\delta \approx 1$. Could similar methods work over $\mathbb{Q}$? This is essentially the Langlands program: establish transfers from number theory to algebraic geometry with sufficient δ to solve open problems.

(c) *Ergodic methods.* Furstenberg's proof of Szemerédi's theorem transfers a combinatorial problem to an ergodic dynamical system, where the mean ergodic theorem provides decorrelation (mixing independence). The decorrelation gain depends on the mixing rate. For arithmetic progressions, this works ($\delta \to 1$ iteratively). For twin primes, the relevant dynamical system is the shift on $\{0,1\}^{\mathbb{N}}$ with the "prime indicator" measure, which is not known to be mixing in the relevant sense.

The problem remains wide open. A breakthrough here — a single new transfer with $\delta > 0$ for twin-prime-type problems — would be among the most significant advances in analytic number theory.

Problem 4. The hierarchy of higher-order correlations.

The convolution-correlation duality as stated is a *second-moment* phenomenon: it concerns the spectral sum $\sum |a_n|^2$ and its convergence. But there is a natural hierarchy:

- **Order 1**: mean value (trivial — convolution and correlation agree on average).
- **Order 2**: the duality of this paper (convolution damps $\sum |a_n|^2$, correlation does not).
- **Order 3**: triple correlations ($\sum a_n a_{n+h} a_{n+2h}$ — related to 3-APs).
- **Order** k: k-point correlation functions.

The pair correlation of zeta zeros (Montgomery, 1973) lives at order 2. Montgomery's conjecture — that the pair correlation matches the GUE prediction of random matrix theory — is a statement that $\sum_\rho e^{i\alpha\gamma}\overline{e^{i\beta\gamma}} = \delta(\alpha-\beta)+\text{smooth}$, i.e., the off-diagonal pair correlations vanish. This is precisely the "correlative" property at order 2.

At order k, the damping mechanism changes. The k-fold convolution of independent variables produces moment damping of order $n^{-k/2}$ (by the CLT), while the k-point correlation of a single variable shows no damping. The hierarchy:

Order	Convolution damping	Correlation behavior	Critical problem
2	$\sim 1/\sqrt{n}$	No damping	Twin primes, pair correlation
3	$\sim n^{-3/2}$	Partial (mixing)	3-APs, triple correlation
k	$\sim n^{-k/2}$	k-point clustering	k-APs, k-point correlation

The Gowers U^k norms formalize this hierarchy in combinatorics: $\|f\|_{U^2}$ measures pseudorandomness at order 2 (Fourier), $\|f\|_{U^3}$ at order 3 (quadratic Fourier), and so on. The Green-Tao theorem uses the full hierarchy: the proof circuit passes through U^k domains for all k, with each higher-order domain providing decorrelation for the corresponding order of obstruction.

Whether this hierarchy has a categorical interpretation — whether **Spec** admits a filtration by "order of correlation" with each stratum having its own transfer structure — is open and would connect the framework to Gowers' higher-order Fourier analysis program.

Problem 5. Structural properties of Spec.

Section 6.7 established that **Spec** is a dagger category enriched over $([0,1], \geq, \cdot)$ with a monoidal product. Several deeper structural questions remain:

(a) *Is* **Spec** *compact?* In the sense: does every infinite sequence of spectral domains (D_n) have a convergent subsequence (in some natural topology on objects)? If yes, the set of achievable decorrelation gains for any problem would be closed, and the supremum $\delta(P)$ would be attained. If not, $\delta(P)$ might be a limit that no finite circuit achieves.

(b) *Is the tensor product symmetric monoidal?* We defined $\mathcal{D}_1 \otimes \mathcal{D}_2$ in §6.7.1. Is $\mathcal{D}_1 \otimes \mathcal{D}_2 \cong \mathcal{D}_2 \otimes \mathcal{D}_1$ canonically (symmetric), and does $(\mathcal{D}_1 \otimes \mathcal{D}_2) \otimes \mathcal{D}_3 \cong \mathcal{D}_1 \otimes (\mathcal{D}_2 \otimes \mathcal{D}_3)$ (associativity)? For Hilbert space tensor products, both hold at the level of spaces. The question is whether the polarity functions are also compatible, i.e., whether the classification of spectral components as convolutive/correlative respects the tensor structure.

(c) *Does* **Spec** *admit a model structure?* In the sense of homotopy theory: is there a class of "weak equivalences" (transfers that preserve decorrelation gain exactly), "fibrations" (transfers that do not decrease gain), and "cofibrations"? A model structure would provide formal tools for proving that two proof circuits are "homotopy equivalent" — that they achieve the same decorrelation by different paths.

Problem 6. Discovery of new transfers.

The explicit graph (§6.8) has 7 vertices and approximately 12 known edges. The number of possible directed edges is 42. The ~ 30 missing edges fall into three categories:

(i) *Structurally impossible*: the domains are too dissimilar for a bounded linear map to exist with the spectral preservation properties required (§6.6, spectral transfer definition).

(ii) *Conjecturally exist*: the Langlands transfer (NT   algebraic geometry), higher-order Fourier transfers (combinatorics   higher-order harmonic analysis), and the geometric Langlands correspondence (number theory   representation theory).

(iii) *Unknown*: we do not have evidence for or against.

The question: **is there a systematic method to detect missing edges?** One approach: if domains $\mathcal{D}_1$ and $\mathcal{D}_2$ have a common sub-object (a spectral domain that embeds into both), then a transfer between them can be constructed as the composite of the two embeddings. Concretely, if number theory and PDE share a common "spectral geometry" domain (via the Selberg trace formula, which connects hyperbolic geometry to both zeta zeros and Laplacian eigenfunctions), this provides a candidate edge.

Problem 7. Automated circuit design.

Given a computable approximation to the transfer graph (finitely many domains, finitely many known transfers with computed δ values), the optimal proof circuit for a problem P is a combinatorial optimization problem:

Maximize $\delta(\gamma) = 1 - \prod_{i=1}^{k}(1 - \delta(\Phi_i))$ *over all paths* $\gamma = (\Phi_1, \ldots, \Phi_k)$ *in the transfer graph starting and ending at* $\mathcal{D}_P$*.*

This is a variant of the shortest path / maximum reliability problem on a weighted directed graph. Standard algorithms (Dijkstra on log-transformed weights, Bellman-Ford for general graphs) apply if the graph is finite.

The catch: the graph is infinite (there are infinitely many spectral domains) and the δ values are problem-dependent. A practical approximation:

1. Fix the seven known domains as vertices.
2. For each known transfer, compute δ for the target problem (this requires mathematical analysis, not just graph theory).
3. Run maximum-reliability path search on the resulting weighted graph.
4. If the best circuit has $\delta < \delta^*$, report the gap and identify which new edge would maximally reduce it.

Step 4 is the strategic output: it tells mathematicians *which new transfer to look for*. For the twin prime problem, the answer (from §6.8) is: a transfer from number theory to algebraic geometry (or to any new domain) with $\delta > 0.7$ for the twin prime correlative structure. This is a precise target for the Langlands program and related efforts.

7.5 Navier-Stokes Regularity as a δ Problem

The 3D incompressible Navier-Stokes regularity problem — the last Clay Millennium Prize — maps onto the **Spec** framework with surprising precision. The Gevrey energy balance [Nagy, 2026g] is

$$\frac{d}{dt}G_\sigma = \underbrace{-2\nu H_\sigma}_{\text{grade-2: viscous}} - \underbrace{2\,b_\sigma(u,u,u)}_{\text{grade-3: advective}}$$

The first term is a decorrelation mechanism: viscous diffusion damps high-frequency Fourier modes at rate $e^{-\nu|k|^2 t}$, providing $\delta_{\text{visc}} > 0$ unconditionally. The second term is a correlation mechanism: the advective nonlinearity $u \cdot \nabla u$ transfers energy between Fourier triads, potentially concentrating it at high wavenumbers. The regularity question is: **does $\delta_{\text{visc}} \geq \delta_{\text{adv}}$?**

2D: regularity from grade structure. In two dimensions, $b_\sigma \equiv 0$ for all Gevrey radii σ (enstrophy conservation forces the grade-3 term to vanish identically). Thus $\delta_{\text{adv}} = 0$, the viscous term dominates unconditionally, and regularity follows. This is Ladyzhenskaya's theorem (1969) through the spectral lens: in 2D, Navier-Stokes is a *purely convolutive* system.

3D: the advective obstruction. In three dimensions, $b_\sigma \neq 0$ for $\sigma > 0$ — vortex stretching injects energy into the Gevrey tail, competing with viscous damping. The triadic interaction $b_\sigma(u,u,u)$ is bounded by $C \cdot H_\sigma^{3/2}/G_\sigma^{1/2}$, which grows with enstrophy. The problem is whether this growth can outpace the viscous decay. In **Spec** terms: the circuit $\mathcal{D}_{\text{energy}} \to \mathcal{D}_{\text{Fourier}} \to \mathcal{D}_{\text{Gevrey}} \to \mathcal{D}_{\text{energy}}$ has a computable δ that depends on the enstrophy.

The parallel with the Riemann Hypothesis. The structural parallel between NS and RH through **Spec** is striking:

Feature	Riemann Hypothesis	Navier-Stokes regularity
Partial result	Density-1 zeros on $\text{Re}(s) = 1/2$	Almost-everywhere regularity (Leray)
Its δ	≈ 0.5 (most zeros accounted for)	≈ 0.5 (enstrophy $\in L^1$)
Full result	All zeros on line	Regularity for all time

Feature	Riemann Hypothesis	Navier-Stokes regularity
Its δ^*	1.0	1.0
Gap mechanism	Rare off-line zeros evade density estimates	Rare vortex stretching events evade enstrophy bounds
What's needed	L^p control with $p > 2$ on zero-counting error	L^p control with $p > 2$ on vorticity

In both cases, the existing best results achieve "density-1" or "almost everywhere" decorrelation ($\delta \approx 0.5$), and the gap to full resolution ($\delta = 1.0$) requires controlling rare extreme events — off-line zeros for RH, vortex stretching singularities for NS. The gap in both cases is ≈ 0.5, and no known technique closes it.

The L^p barrier at $p = 2$. The parallel is not metaphorical but structural. In both problems, the obstruction is a failure of L^p integrability at the critical exponent $p = 2$:

For RH: The zero-density estimate gives $N(\sigma, T) \leq CT^{A(1-\sigma)}(\log T)^B$ for the number of zeros with $\mathrm{Re}(\rho) > \sigma$. For $\sigma = 1/2 + \epsilon$, this controls *most* zeros — but "most" is an L^2 statement. To prove RH, one needs the pointwise statement "no zeros off the line," which requires L^∞ control. The gap between L^2 (density estimate) and L^∞ (RH) is the decorrelation gap. Every known technique for counting zeros — the Selberg moment method, the Korobov-Vinogradov estimates, density hypotheses — uses L^2 averaging. The correlative obstruction is that rare zeros can cluster off the line in configurations invisible to second-moment methods.

For NS: The enstrophy estimate gives $\int_0^T \|\nabla u\|^2 \, dt \leq C$ — an L^2-in-time bound on the gradient. For regularity, one needs the pointwise statement "$\|\nabla u(t)\| < \infty$ for all t," which requires L^∞-in-time control. The gap between L^2 (Leray's energy inequality) and L^∞ (regularity) is the decorrelation gap. Every known technique for bounding vorticity — the Beale-Kato-Majda criterion, the Prodi-Serrin conditions, the Escauriaza-Seregin-Šverák result — uses integral averaging. The correlative obstruction is that rare vortex stretching events can produce singularities invisible to L^2 enstrophy bounds.

The structural identity:

Quantity	RH	NS
L^2 controlled	Zero density in vertical strips	Enstrophy in time intervals
L^∞ needed	Pointwise location of every zero	Pointwise regularity at every time
Correlative residue	Off-line zero clusters	Vortex stretching events
Spectral decomposition	Zeta zeros ρ	Fourier modes $\hat{u}(k)$
Convolutive term	$\sum 1/\|\rho\|^2$ (convergent)	$\nu\|k\|^2\|\hat{u}\|^2$ (viscous damping)
Correlative term	$\sum 1/\|\rho\|$ (divergent)	$\sum_{k_1+k_2=k} \hat{u}(k_1) \cdot k_2 \, \hat{u}(k_2)$ (advective)

The convolutive-correlative decomposition is exact in both cases: viscous damping IS convolution (with the heat kernel), and vortex stretching IS correlation (the velocity field multiplied by its own gradient — the same field reused). The duality predicts exactly which term controls regularity and which term threatens blowup.

An open question. Both RH and NS have $\Delta \approx 0.5$. Is this a coincidence, or is there a deeper reason? One candidate: both problems sit at the critical exponent of a Sobolev embedding. For RH, the relevant embedding is $H^{1/2}(\text{critical line}) \hookrightarrow L^\infty$ (which fails in 1D). For NS, the relevant embedding is $H^1(\mathbb{R}^3) \hookrightarrow L^\infty$ (which fails in 3D). Both failures are borderline — the embedding holds in lower dimensions (this is why GRH for function fields is proved, and why 2D NS is regular). The decorrelation gap $\Delta \approx 0.5$ may reflect the universal margin by which critical Sobolev embeddings fail.

The companion papers [Nagy, 2026g; Nagy, 2026h] develop the NS analysis in full: 1,156 Platonic-verified theorems for the Gevrey approach, 400+ for the direction coherence approach. The grade decomposition in both approaches separates the convolutive (viscous) and correlative (advective) contributions in exactly the manner predicted by the duality.

Acknowledgments

The number-theoretic instance of this duality was developed in the course of the Goldbach Latent research program [Nagy, 2026d], where it emerged as Theorem T150. The formalization of the original theorem (190 verified theorems) was carried out using the Platonic proof language with Lean 4 verification.

AI Disclosure

This paper was drafted with AI assistance (Claude, Anthropic) under human direction. All mathematical content — the identification of the duality as a universal principle, the selection of instances, the formulation of the general theorem, and the identification of open questions — reflects the author's research program. The formal verification of the number-theoretic instance uses the Platonic proof language backed by a Lean 4 type checker.

References

1. Nagy, T. (2026a). "The Latent Is All You Need: A Finite Object Behind Every Smooth System." Working paper.
2. Nagy, T. (2026b). "D_∞ and the Goldbach Convergence Hierarchy." Working paper.
3. Nagy, T. (2026d). "The Goldbach Latent: Twenty-Five Paths and the Convolution Convergence Theorem." Working paper. 190 formally verified theorems.
4. Hardy, G. H. and Littlewood, J. E. (1923). "Some problems of 'Partitio Numerorum'; III: On the expression of a number as a sum of primes." *Acta Mathematica*, 44, 1–70.
5. Vinogradov, I. M. (1937). "Representation of an odd number as a sum of three primes." *Doklady Akademii Nauk SSSR*, 15, 169–172.
6. Wiener, N. (1930). "Generalized harmonic analysis." *Acta Mathematica*, 55, 117–258.

7. Khinchin, A. (1934). “Korrelationstheorie der stationären stochastischen Prozesse.” *Mathematische Annalen*, 109, 604–615.

8. Plünnecke, H. (1970). “Eine zahlentheoretische Anwendung der Graphentheorie.” *Journal für die reine und angewandte Mathematik*, 243, 171–183.

9. Ruzsa, I. Z. (1996). “Sums of finite sets.” In *Number Theory: New York Seminar*, Springer, 281–293.

10. Montgomery, H. L. (1973). “The pair correlation of zeros of the zeta function.” *Proc. Sympos. Pure Math.*, 24, 181–193.

11. Tao, T. and Vu, V. H. (2006). *Additive Combinatorics*. Cambridge University Press.

12. Shannon, C. E. (1948). “A Mathematical Theory of Communication.” *Bell System Technical Journal*, 27, 379–423, 623–656.

13. Fenton, L. F. (1960). “The Sum of Log-Normal Probability Distributions in Scatter Transmission Systems.” *IRE Transactions on Communications Systems*, 8(1), 57–67.

14. Fang, F. and Oosterlee, C. W. (2008). “A Novel Pricing Method for European Options Based on Fourier-Cosine Series Expansions.” *SIAM Journal on Scientific Computing*, 31(2), 826–848.

15. Nagy, T. (2013). *Valuation of Emission Permits and Power Plants under Uncertain Regulation.* PhD thesis, Corvinus University of Budapest.

16. Rosenblatt, M. (1952). “Remarks on a Multivariate Transformation.” *The Annals of Mathematical Statistics*, 23(3), 470–472.

17. Selberg, A. (1947). “An Elementary Proof of the Prime-Number Theorem.” *Annals of Mathematics*, 50(2), 305–313.

18. Nagy, T. (2026e). “Proof Circuits: Cross-Domain Transfer as Mathematical Method.” Working paper. Companion to the present paper — full categorical treatment of **Spec**, circuit analysis of historical proofs, and the missing edge problem.

19. Wilkinson, R. I. (1934). Unpublished memorandum, Bell Telephone Laboratories. Cited in Marlow, N. A. (1967), “A Normal Limit Theorem for Power Sums of Independent Random Variables,” *Bell System Technical Journal*, 46, 2081–2089.

20. Nagy, T. (2026c). “Platonic Kernel Formalization of Spectral Transfer Theory.” Working paper. 67 verified theorems in the Platonic proof language.

21. Nagy, T. (2026g). “Navier-Stokes Regularity via Gevrey Energy Balance and the Convolution-Correlation Duality.” Working paper. 1,156 Platonic-verified theorems.

22. Nagy, T. (2026h). “Direction Coherence for Navier-Stokes: A Spectral Approach.” Working paper. 400+ Platonic-verified theorems.

23. Nagy, T. (2026f). “The Decorrelation Index: A Single Number Measuring Distance to Proof.” Working paper. Develops $\delta(P)$ as a mathematical invariant with computability, monotonicity, and proof complexity results.

24. Ladyzhenskaya, O. A. (1969). *The Mathematical Theory of Viscous Incompressible Flow* (2nd ed.). Gordon and Breach.

25. Nagy, A., Savona, V. (2011). "Exploring phase transitions by finite-entanglement scaling of MPS in the 1D ANNNI model." arXiv:1102.4264.

26. Nagy, A., Romhányi, J., Shi, J.-D. (2011). "Simulating quantum systems on the Bethe lattice by translationally invariant infinite-Tree Tensor Network." arXiv:1106.3054.

27. Nagy, A., Akbari, A., Cavaliere, F. (2010). "Three qubits in a symmetric environment: dissipatively generated asymptotic entanglement." arXiv:1009.1476.

28. Nagy, A., Cavaliere, F., Ferraro, E. (2011). "Asymptotic Entanglement and Lindblad Dynamics: a Perturbative Approach." arXiv:1103.3790.

Appendix A: The Beta Function Mechanism in Detail

The spectral damping in the number-theoretic instance arises from the Beta function integral. For completeness, we derive the key identity.

Lemma. *Let* $Re(\rho_1), Re(\rho_2) > 0$*. Then*

$$\int_0^1 x^{\rho_1}(1-x)^{\rho_2}\,dx = B(\rho_1+1, \rho_2+1) = \frac{\Gamma(\rho_1+1)\Gamma(\rho_2+1)}{\Gamma(\rho_1+\rho_2+2)}.$$

For $|\rho_2| \to \infty$ *with* $Re(\rho_2) > 0$*:*

$$|B(\rho_1+1, \rho_2+1)| \sim \frac{|\Gamma(\rho_1+1)|}{|\rho_2|^{\mathrm{Re}(\rho_1)+1}}.$$

In particular, under RH ($Re(\rho) = 1/2$*),* $|B(\rho_1+1, \rho_2+1)| \sim C/|\rho_2|^{3/2}$*, and the total contribution to the Goldbach explicit formula is bounded by* $\sum_{\rho_2} C/|\rho_2|^{3/2}$*, which converges.*

The integration over x — the independent variable — is what produces the Gamma function in the denominator, and hence the damping. If x were fixed (as in the twin prime problem), no integral occurs, no Gamma function appears, and the spectral coefficient remains $1/|\rho|$.

This is the Beta function mechanism: the *existence of an independent integration variable* converts a Gamma function from the numerator (spectral amplitude) to the denominator (spectral damping). The mechanism is the Fourier-analytic manifestation of the general duality.