

Renormalization as Projected Ontology: The Theory That Was Never Divergent

Engin Atik¹

¹Kleis Research, <https://kleis.io>

Abstract

Volumes I–VI of this series established that physical observables are images of ontological flows under admissible Green’s function kernels, with gauge freedom identified as the kernel’s null space. This paper extends the framework to renormalization in quantum field theory. We argue that the regularization procedures used since 1948 — zeta regularization, heat kernel methods, Pauli–Villars, dimensional continuation — are projection kernels in the precise sense of the Projected Ontology (POT) framework. The finite part extracted by Hadamard’s constant-term operator is the image of the projection; the divergent terms are the null space; the choice of regularization scheme is gauge freedom. We formalize the **gauge group of admissible regulators** and prove that gauge-equivalent regulators yield identical observables. The central result is that the path integral and renormalization compose into a single integral transform $K_{\text{QFT}} = \text{FP} \circ K_{\text{ren}} \circ K_{\text{path}}$ mapping sources directly to observables. This composition is legitimized by the Integral Transform Composition Method (ITCM) of Sitnik and collaborators, which provides explicit kernel representations for composed integral transforms. Applied to QED, the ITCM identifies dimensional regularization as a transmutation operator with weight $w(t) = (-t^2)^s$, yielding a closed-form composite kernel via hypergeometric functions. Renormalization is therefore not an add-on procedure that fixes a broken theory; it is a constitutive factor of the integral transform that defines the theory. The heat kernel and propagator used for regularization are the same operators that govern physical phenomena; mathematical devices like the zeta kernel produce correct results because they are gauge-equivalent to a physical kernel. A literature survey confirms that every major application of zeta regularization in the last fifty years accessed the heat kernel through a mathematical intermediary. The claim is not that QFT is wrong, but that the essential structure underlying the formalization sought by the Clay Millennium Problem may already have been present in the original 1948 formulation, though not expressed as a single composite operator. We conclude with a discussion of what the kernel decomposition tempts us to infer about the ontological Hilbert space \mathcal{H}_{ont} , and why the framework’s own epistemic constraint — established in Volume II — prevents us from acting on that temptation. All structural results are machine-verified by the Z3 SMT solver via 40 axiomatic examples in the Kleis formal verification language.

Keywords: renormalization, projection kernel, regularization, zeta function, heat kernel, gauge freedom, finite part operator, quantum field theory, projected ontology, integral transform composition method, transmutation operator, composite kernel, formal verification, Z3

1 Introduction

Quantum field theory is the most precisely tested framework in the history of science. The anomalous magnetic moment of the electron is calculated to twelve decimal places and agrees with experiment. Yet QFT has no rigorous mathematical foundation. The Clay Mathematics Institute lists the existence and mass gap of Yang–Mills theory as one of seven Millennium Prize Problems. The standard assumption is that new mathematics is needed — constructive field theory, new axiom systems, new functional analysis.

This paper argues that the foundation was laid in 1948 and has been used successfully ever since. It was mislabeled as a computational technique.

When Tomonaga, Schwinger, and Feynman introduced renormalization to handle the divergent integrals of quantum electrodynamics, they described it as a procedure: subtract the infinities, keep the finite part, check against experiment. The procedure works spectacularly well. But a procedure is not a mathematical object. You cannot prove theorems about it, compose it with other procedures, or classify it. The Projected Ontology (POT) framework, developed in Volumes I–VI of this series, provides the missing language: renormalization is a **projection kernel**. Its image is the space of physical observables. Its null space is the space of divergent terms. The choice of regularization scheme — cutoff, exponential damping, heat kernel, zeta function, dimensional continuation — is gauge freedom within the projection.

This is not a new proposal for the foundations of QFT. It is the recognition that QFT’s own computational practice **already embodies** the projection-kernel structure. Every time a physicist subtracts a pole, takes a finite part, or checks scheme independence, they are performing kernel-based projection without naming it as such.

The paper is organized as follows. Section 2 recapitulates the POT framework established in Volumes I–VI. Section 3 introduces divergent objects and the regularization kernel. Section 4 treats the paradigmatic case of zeta regularization and $\zeta(-1) = -1/12$. Section 5 treats the product case and the connection to the Gamma function and $\sqrt{2\pi}$. Section 6 formalizes the gauge group of admissible regulators. Section 7 identifies the finite part as the gauge-invariant observable. Section 8 shows that regularization kernels are physical kernels — the same operators that govern heat diffusion, particle exchange, and spectral theory. Section 9 draws the implications for QFT, including the composite kernel $K_{\text{QFT}} = \text{FP} \circ K_{\text{ren}} \circ K_{\text{path}}$, its connection to the Integral Transform Composition Method (ITCM), and the explicit application to QED. Section 10 concludes.

2 Recapitulation of the POT Framework

We summarize the results of Volumes I–VI that this paper extends.

2.1 Admissible Kernels (Volumes I–III)

A projection kernel K is admissible if it satisfies three conditions: (1) linearity — $K(A + B) = K(A) + K(B)$; (2) scalar compatibility — $K(cA) = cK(A)$; (3) zero preservation — $K(0) = 0$. These conditions formalize Occam’s razor for projections: an admissible kernel projects without adding structure.

Every physical law in the POT framework is a property of an admissible kernel. Observable quantities are the kernel's image. Gauge freedom is the kernel's null space. Volume I derived galactic rotation curves from the gravitational kernel. Volume II showed that Bell-inequality violation arises from projecting non-separable ontic states through an admissible kernel. Volume III identified the electromagnetic field strength $F = dA$ as the image of the exterior derivative kernel, with gauge potentials differing by exact forms as the null space.

2.2 The Admissibility Defect and Confinement (Volume IV)

For non-abelian gauge theories, the kernel is not admissible. The defect $\Delta(A, B) = K(A + B) - K(A) - K(B) = [A, B]$ is the Lie bracket of the gauge algebra. Volume IV proved that a non-admissible kernel produces fiber non-invariance: charge becomes a predicate that cannot be determined from gauge-invariant observables. This is structural confinement.

2.3 Admissibility Restoration (Volume V)

Volume V resolved the puzzle of the weak interaction, where $SU(2)$ is non-abelian but charges **are** observed. A non-admissible kernel can become effectively admissible through coupling to a restoring field that compensates the Lie bracket defect. The structural constraints on any such field — gauge-charged, non-zero preferred value, characteristic scale — match the defining properties of the Higgs. Three forces, one classification principle: admissible (electromagnetism), restored (weak), unrestored (strong).

2.4 The Four-Sentence Axiom (Epilogue, Volume VI)

The Epilogue distilled the POT framework into four sentences: (1) Physical laws are properties of the kernel. (2) Observable quantities are the image. (3) Gauge freedom is the null space. (4) There is nothing else.

This paper applies the same four sentences to renormalization.

3 Divergent Objects and the Regularization Kernel

A **divergent object** X is a mathematical entity — a formal series, a spectral sum, an asymptotic integral — that does not converge in the ordinary sense. The expression $1 + 2 + 3 + 4 + \dots$ is a divergent object. The infinite product $1 \times 2 \times 3 \times 4 \times \dots$ is a divergent object. The loop integrals of QFT are divergent objects.

A **regularization kernel** K_R maps a divergent object X to a parameterized family of values $K_R[X](\varepsilon)$ that is well-defined for ε in some domain. The parameter ε is the regularization parameter: it could be a complex variable s (zeta), a proper time t (heat kernel), a cutoff Λ (momentum), or a dimension $d = 4 - \varepsilon$ (dimensional continuation).

Definition. A regularization kernel K_R is **admissible** if it satisfies two conditions:

(1) **Consistency:** For any convergent object X , the regularized value agrees with the ordinary value.

(2) **Zero preservation:** The regularization of the zero object is zero.

These conditions are the regularization analogues of the admissibility conditions in Volumes I–III. An admissible regularization kernel does not create observables where there are none, and it agrees with ordinary summation where ordinary summation is defined.

The formal structures are defined in `RegularizationKernel` and verified by `Z3` (Examples 1–3 in the theory file).

4 The Sum Case: Spectral Zeta Regularization

The Riemann zeta function $\zeta(s) = \sum_{n=1}^{\infty} n^{-s}$ converges for $\text{Re}(s) > 1$ and admits analytic continuation to $\mathbb{C} \setminus \{1\}$. The regularization kernel is $K(s, n) = n^{-s}$: it suppresses the growth of n by a power law controlled by s .

At $s = -1$, the formal series $\sum n^{-(-1)} = \sum n = 1 + 2 + 3 + \dots$ diverges. But the analytically continued value is $\zeta(-1) = -1/12$.

In POT language: the divergent object $X = (1, 2, 3, \dots)$ is the raw ontic structure. The zeta kernel K_{ζ} maps it to a meromorphic function $\zeta(s)$. The **image** of the projection — the finite part at $s = -1$ — is $-1/12$. The divergent partial sums $S_N = N(N+1)/2 \rightarrow \infty$ are the **null space**: they are annihilated by the projection.

The value $-1/12$ is not an arithmetic identity. It is a projected observable: the unique finite value extracted from a divergent structure by an admissible kernel. The same value appears in the Casimir effect (vacuum energy between parallel plates) and in string theory (the critical dimension $d = 26$ follows from $\zeta(-1) = -1/12$). These are not coincidences. They are instances of the same projection.

The `SpectralZeta` structure in the theory file asserts that the zeta kernel is admissible and that $\zeta(-1) = -1/12$ is the projected observable of the integer spectrum. Both claims are verified by `Z3` (Examples 7–9).

5 The Product Case: Zeta-Regularized Products and $\sqrt{2\pi}$

The infinite product $1 \times 2 \times 3 \times \dots$ diverges. But the zeta-regularized product, defined via the derivative of the zeta function at zero, yields a finite value.

The zeta-regularized product of a sequence λ_n with spectral zeta function $\zeta_A(s) = \sum \lambda_n^{-s}$ is:

$$\prod_{\text{reg}} \lambda_n = \exp(-\zeta'_A(0))$$

For the integer spectrum $\lambda_n = n$, this gives:

$$\prod_{\text{reg}} n = \exp(-\zeta'(0)) = \sqrt{2\pi}$$

since $\zeta'(0) = -1/2 \log(2\pi)$.

The constant $\sqrt{2\pi}$ is one of the most recurring values in mathematics. It appears as: the normalization of the Gaussian integral $\int_{-\infty}^{\infty} e^{-x^2/2} dx = \sqrt{2\pi}$; the leading term of Stirling's approximation $n! \sim \sqrt{2\pi n}(n/e)^n$; the area under the standard normal distribution; and the regularized product of all positive integers.

In POT language: these are not four different facts. They are four manifestations of the same projected observable. The zeta kernel, the Laplace/saddle-point kernel, and the Fourier kernel all extract the same finite value from different divergent structures because they are gauge-equivalent projections (Section 6).

The connection to the Gamma function makes this concrete. The Gamma function $\Gamma(n+1) = n! = \int_0^{\infty} x^n e^{-x} dx$ can be evaluated by Laplace's method for large n : the integrand is dominated by a peak near $x = n$, and expanding to second order (the **quadratic germ** or **parabolic approximation**) produces $\sqrt{2\pi n}$ as the area of the Gaussian envelope. The saddle-point kernel extracts the same finite value as the zeta derivative at zero — because both are admissible regularization kernels projecting onto the same image.

This is formalized in `SpectralZeta` (axiom `zeta_product_via_derivative`, Example 10 in the theory file).

6 The Gauge Group of Admissible Regulators

Different regularization schemes modify a divergent object in different ways but extract the same finite part. In POT language, they are gauge transformations.

Definition. Two admissible regulators R and R' are **gauge-equivalent**, written $R \sim R'$, if for all divergent objects X :

$$\text{FP}[K_R[X]] = \text{FP}[K_{R'}[X]]$$

where `FP` denotes the finite part operator (Section 7).

Theorem 3 (Equivalence relation). Gauge equivalence is reflexive, symmetric, and transitive. The proof is immediate from the definition and is verified by Z3 (Examples 4–6).

Theorem 4 (Gauge invariance). If $R \sim R'$, then for every divergent object X , the projected observable is the same:

$$\text{FP}[K_R[X]] = \text{FP}[K_{R'}[X]]$$

This is the regulator analogue of gauge invariance in Volumes III–IV: different representatives of the same equivalence class produce the same observable. The divergent terms — the pole structure, the power-law blowup, the logarithmic divergence — are the null space. They differ between schemes. The finite part is the image. It is scheme-independent.

Example. The heat kernel $K(t, \lambda) = e^{-t\lambda}$ and the zeta kernel $K(s, n) = n^{-s}$ are gauge-equivalent. They are connected by the Mellin transform:

$$\zeta_A(s) = \frac{1}{\Gamma(s)} \int_0^\infty t^{s-1} \operatorname{Tr}(e^{-tA}) dt$$

This integral relation means that the spectral zeta function is determined by the heat trace. Since both encode the same spectral information, their finite parts agree. Schwinger’s proper-time method (1948) and Hawking’s zeta regularization (1977) give the same physics — because they are gauge-equivalent kernels. (Section 8 shows that the heat kernel is the physically fundamental member of this equivalence class, with the zeta kernel as a mathematical intermediary.)

The formal proof of gauge equivalence between heat and zeta regulators, together with the Mellin bridge, is verified in `HeatKernelRegularization` (Examples 11–13).

Remark (The gauge group). The set of admissible regulators modulo gauge equivalence forms a group G_{adm} acting on divergent objects. The group operation is composition of regulators; the identity is the trivial regulator (which agrees with ordinary summation on convergent objects). This group is the analogue of the gauge group in Volumes III–V, with regulator equivalence classes playing the role of gauge orbits (projection fibers). The formal axioms for this group — including identity, reflexivity, symmetry, and transitivity — are in `RegulatorGaugeGroup`.

7 The Finite Part as Gauge-Invariant Observable

The finite part operator, due to Hadamard, extracts the constant term from the asymptotic expansion of a regularized value.

Definition. Given a regularized value $F(\varepsilon)$ with asymptotic expansion

$$F(\varepsilon) = S(\varepsilon) + C + o(1) \quad \text{as } \varepsilon \rightarrow 0$$

where $S(\varepsilon)$ belongs to a prescribed space of singular germs (poles, logarithms, power-law divergences), the **finite part** is:

$$\text{FP}[F] := C$$

For meromorphic functions, this is the Laurent extraction: if $F(s) = \sum a_k (s - s_0)^k$, then $\text{FP} := a_0$.

Theorem 5 (Linearity). The finite part operator is linear:

$$\text{FP}[F + G] = \text{FP}[F] + \text{FP}[G] \quad \text{FP}[cF] = c \cdot \text{FP}[F]$$

Theorem 6 (Zero preservation). The finite part of zero is zero: $\text{FP}[0] = 0$.

These properties make the finite part operator a **projection** in the algebraic sense. Combined with the gauge invariance of Section 6, the finite part is the unique quantity that is both scheme-independent and extracted by a linear, zero-preserving operator. It is the POT observable.

The finite part is the **image** of the projection. The singular part $S(\varepsilon)$ is the **null space**: it is annihilated by FP. The decomposition $F = S + C + o(1)$ is the null-space/image decomposition of Volume III applied to regularization.

Regularization kernel	Status	What it does
Heat kernel $e^{-t\lambda}$	Physical	Governs heat diffusion, quantum propagation
Propagator $1/(k^2 + m^2)$	Physical	Governs massive particle exchange (Yukawa)
Biot–Savart $\mathbf{x}/ \mathbf{x} ^3$	Physical	Governs electromagnetism, fluid mechanics
Zeta kernel n^{-s}	Mathematical	Analytic continuation of spectral sums
Dimensional continuation $d = 4 - \varepsilon$	Mathematical	Interpolation of spacetime dimension

Table 1: Physical kernels govern real processes; mathematical kernels produce correct results because they are gauge-equivalent to a physical kernel.

All properties are formalized in `FinitePartOperator` and verified by `Z3` (Examples 1–3).

8 Regularization Kernels Are Physical Kernels

Some regularization kernels used in QFT are not abstract mathematical devices borrowed to fix divergences — they are the same kernels that govern physical phenomena. Others are mathematical tools that work because they are gauge-equivalent to a physical kernel. The distinction matters.

8.1 The heat kernel is physical

When Schwinger used the heat kernel to regularize QED, he was not performing an abstract mathematical trick. He was applying the kernel for heat diffusion — the same operator that Fourier introduced to describe the conduction of heat in solids. The heat kernel e^{-tA} is the fundamental solution of the heat equation $\partial_t u = -Au$. Schwinger’s proper-time method applies this physical kernel to the operator $A = -D^2 + m^2$ of the quantum field. The regularization is not borrowed from physics — it **is** physics.

Similarly, the Feynman propagator $1/(k^2 + m^2)$ is the Green’s function for the Klein–Gordon equation. It governs massive particle exchange and yields the Yukawa potential at long distances. When it appears in loop integrals, it is simultaneously the physical content of the theory and the mechanism that regularizes (partially) the divergence.

The Biot–Savart kernel, treated in the Epilogue, governs both electromagnetism (source = current density) and fluid mechanics (source = vorticity). It is the POT exemplar of a single kernel giving rise to different physics through different source spaces.

These are **physical** kernels: they describe measurable processes. Their use as regularizers is not a coincidence — it is physics recognizing its own projection structure.

8.2 The zeta kernel is a mathematical device

The zeta kernel n^{-s} is not physical in the same sense. The Riemann zeta function is an analytic continuation of a spectral sum — a mathematical operation, not a physical process. Nobody measures $\zeta(-1)$; the number $-1/12$ is extracted by continuing a function beyond its domain of convergence.

When Hawking (1977) used zeta regularization to compute the partition function of quantum fields in curved spacetime, his physical problem **was** physical: he needed the one-loop effective action for quantum fields propagating on a black hole background. The question was: what is the energy-momentum tensor of quantum fields near a Schwarzschild black hole? The answer determines Hawking radiation — the thermal emission from black holes.

The physical content of Hawking’s calculation is in the heat kernel. The heat trace $\text{Tr}(e^{-tA})$, where $A = -\nabla^2 + m^2 + \xi R$ is the curved-space wave operator, describes the diffusion of a scalar field on the black hole geometry. The short-time expansion of this heat trace encodes the local geometry (the Seeley–DeWitt coefficients a_0, a_1, a_2, \dots), and its integral gives the one-loop effective action. This is a physical computation: the heat kernel propagates the field on the curved background, and its trace integrates over all field modes.

Hawking then applied the Mellin transform:

$$\zeta_A(s) = \frac{1}{\Gamma(s)} \int_0^\infty t^{s-1} \text{Tr}(e^{-tA}) dt$$

to convert the heat trace into a zeta function, because the zeta function is easier to analytically continue. The zeta-regularized determinant $\det(A) = \exp(-\zeta'_A(0))$ gives the effective action in closed form. But the physics was always in the heat trace. The zeta function is a mathematical intermediary — an elegant bookkeeping device for the spectral data that the heat kernel encodes.

Hawking’s results are correct not because zeta regularization is physical, but because zeta regularization is **gauge-equivalent** to the heat kernel (Section 6). The Mellin bridge guarantees they extract the same finite parts. The zeta function inherits its physical correctness from the heat kernel.

8.3 The preferred representative

This observation sharpens the gauge group of regulators (Section 6). Not all gauge-equivalent regulators are equally fundamental. The gauge group has a **preferred representative**: the physical kernel. The heat kernel is the regulator that directly describes the underlying physics — diffusion, propagation, heat flow. The zeta kernel, dimensional regularization, and other mathematical schemes work because they are in the same gauge orbit as the heat kernel.

The analogy to standard gauge theory is precise. In electrodynamics, all gauges ($\nabla \cdot \mathbf{A} = 0$, $A_0 = 0$, etc.) are physically equivalent, but the Coulomb gauge reveals the instantaneous interaction most directly. Similarly, all admissible regulators yield the same observables, but the heat kernel reveals the physical mechanism most directly. The others are computational gauges’ — mathematically convenient, physically derivative.

Theorem 7 (Physical kernels are admissible). If a regulator R governs a physical phenomenon, then R is admissible.

This codifies the empirical observation that every regularization scheme known to produce correct physics satisfies the admissibility conditions. The heat kernel and the propagator satisfy consistency (they agree with ordinary summation on convergent cases) and zero preservation. Mathematical devices like the zeta kernel also satisfy these conditions — not because they govern a physical process, but because the Mellin bridge transfers the admissibility of the heat kernel to the zeta kernel. Admissibility is inherited via gauge equivalence.

The formal verification is in `RegulatorPhysicalKernel` (Examples 14–16).

The double-duty thesis. The physical kernel does double duty: it defines the physical law **and** projects out the observable. When QFT uses the heat kernel to regularize, it is recognizing that the heat kernel is part of the theory’s own projection structure. When QFT uses the zeta kernel instead, it gets the same answer — not because analytic continuation is physics, but because gauge equivalence guarantees it. The physics was always in the heat kernel. The regularization trick was never a trick — it was physics, sometimes accessed directly (Schwinger), sometimes accessed through a mathematical intermediary (Hawking).

8.4 Fifty years of evidence: the zeta kernel always accessed the heat kernel

The claim that the heat kernel is the preferred representative is not a philosophical preference — it is an empirical observation supported by every major application of zeta regularization in the literature.

The mathematical foundation was laid by Minakshisundaram and Pleijel (1949), who defined the spectral zeta function of the Laplacian as the Mellin transform of the heat trace, and by Seeley

Author(s)	Year	Problem	Underlying kernel
Schwinger	1951	QED vacuum polarization	Heat kernel (used directly)
McKean–Singer	1967	Index theorem	Heat kernel (used directly)
Seeley	1967	Complex powers of elliptic operators	Proved $\zeta_A(s) = \mathcal{M}[\text{Tr}(e^{-tA})]$
Ray–Singer	1971	Analytic torsion	Laplacian heat kernel via Mellin
Dowker–Critchley	1976	de Sitter effective action	Heat kernel of curved-space wave operator
Hawking	1977	Black hole radiation	Heat kernel of $-\nabla^2 + m^2 + \xi R$
Voros	1987	Quantum spectral theory	Partition function $\text{Tr}(e^{-\beta H})$
Elizalde	1994–	Casimir, gravity, strings (10 applications)	Laplacian heat kernel in each case
String theory	various	Critical dimension, anomalies	Worldsheet Laplacian heat kernel
Kaluza–Klein	various	Extra-dimensional corrections	Heat kernel on compact dimensions

Table 2: Every major application of zeta regularization in physics. In each case, the physical content resides in the heat kernel; the zeta function is a mathematical intermediary derived via the Mellin transform.

(1967), who proved that $\text{Tr}(A^{-s})$ is meromorphic with poles and residues determined by the heat kernel coefficients. The direction of derivation is always heat kernel \rightarrow zeta function, never the reverse.

Schwinger (1951) used the heat kernel directly. His proper-time method never invokes analytic continuation; it computes the effective action from the heat trace $\text{Tr}(e^{-tA})$ by integrating over proper time t . No zeta function appears.

Dowker and Critchley (1976) computed the effective Lagrangian for a scalar field in de Sitter space. Their wave operator $A = -(\square)^2 + \xi R + m^2$ is a Laplacian on curved spacetime. The heat kernel e^{-tA} describes physical field propagation on this background. They chose zeta regularization for notational convenience — the heat kernel method gives identical results.

Ray and Singer (1971) defined analytic torsion as $\exp(-\zeta'_\Delta(0))$, a topological invariant of the Laplacian on differential forms. The spectral data is encoded by the heat trace. Later formulations define torsion directly from the heat kernel; the zeta approach was an elegant shortcut.

In bosonic string theory, the critical dimension $D = 26$ requires $\zeta(-1) = -1/12$. The underlying operator is the worldsheet Laplacian, whose heat kernel $e^{-t\Delta_{\text{ws}}}$ describes diffusion on the string worldsheet. The conformal anomaly that fixes D can be computed entirely via heat kernel coefficients.

For the Casimir effect, the vacuum energy between conducting plates involves the eigenvalues of the Laplacian with boundary conditions. Elizalde (2012) explicitly demonstrated that the heat kernel method and zeta regularization yield identical Casimir energies for every geometry studied.

In Kaluza–Klein models, quantum corrections from the infinite tower of KK modes require regularization of spectral sums on compact extra dimensions. The heat kernel on S^N provides the spectral data; zeta regularization processes it.

The pattern is without exception: **nobody ever used the zeta kernel because the physics demanded it. In every case, the physics demanded the heat kernel, and the zeta function was a mathematical convenience for processing the spectral data that the heat kernel encodes.** The gauge equivalence of Section 6 is not merely a formal observation — it is the mechanism by which fifty years of zeta-regularized calculations inherited their physical correctness from the heat kernel.

9 Implications for Quantum Field Theory

Feynman’s path integral defines QFT as:

$$Z[J] = \int \mathcal{D}\varphi e^{iS[\varphi] + i \int J\varphi}$$

This is a sum over all field configurations weighted by the action. The definition is complete, elegant, and produces divergent quantities. Renormalization appears nowhere in it.

The actual practice of QFT proceeds in four steps: (1) **Definition** — the path integral produces divergent correlation functions. (2) **Computation** — Feynman diagrams organize the perturba-

tive expansion; divergences appear as loop integrals. (3) **Renormalization** — divergences are subtracted and finite parts retained. (4) **Comparison with experiment** — the finite results agree to extraordinary precision.

The gap is between steps 1 and 3. The definition does not include the mechanism that makes the theory produce finite observables. Renormalization is bolted on as a computational necessity, not derived from the formulation.

9.1 The composite kernel: $K_{\text{QFT}} = \text{FP} \circ K_{\text{ren}} \circ K_{\text{path}}$

In POT language, the path integral defines a kernel K_{path} that maps source currents J to divergent correlation functions $Z[J]$. The renormalization kernel K_{ren} maps divergent objects to regularized families. The finite part operator FP extracts the image. The complete theory is the **composition**:

$$K_{\text{QFT}} = \text{FP} \circ K_{\text{ren}} \circ K_{\text{path}} : \text{sources} \rightarrow \text{observables}$$

This is not three separate procedures applied sequentially. It is a **single integral transform** from sources to observables. The path integral alone is not the theory — it is one factor of the theory. Renormalization is not a fix applied after the fact — it is another factor. The composed object K_{QFT} is the theory.

This is the Volume V pattern applied to divergences. In Volume V, the gauge kernel K_{gauge} alone is non-admissible (the Lie bracket defect prevents linearity), but coupling to the Higgs field φ restores admissibility: $K_{\text{eff}} = K_{\text{gauge}} \circ \varphi$ is admissible. Here, K_{path} alone produces non-admissible output (divergent), but composing with K_{ren} restores admissibility: K_{QFT} maps sources to finite, gauge-invariant observables.

Feynman gave us K_{path} . Tomonaga, Schwinger, and Feynman (again) gave us K_{ren} . Nobody wrote down $K_{\text{QFT}} = \text{FP} \circ K_{\text{ren}} \circ K_{\text{path}}$ as the **definition** of the theory. That composition IS the theory.

9.2 Integral Transform Composition Method (ITCM)

The claim that integral transforms compose into a single well-defined integral transform is not a POT invention. It is a theorem of transmutation theory, formalized by Sitnik and collaborators as the **Integral Transform Composition Method (ITCM)**.

The ITCM constructs transmutation operators as compositions of generalized Fourier transforms:

$$S = F_B^{-1} \frac{1}{w(t)} F_A, \quad P = F_A^{-1} w(t) F_B$$

where F_A and F_B are integral transforms connected to operators A and B , and $w(t)$ is a weight function. Sitnik and Jebabli prove that these compositions yield explicit integral operators with closed-form representations involving hypergeometric functions, Bessel functions, and Gamma factors. All known classes of transmutation operators — Poisson, Sonine, Vekua–Erdélyi–Lowndes, Buschman–Erdélyi — arise as special cases.

ITCM component	QED realization
F_A (generalized Fourier)	Fourier transform to momentum space
F_A^{-1} (inverse transform)	Inverse Fourier to position space
$w(t)$ (weight function)	Renormalization weight: counterterms + running coupling
$T = F^{-1}wF$ (transmutation)	Complete renormalized propagator
Index shift $TB_\mu = B_\nu T$	Running of coupling constants with energy scale

Table 3: The ITCM-QED correspondence. Each component of the transmutation operator maps directly to a standard QED operation.

The ITCM result is decisive for the POT thesis: **the composition of integral transforms is itself an integral transform**. This is not an approximation, not a formal manipulation — it is a rigorous theorem with explicit integral representations (Sitnik and Shishkina, 2020; Sitnik and Jebabli, 2024).

9.3 Application to QED

Now apply the ITCM framework to quantum electrodynamics. The QED procedure is:

- (1) Transform field configurations to momentum space (Fourier transform F).
- (2) Apply the renormalization weight $w(k)$ — counterterms, self-energy $\Sigma(p)$, vacuum polarization $\Pi(q^2)$, vertex corrections.
- (3) Transform back to position space (inverse Fourier F^{-1}).

This is the ITCM formula $T = F^{-1} \circ w(k) \circ F$ **verbatim**. The QED-to-ITCM dictionary is:

For one-loop QED, the weight function takes the form:

$$w(k) \sim 1 + \frac{\alpha}{2\pi}[\text{loop corrections}] + \mathcal{O}(\alpha^2)$$

where $\alpha \approx 1/137$ is the fine-structure constant. The **running** of α with energy scale — the renormalization group flow — is the scale-dependence of $w(k, \mu)$ as a function of the renormalization point μ . In ITCM language, this is an **index shift**: the transmutation operator shifts the spectral index of the Bessel-type kernel from μ to ν , exactly as in Sitnik and Jebabli’s Theorems 1–4.

The crucial insight is that dimensional regularization corresponds to the ITCM with $w(t) = (-t^2)^s$ where $s = \varepsilon/2$ and $d = 4 - 2\varepsilon$. The resulting transmutation operator has the explicit integral representation given by Theorem 5 of Sitnik and Jebabli, involving Gauss hypergeometric functions ${}_2F_1$. This provides a **closed-form kernel** for the renormalized QED propagator — something the standard QFT literature treats only perturbatively.

The fine-structure constant α itself may be understood as a projected observable: the finite part extracted from the electromagnetic kernel’s spectral data at a given energy scale. Its value $\alpha^{-1} \approx 137$ at low energies, and its running to $\alpha^{-1} \approx 128$ at the Z -boson mass, is the ITCM index shift in action.

9.4 The explicit hypergeometric kernel for QED

The ITCM does not merely assert that the composition exists — it provides the closed-form integral representation. By Theorem 5 of Sitnik and Jebabli (2024), the composition of Hankel transforms H_ν and H_μ with the power weight $w(t) = (-t^2)^s$ yields:

$$\begin{aligned} [H_\nu(-t^2)^s H_\mu f](x) &= c_1 x^{-\mu-2s-3/2} \int_0^x {}_2F_1(a_1, b_1; c_1; y^2/x^2) y^{\mu+1/2} f(y) dy \\ &+ c_2 x^{\nu+1/2} \int_x^\infty {}_2F_1(a_2, b_2; c_2; x^2/y^2) y^{-\nu-2s-3/2} f(y) dy \end{aligned}$$

where the hypergeometric parameters are:

$$a_1 = \frac{\mu + \nu}{2} + s + 1, \quad b_1 = \frac{\mu - \nu}{2} + s + 1, \quad c_1 = \mu + 1$$

and $c_1(s, \mu, \nu)$, $c_2(s, \mu, \nu)$ are explicit constants involving Gamma functions:

$$c_1 = (-1)^s \cdot 2^{2s+1} \cdot \frac{\Gamma((\mu + \nu)/2 + s + 1)}{\Gamma((\nu - \mu)/2 - s) \cdot \Gamma(\mu + 1)}$$

For QED with dimensional regularization $d = 4 - 2\varepsilon$, the parameter identification is: $s = \varepsilon/2$ (the regularization parameter), ν is the free-field spectral index (determined by spin and mass), and μ is the interacting-field spectral index (shifted by the electromagnetic coupling). The transmutation operator maps free-field Green's functions to renormalized interacting-field Green's functions in a **single integral operation** — not as a perturbative series in α , but as a closed-form integral with hypergeometric kernel.

This is a non-perturbative result. Standard QED provides renormalized amplitudes as power series in α ; the ITCM provides the **exact integral operator** whose Taylor expansion reproduces those series. The kernel decomposes into two terms: an inner integral over $[0, x]$ and an outer integral over $[x, \infty)$, each weighted by a Gauss hypergeometric function ${}_2F_1$. The two terms correspond to two integration regions — forward and reflected channels — and their specific hypergeometric weights encode all the information that perturbative QED distributes across Feynman diagrams of every loop order.

9.4.1 Symmetries of the hypergeometric kernel

The kernel has three symmetries that dramatically simplify its structure.

Exchange symmetry ($\mu \leftrightarrow \nu$). The second-term parameters (a_2, b_2, c_2) are obtained from the first by swapping μ and ν : $a_2 = a_1$ (it is symmetric in μ, ν), $b_2 = (\nu - \mu)/2 + s + 1$, and $c_2 = \nu + 1$. The relation $b_1 + b_2 = 2(s + 1)$ holds. The two causal regions are therefore **spectral duals**: they exchange the free-field index ν and interacting-field index μ . The full kernel is invariant under simultaneous $\mu \leftrightarrow \nu$ and transposition of the integration regions.

Euler factorization. Applying Euler's hypergeometric transformation ${}_2F_1(a, b; c; z) = (1 - z)^{c-a-b} {}_2F_1(c - a, c - b; c; z)$ to the first term, the Euler exponent is:

$$c_1 - a_1 - b_1 = (\mu + 1) - ((\mu + \nu)/2 + s + 1) - ((\mu - \nu)/2 + s + 1) = -1 - 2s$$

The kernel factors into a **universal pole** and a **regular correction**:

$${}_2F_1(a_1, b_1; c_1; y^2/x^2) = (1 - y^2/x^2)^{-1-2s} \cdot {}_2F_1((\mu - \nu)/2 - s, (\mu + \nu)/2 - s; \mu + 1; y^2/x^2)$$

The singular factor $(1 - y^2/x^2)^{-1-2s}$ has a branch point at $y = x$ (the lightcone). The remaining ${}_2F_1$ is **regular** for all $y^2/x^2 \in [0, 1]$. This factorization separates the **universal singularity structure** — which generates the divergences when $s \rightarrow 0$ — from the **theory-specific content**, which determines the finite parts.

At $s = 0$ (the physical limit), the singular factor becomes $(1 - y^2/x^2)^{-1} = x^2/(x^2 - y^2)$, a simple Cauchy-type pole. The full kernel reduces to:

$$K(x, y)|_{s=0} = \frac{x^2}{x^2 - y^2} \cdot {}_2F_1((\mu - \nu)/2, (\mu + \nu)/2; \mu + 1; y^2/x^2)$$

Green's function structure of the Cauchy pole. The denominator factors as $x^2 - y^2 = (x - y)(x + y)$, revealing the propagator content:

$$\frac{x^2}{x^2 - y^2} = \frac{x^2}{(x - y)(x + y)} = \frac{x}{2} \cdot \left[\frac{1}{x - y} + \frac{1}{x + y} \right]$$

The pole $1/(x - y)$ is the **direct channel**: singular when $y = x$ (the source coincides with the observation point, forward propagation). The pole $1/(x + y)$ is the **reflected channel**: singular at $y = -x$ (backward propagation, the crossed channel). Their sum realizes a propagator-like decomposition into direct and reflected channels, structurally analogous to the momentum-space scalar propagator $1/(k^2 - m^2) = 1/((k - m)(k + m))$. The ITCM does not assume this propagator structure — it **derives** it from the composition of integral transforms.

The full kernel at $s = 0$ then reads:

$$K(x, y)|_{s=0} = \underbrace{\frac{x^2}{(x - y)(x + y)}}_{\text{free propagator}} \cdot \underbrace{{}_2F_1((\mu - \nu)/2, (\mu + \nu)/2; \mu + 1; y^2/x^2)}_{\text{interaction dressing}}$$

The interaction dressing is the regular ${}_2F_1$ correction: it modifies the free propagator's pole structure without introducing new singularities. The mass gap question, in this language, is whether the interaction dressing shifts the propagator's spectral support away from zero — whether R^{YM} converts the gapless Green's function $1/(x^2 - y^2)$ into a massive one.

Free-field limit ($\mu \rightarrow \nu$). When $\mu = \nu$ (coupling turned off), the parameters simplify: $b_1 \rightarrow s + 1$ and $c_1 \rightarrow \nu + 1$. At $s = 0$: ${}_2F_1(\nu + 1, 1; \nu + 1; z) = (1 - z)^{-1}$, since ${}_2F_1(a, b; a; z) = (1 - z)^{-b}$ for any a . The entire kernel collapses to the Cauchy kernel $1/(1 - z)$ — which is the **free-field propagator**. The hypergeometric correction vanishes, and only the universal pole survives.

Summary of the factorization. The QED kernel has the schematic form:

$$K_{\text{QED}}(x, y) = [\text{Cauchy pole}]^{1+2s} \times [\text{regular hypergeometric correction}] \times [\text{power-law weights}]$$

The Cauchy pole is universal (every QFT has it). The hypergeometric correction encodes the coupling. The power-law weights track the spectral indices. The factorization makes the physics transparent: the pole generates the loop integrals, the correction determines their finite parts, and the weight functions encode the theory's specific quantum numbers.

The theory file formalizes this as the `QEDHypergeometricKernel` structure with six hypergeometric parameters $(a_1, b_1, c_1, a_2, b_2, c_2)$ and axioms for kernel decomposition, exchange symmetry, and gauge invariance. The `ITCMtransmutation` structure establishes that the ITCM transmutation equals the composite QED kernel (Example 21), and that the index shift maps the free-field spectral index to the interacting-field index (Example 22), formalizing the intertwining property $TA = BT$ that underlies the running of coupling constants.

Kernel Decomposition Principle (proposed). Any ITCM-generated transmutation kernel admits a decomposition into a universal singular Green’s kernel and a regular dressing factor, with spectral properties determined solely by the latter after finite-part extraction.

This principle unifies the results of this subsection. The Euler factorization (derived from the composition theorem) separates the universal pole from the regular correction. The partial-fraction decomposition reveals the pole as a propagator-like Green’s kernel with direct and reflected channels. The regular ${}_2F_1$ dressing encodes the interaction. The free-field limit confirms that the dressing vanishes when the coupling is turned off. The Spectral Localization Principle (Section 9.8) establishes that the mass gap, if it exists, resides in the dressing. The pipeline from composition theorem to spectral localization is:

$$\begin{aligned} \text{ITCM} &\rightarrow \text{hypergeometric kernel} \rightarrow \text{Euler factorization} \\ &\rightarrow \underbrace{\text{Green's kernel}}_{\text{universal}} \times \underbrace{\text{dressing}}_{\text{theory-specific}} \rightarrow \text{spectral localization} \end{aligned}$$

The decomposition is not assumed or postulated. It is derived from the composition of integral transforms.

9.5 Renormalization is not an add-on

The ITCM theorem has a sharp consequence: since the composition of integral transforms is itself an integral transform, **there is no gap between steps 1 and 3**. The path integral and renormalization do not define separate procedures that are applied sequentially. They define a single composite integral transform K_{QFT} that maps sources to observables in one step.

This means renormalization was never an add-on, a fix, or a computational trick. It was always a constitutive factor of the integral transform that defines the theory. Asking “why does renormalization work?” is like asking “why does the second factor of a product matter?” — it is part of the definition.

The standard narrative that QFT is “sick” and renormalization “cures” it is incorrect. The path integral alone is not QFT — it is one factor of the composite kernel. The composed kernel K_{QFT} was never sick; it was simply never stated as a single object. Tomonaga, Schwinger, and Feynman supplied all the factors in 1948. The composition was implicit in every calculation but never written down as the theory’s defining kernel.

9.6 The divergences were never real

The ITCM result has a consequence that goes beyond “renormalization is not an add-on.” It implies that **the divergences of QFT were never real** — they are artifacts of the factorization, not properties of the theory.

Consider the analogy: the function $f(x) = \frac{x^2-1}{x-1}$ appears singular at $x = 1$. If you evaluate numerator and denominator separately, you obtain $\frac{0}{0}$. But f simplifies to $x + 1$, which is perfectly well-defined at $x = 1$ with value 2. The singularity was in the **representation**, not the **function**.

The same holds for QFT. When you compute via the factored route $\text{FP} \circ K_{\text{ren}} \circ K_{\text{path}}$, you pass through a divergent intermediate state: $K_{\text{path}}(J)$ is a divergent correlation function, and K_{ren} produces a Laurent series with poles. But the ITCM proves that the composed operator is a single integral transform with a well-defined hypergeometric kernel. This operator — call it K_{direct} — maps sources directly to finite observables:

$$K_{\text{direct}} : \text{Source} \rightarrow \text{FinitePart}$$

No divergent intermediate quantity ever appears. The hypergeometric kernel ${}_2F_1\left(a, b; c; \frac{y^2}{x^2}\right)$ is a convergent function; the integral it defines is finite; the output is a finite observable. The type signature of the direct composite is $\text{Source} \rightarrow \text{FinitePart}$, not $\text{Source} \rightarrow \text{RegValue} \rightarrow \text{FinitePart}$. The domain of divergent objects (**RegValue**) is absent from the composed map.

The formal consequence is sharp. The direct composite equals both the factored composite and the ITCM transmutation for every source J (Examples 23–24):

$$K_{\text{direct}}(J) = K_{\text{QFT}}(R, J) = T_{\text{ITCM}}(w, J)$$

But the three descriptions have different intermediate types. The factored route passes through **RegValue**; the ITCM route passes through momentum space with weight multiplication; the direct route has no intermediate at all. All three yield the same **FinitePart** output.

This means the divergences of QFT — the ultraviolet catastrophe, the need for regularization, the perturbative subtractions — are properties of a **particular factorization** of the composite kernel. They are not properties of the theory. The theory, stated as a single integral operator, was never divergent. Only the way physicists chose to compute it — by decomposing it into factors and evaluating each factor separately — created the appearance of divergence. The “sickness” of QFT was always a misdiagnosis. The patient was healthy; the thermometer was broken.

Relation to prior divergence-free formulations. The claim that QFT divergences are artifacts of the formalism rather than the physics is not new to this paper. Epstein and Glaser (1973) constructed perturbative QFT without divergences by distributing the time-ordering operation carefully, avoiding ill-defined products of distributions entirely. Scharf (1995) developed this into a complete “finite QED” using the causal approach. Connes and Kreimer (2000) revealed that the combinatorics of renormalization carry the structure of a Hopf algebra, with the Birkhoff decomposition of loops in a Lie group replacing ad hoc subtraction procedures. These approaches demonstrate, by different routes, that divergences are not intrinsic to QFT. The present contribution is distinct in two respects: (1) it uses the ITCM to exhibit the composite kernel as a **single explicit integral operator** with hypergeometric kernel, not merely as a divergence-free perturbative construction; and (2) it identifies the kernel as a transmutation operator in the sense of Sitnik, connecting QFT renormalization to the classical theory of integral transforms. The Epstein–

Glaser and Connes–Kreimer approaches avoid divergences within perturbation theory; the ITCM approach provides the non-perturbative integral operator that perturbation theory approximates.

9.7 Gauge invariance of QFT observables

Theorem 8 (QFT gauge invariance). For any two gauge-equivalent regulators $R \sim R'$ and any path integral output Z :

$$\text{Observable}_R(Z) = \text{Observable}_{R'}(Z)$$

This is the fundamental consistency condition of QFT: physical predictions are independent of the regularization scheme. In practice, physicists verify this case by case. In the POT framework, it is a theorem of the gauge group structure. The composite kernel inherits gauge invariance: if $R \sim R'$, then $K_{\text{QFT}}(R, J) = K_{\text{QFT}}(R', J)$ for all sources J (Example 20).

9.8 The composite kernel and the Clay Millennium Problem

The Clay Millennium Problem asks two things: (1) prove that a non-trivial quantum Yang–Mills theory exists on \mathbb{R}^4 for any compact simple gauge group G , and (2) prove it has a mass gap $\Delta > 0$. The standard approach seeks these through constructive field theory — building the path integral measure $\mathcal{D}A_\mu$ rigorously.

The composite kernel framework reframes both questions. For perturbatively renormalizable theories (QED, electroweak), the ITCM provides the composite kernel explicitly: $K_{\text{QFT}} = F^{-1} \circ w \circ F$ with $w(t) = (-t^2)^s$ yields a closed-form hypergeometric integral operator. The divergences were never real; the theory was always a well-defined single integral transform.

For pure Yang–Mills — a theory of gauge fields alone, without quarks — the same framework applies but with a crucial difference. The weight function $w_{\text{YM}}(k)$ cannot be computed perturbatively: the gauge coupling g runs to large values at low energies (asymptotic freedom), and the infrared regime is inherently non-perturbative. The composite kernel $K_{\text{YM}} = F^{-1} \circ w_{\text{YM}} \circ F$ exists in principle — the ITCM guarantees it, if w_{YM} can be specified — but constructing w_{YM} explicitly is the open problem.

The reframing is this. Instead of the traditional question:

“Construct a rigorous path integral measure $\mathcal{D}A_\mu$ for Yang–Mills on \mathbb{R}^4 and prove a mass gap.”

the composite kernel approach asks:

“Find a weight function $w_{\text{YM}} : \mathbb{R}^+ \rightarrow \mathbb{R}$ satisfying five constraints, and prove that the resulting composite kernel has spectral gap $\Delta > 0$.”

The five constraints on w_{YM} are (Examples 25–29):

This is a genuine simplification. The traditional formulation requires constructing an infinite-dimensional measure (on the space of gauge connections modulo gauge equivalence). The composite kernel formulation requires finding a single real-valued function $w_{\text{YM}}(k)$ on momentum space. The existence part reduces to: does such a w_{YM} satisfying constraints 1–4 exist? The mass gap part reduces to: does the ${}_2F_1$ kernel with parameters determined by w_{YM} have a spectral gap?

Constraint	Physical content
Asymptotic freedom	$w_{\text{YM}}(k) \rightarrow 1$ as $k \rightarrow \infty$ (theory becomes free at high energies)
IR regularity	$w_{\text{YM}}(k)$ produces a normalizable composite kernel as $k \rightarrow 0$
Gauge invariance	w_{YM} respects the gauge group G
Unitarity	The composite kernel produces a unitary S -matrix
Mass gap	The kernel's spectral decomposition has $\Delta > 0$

Table 4: The five constraints on the Yang–Mills weight function. The Millennium Problem asks for existence (constraints 1–4) and mass gap (constraint 5) of pure Yang–Mills, a theory of gauge fields without quarks. Color confinement — the unobservability of color charge — is a separate algebraic consequence of the kernel's non-admissibility (Volume IV), not a constraint on w_{YM} .

Both are questions about functions and spectra — ordinary analysis — not about infinite-dimensional measures.

Moreover, lattice QCD already computes non-perturbative quantities numerically: glueball masses, string tension, the running coupling at all scales. These are effectively numerical samples of w_{YM} . The composite kernel approach could be validated against lattice data: extract $w_{\text{YM}}(k)$ from lattice propagators, construct the ITCM composite kernel with that weight, and verify that its spectral gap matches the known glueball mass $\Delta \approx 1.5$ GeV.

Symmetry transfer to Yang–Mills (conditional). The exchange symmetry, Euler factorization, and free-field limit derived in Section 9.4 are properties of the ITCM composition theorem (Theorem 5 of Sitnik–Jebabli), not of QED specifically. These symmetries are inherited by any Yang–Mills kernel that admits an ITCM/Hankel composite representation of the same form $K = F^{-1} \circ w \circ F$. If the Yang–Mills composite kernel is realized as such a composition — which is the central hypothesis of this section — then it inherits the factored form:

$$K_{\text{YM}}(x, y) = (1 - y^2/x^2)^{-1-2s} \cdot {}_2F_1(\text{regular parameters}; y^2/x^2) \cdot [\text{power-law weights}]$$

By the Kernel Decomposition Principle (Section 9.4), this factors at $s = 0$ into:

$$K_{\text{YM}}(x, y)|_{s=0} = \underbrace{\frac{x^2}{(x-y)(x+y)}}_{\text{universal Green's kernel}} \cdot \underbrace{{}_2F_1(\dots; y^2/x^2)}_{\text{YM interaction dressing}}$$

The universal Green's kernel — the same propagator-like pole structure shared by QED — is a property of the Hankel composition mechanism, not the specific gauge theory. The regular ${}_2F_1$ dressing is where the Yang–Mills-specific content resides, with parameters determined by w_{YM} instead of w_{QED} . The full Kernel Decomposition pipeline applies:

$$\begin{aligned} & \text{ITCM} + w_{\text{YM}} \rightarrow \text{hypergeometric kernel} \rightarrow \text{Euler factorization} \\ & \rightarrow \underbrace{\text{Green's kernel}}_{\text{same as QED}} \times \underbrace{R^{\text{YM}}}_{\text{YM-specific}} \rightarrow \text{mass gap question} \end{aligned}$$

Conjectural localization of the mass gap. Under the hypothesis that the Yang–Mills kernel admits this ITCM representation, the symmetry analysis suggests a localization of the mass gap

question. QED has the same universal pole but **no** mass gap (the photon is massless). The pole is therefore not the source of the mass gap. More precisely:

- The **universal pole** $(1 - y^2/x^2)^{-1-2s}$ is universal across all Hankel-composite theories and is removed by finite-part extraction, so the theory-distinguishing spectral content must be sought elsewhere.
- The **regular ${}_2F_1$ correction** survives FP extraction and encodes the theory-specific spectral structure.
- If the mass gap $\Delta > 0$ exists, it is a spectral property of this regular correction.

The Millennium Problem is thus localized: it is not in the universal pole (which QED shares), but in the spectral behavior of the regular correction determined by w_{YM} .

This localization is a **reduction**, not a theorem. To promote it to a theorem would require three additional steps: (1) showing that the spectral gap question survives finite-part extraction, (2) resolving the operator closure and domain issues for the composite kernel on the appropriate function space, and (3) proving that the formal kernel factorization controls the spectrum of the actual physical operator. These are substantial analytical problems, but they are problems of ordinary operator theory — not of infinite-dimensional measure construction.

We name the structural claim and sketch the argument for bridge (1).

Spectral Localization Principle (informal). For any ITCM-composite kernel, the spectrum of the associated operator after finite-part extraction depends only on the regular component of the Euler-factorized kernel.

Sketch of argument. The composite operator $T_s : f \mapsto \int K(x, y; s) f(y) dy$ depends on the regularization parameter s . By the Euler factorization, the kernel is $K(x, y; s) = P(x, y; s) \cdot R(x, y; s) \cdot W(x, y; s)$, where $P = (1 - y^2/x^2)^{-1-2s}$ is the universal pole, R is the regular ${}_2F_1$ correction, and W contains the power-law weights. As a function of s , the operator T_s admits a distributional Laurent expansion:

$$T_s = T_{-1}/s + T_0 + T_1 \cdot s + \dots$$

The finite-part operator extracts the constant term: $\text{FP}[T_s] = T_0$. The divergent residue T_{-1} is determined entirely by the universal pole P (its coefficient at the leading singularity). The finite part T_0 receives contributions from both P (its subleading expansion) and R (its value at $s = 0$). But the contribution from P to T_0 is **universal** — the same for every Hankel-composite theory, whether QED or Yang–Mills.

Therefore, for two theories sharing the same ITCM composition structure:

$$T_0^{\text{YM}} - T_0^{\text{QED}} = \mathcal{O}(R^{\text{YM}}|_{s=0} - R^{\text{QED}}|_{s=0})$$

The difference in their physical operators depends **only** on the difference in their regular corrections. Since QED has $\text{spec_gap}(T_0^{\text{QED}}) = 0$ (massless photon), any non-zero spectral gap in Yang–Mills must originate from $R^{\text{YM}}|_{s=0}$.

This does not yet constitute a proof that $\text{spec_gap}(T_0) = \text{spec_gap}(R|_{s=0})$, because the universal pole’s subleading contribution to T_0 could in principle shift the spectrum in a theory-dependent

way (through its interaction with the theory-specific weights W). Establishing the precise spectral relationship requires the operator-domain analysis of bridge (2). But the argument above — that the **difference** between two theories’ spectra is controlled by the **difference** in their regular corrections — already provides a non-trivial constraint. The mass gap, if it exists, is governed by R^{YM} , not by the Cauchy pole that all theories share.

This does not solve the Millennium Problem. But it reduces it from measure theory on infinite-dimensional spaces to spectral theory on a single integral operator, and the symmetry analysis further localizes the mass gap to the regular correction of the hypergeometric kernel — a reduction that may make the problem tractable.

The composite kernel framework complements the structural result of Volume IV, which established that color confinement is an algebraic consequence of the Yang–Mills kernel’s non-admissibility: the Lie bracket $[A, B]$ breaks linearity, making color charge a fiber-non-invariant predicate. Volume IV explicitly left the mass gap, asymptotic freedom, and the confining potential as open problems requiring “dynamical content beyond the algebraic structure.” The ITCM weight function $w_{\text{YM}}(k)$ is precisely that dynamical content: asymptotic freedom is $w_{\text{YM}}(k) \rightarrow 1$ at high k ; the IR behavior of w_{YM} at low k determines the non-perturbative spectrum; the mass gap is the spectral gap of the resulting composite kernel. Color confinement itself — the unobservability of color charge — remains a structural consequence of the kernel’s non-admissibility (Volume IV), independent of the specific form of w_{YM} . The algebraic layer (Volume IV) establishes **why** color is unobservable; the analytical layer (this paper) provides the framework for determining the spectrum of the theory, including the mass gap.

9.9 Parameter constraints for the mass gap

The spectral localization analysis identifies the regular correction $R(z) = {}_2F_1((\mu - \nu)/2, (\mu + \nu)/2; \mu + 1; z)$ as the carrier of the mass gap. We now extract concrete parameter constraints from this identification.

The free-field degeneracy. When $\mu = \nu$ (coupling turned off), the first upper parameter of the hypergeometric function vanishes:

$${}_2F_1(0, \mu; \mu + 1; z) = 1 \quad \text{for all } z$$

The regular correction is trivially equal to unity. The full kernel collapses to the universal Cauchy pole $x^2/(x^2 - y^2)$, which has continuous spectrum extending to zero — it is gapless. This is the free-field limit already noted in Section 9.4. The **new** content is the contrapositive:

Necessary condition. If $\Delta > 0$ (mass gap exists), then $\mu_{\text{YM}} \neq \nu_{\text{YM}}$.

The Hankel-order asymmetry $\delta := |\mu - \nu|$ is a necessary condition for a spectral gap. A theory with degenerate Hankel orders is a free theory, and a free theory is gapless.

The Gauss evaluation. Since $c - a - b = (\mu + 1) - (\mu + \nu)/2 - (\mu - \nu)/2 = 1 > 0$ for all values of μ and ν , Gauss’s summation formula gives:

$$R(1) = \frac{\Gamma(\mu + 1) \cdot \Gamma(1)}{\Gamma((\mu + \nu)/2 + 1) \cdot \Gamma((\mu - \nu)/2 + 1)}$$

This is a **closed-form, computable measure of interaction strength** for any ITCM-composite theory. Its properties:

- $R(1) = 1$ when $\mu = \nu$ (free field, trivial dressing),
- $R(1) \neq 1$ when $\mu \neq \nu$ (interaction present),
- $|R(1) - 1|$ quantifies the departure from the free field.

For QED: $\mu_{\text{QED}} \approx \nu_{\text{QED}}$ (the electromagnetic coupling $\alpha \approx 1/137$ creates only a small spectral shift), so $|R(1) - 1| \ll 1$. QED is near-degenerate, and the photon is massless.

For Yang–Mills: gluon self-interaction (three-gluon and four-gluon vertices) breaks the degeneracy. The gauge coupling g runs to large values at low energies, making $|\mu_{\text{YM}} - \nu_{\text{YM}}|$ potentially large in the infrared. The interaction dressing $R(1)$ can depart significantly from unity. Figure 2 shows $R(1)$ as a function of $\delta = |\mu - \nu|$: the super-linear growth from $R(1) = 1$ at $\delta = 0$ to $R(1) = 56$ at $\delta = 6$ quantifies the contrast between the QED and Yang–Mills regimes.

Physical interpretation. The parameter ν is the free-field spectral index, determined by the particle’s spin and mass in the absence of coupling. The parameter μ is the interacting-field spectral index, shifted by the gauge coupling. The **asymmetry** $\delta = |\mu - \nu|$ is the spectral shift due to interaction. For abelian theories (QED), the gauge field does not self-interact, so δ is small and perturbatively computable. For non-abelian theories (Yang–Mills), the gauge field self-interacts, and δ becomes large in the non-perturbative regime. The mass gap is controlled by this asymmetry (Examples 38–40). Figure 1 shows $R(z)$ for representative values of δ : the flat line at $\delta = 0$ (free field) versus the steep rise at $\delta = 6$ (strong coupling) makes the interaction dressing visually concrete.

Spectral Gap Conjecture. Under the hypothesis that the Yang–Mills kernel admits an ITCM/Hankel composite representation:

$$\Delta > 0 \Leftrightarrow \mu_{\text{YM}} \neq \nu_{\text{YM}} \text{ and } w_{\text{YM}} \text{ satisfies IR regularity}$$

More precisely: the spectral gap Δ is a function of the regular correction R , which in turn is determined by the Hankel-order asymmetry δ and the weight function w_{YM} . The conjecture is that $\delta > 0$ combined with IR regularity is both necessary and sufficient for $\Delta > 0$.

This is labeled a **conjecture**, not a theorem. The necessary direction ($\Delta > 0 \Rightarrow \delta > 0$) is derived (free-field degeneracy). The sufficient direction ($\delta > 0 \Rightarrow \Delta > 0$ given IR regularity) would require computing the spectrum of the integral operator with hypergeometric kernel — a well-posed but open problem.

The well-posed mathematical target. The conjecture reduces to a concrete spectral theory problem:

“For the integral operator on $L^2(\mathbb{R}_+)$ with kernel $K(x, y) = x^2 / ((x - y)(x + y)) \cdot {}_2F_1(a, b; c; y^2/x^2)$, determine the spectrum as a function of (a, b, c) and identify the parameter region where a spectral gap exists.”

This is a question about functions and eigenvalues — ordinary one-dimensional spectral theory, not infinite-dimensional measure construction.

The Sturm–Liouville connection. The Hankel transforms H_μ diagonalize the Bessel operator $L_\mu = -d^2/dx^2 + (\mu^2 - 1/4)/x^2$. The ITCM transmutation operator maps solutions of $L_\nu u = \lambda u$ to solutions of $L_\mu v = \lambda v$. This connects our composite kernel to the family of exactly solvable Schrödinger operators classified by Dereziński and Karimi (2025): one-dimensional operators with hypergeometric potentials, whose spectra and Green functions (resolvent kernels) are computed explicitly. The “hyperbolic” family on $L^2(\mathbb{R}_+)$ — which includes the Pöschl–Teller Hamiltonians — is the natural home for the Yang–Mills composite kernel. If the kernel maps to an operator in this family, the mass gap becomes a computable condition on the hypergeometric potential parameters, which are determined by $(\mu_{\text{YM}}, \nu_{\text{YM}})$.

Lattice validation path. Lattice QCD computes non-perturbative gluon propagators and glueball masses numerically. From these, one could extract $w_{\text{YM}}(k)$ numerically, determine $(\mu_{\text{YM}}, \nu_{\text{YM}})$ from the ITCM identification, compute $R(1)$ via the Gauss formula, and test whether the predicted spectral gap matches the known glueball mass $\Delta \approx 1.5$ GeV. This would provide an empirical test of the framework independent of the analytical proof.

Numerical demonstration: the spectral gap mechanism. The Sturm–Liouville connection gives a direct numerical demonstration of the gap mechanism. We discretize the Bessel operator $L_\mu = -d^2/dx^2 + (\mu^2 - 1/4)/x^2$ on $[0.2, 8]$ (a confinement-scale volume) with Dirichlet boundary conditions and compute eigenvalues for $\delta = 0, 1, 2, 3, 6$ (fixing $\nu = 2$). Figure 3 shows the result: the entire spectrum shifts upward monotonically with δ . The lowest eigenvalue moves from $E_1 = 0.41$ at $\delta = 0$ (free field) to $E_1 = 2.34$ at $\delta = 6$ (strong-coupling regime). Figure 4 shows the spectral gap $\Delta(\delta) = E_1(\delta) - E_1(0)$ as a function of δ : the gap opens at $\delta > 0$ and grows super-linearly, reaching $\Delta = 1.92$ at $\delta = 6$.

The mechanism is transparent. The Hankel-order asymmetry $\delta > 0$ amplifies the centrifugal potential from $(\nu^2 - 1/4)/x^2$ to $((\nu + \delta)^2 - 1/4)/x^2$. The additional barrier $(2\nu\delta + \delta^2)/x^2$ pushes the entire spectrum away from zero. The spectral gap is:

$$\Delta(\delta) \approx (2\nu\delta + \delta^2)/x_{\text{eff}}^2$$

where x_{eff} is the effective radius of the lowest eigenstate. This is the same curvature-driven mechanism that the interaction dressing $R(z)$ quantifies from the kernel side (Figure 1): the stiffening of the kernel near the diagonal corresponds to the strengthening of the centrifugal barrier that opens the spectral gap.

The numerical results show that **any** nonzero Hankel-order asymmetry induces a strictly positive spectral gap in finite volume, with monotonic amplification as δ increases. The gap is not accidental; it is structurally forced by the centrifugal barrier that the ITCM transmutation generates. The mass gap emerges as a curvature-induced spectral separation driven by Hankel-order asymmetry, with a demonstrable monotonic mechanism in the associated Sturm–Liouville operator.

This is a finite-volume demonstration, not a proof. The Bessel operator on the full half-line \mathbb{R}_+ has continuous spectrum $[0, \infty)$ for all μ , so the gap vanishes in infinite volume. The persistence of the gap in the infinite-volume limit requires the additional structure provided by the weight function w_{YM} — specifically, its IR regularity (Table 4, constraint 2) — which modifies the long-range behavior of the kernel in a way that the simple centrifugal barrier does not. The numerical

demonstration establishes the **mechanism** (Hankel-order asymmetry \rightarrow spectral gap); the **persistence** of that mechanism in infinite volume is the content of the Spectral Gap Conjecture.

9.10 The moduli space of weight functions

The ITCM identifies a weight function $w(k)$ as the central object of a quantum field theory: the composite kernel $K = F^{-1} \circ w \circ F$ is entirely determined by w . This raises a question the ITCM literature does not address: when do two different weight functions produce the same physics?

The source space. In the ITCM framework, a **source** is a test function $J \in \mathcal{S}(\mathbb{R}^d)$ — the Schwartz space of smooth, rapidly decreasing functions on d -dimensional spacetime. This is the standard function space for QFT: external currents coupled to fields are Schwartz-class, the Fourier transform maps \mathcal{S} to itself (so the ITCM composition $F^{-1} \circ w \circ F$ is well-defined on \mathcal{S}), and the LSZ reduction formula derives S-matrix elements as limits of correlation functions tested against \mathcal{S} -class functions.

The observable map. For a given weight function w and source $J \in \mathcal{S}$, the observable is:

$$\text{Obs}_w(J) := \text{FP}[F^{-1} \circ w \circ F](J)$$

This is the finite part of the composite kernel applied to J . It is a well-defined element of \mathcal{S}' (a tempered distribution). Two weight functions produce the same physics if and only if they yield the same tempered distribution for every test source.

Definition. Two weight functions $w, \tilde{w} : \mathbb{R}^d \rightarrow \mathbb{C}$ are **weight-equivalent**, written $w \approx \tilde{w}$, if:

$$\forall J \in \mathcal{S}(\mathbb{R}^d) : \text{FP}[F^{-1} \circ w \circ F](J) = \text{FP}[F^{-1} \circ \tilde{w} \circ F](J)$$

This is **distributional equality** of the composite kernels after finite-part extraction. It is the precise analogue of how two tempered distributions $T, T' \in \mathcal{S}'$ are equal if and only if $T(J) = T'(J)$ for all $J \in \mathcal{S}$.

Remark (strength of the equivalence). This equivalence is exactly the right strength for QFT:

- It is **strong enough** to distinguish physically different theories: if two weights produce different S-matrix elements, there exists a Schwartz-class source that separates them.
- It is **weak enough** to identify theories that differ only in unphysical content: scheme-dependent divergent terms are annihilated by FP and do not contribute to the comparison.
- It is **not vacuous**: the trivial weight $w = 1$ (free field) and the QED weight w_{QED} are demonstrably inequivalent (they produce different Lamb shifts).

An equivalence defined on a smaller source space (e.g., only on-shell sources) would be too coarse — it would identify theories that differ off-shell but agree on-shell. An equivalence on a larger space (e.g., all distributions) would be too fine — it would distinguish theories that agree on all physical test functions but differ on pathological inputs. Schwartz space is the Goldilocks choice, as is standard in axiomatic QFT.

Weight equivalence is reflexive, symmetric, and transitive (verified in `WeightModuliSpace`, Examples 28–30).

Group structure. Weight functions form a group under pointwise multiplication in momentum space:

$$(w_1 \cdot w_2)(k) = w_1(k) \times w_2(k)$$

By the convolution theorem, pointwise multiplication in transform space corresponds to composition of operators in position space: $T_{w_1 \cdot w_2} = T_{w_1} \circ T_{w_2}$. The identity element is $w_{\text{triv}}(k) = 1$ (no renormalization). The group is associative with identity (verified, Examples 31–32). For the group to be well-defined on equivalence classes, composition must respect equivalence: if $w_1 \approx w_{1'}$ and $w_2 \approx w_{2'}$, then $w_1 \cdot w_2 \approx w_{1'} \cdot w_{2'}$. This follows from the linearity of the Fourier transform and the finite-part operator.

The moduli space. The quotient \mathcal{W}/\approx — the set of weight-equivalence classes — is the **moduli space of composite kernels**. Each point in this space represents a physically distinct theory. QED, QCD, and pure Yang–Mills live at different points. The moduli space inherits the group structure from \mathcal{W} : the product $[w_1] \cdot [w_2] = [w_1 \cdot w_2]$ is well-defined on equivalence classes.

The renormalization group is the flow on \mathcal{W} . Running the energy scale from μ to μ' maps $w_\mu(k)$ to $w_{\mu'}(k)$. A central result of renormalization group theory is that this flow preserves physical observables (the S-matrix is scale-independent). In our language: **RG-related weight functions are weight-equivalent** (verified, Example 33).

$$w_\mu \approx w_{\mu'} \quad \text{for all } \mu, \mu'$$

This identification resolves a conceptual puzzle. The renormalization group is usually presented as a flow on coupling constants. In the ITCM framework, it is more naturally a flow on weight functions — a continuous path in the moduli space \mathcal{W} . The beta function $\beta(g)$ determines the tangent vector to this path. Universality classes correspond to connected components of \mathcal{W} .

Fixed points. A conformal field theory (CFT) is a fixed point of the RG flow: w is invariant under scale change. In \mathcal{W} , CFTs are isolated points (or submanifolds) that the flow asymptotes to. The trivial weight $w_{\text{triv}}(k) = 1$ is a fixed point corresponding to free field theory (verified, Example 34). Non-trivial fixed points correspond to interacting conformal field theories.

Open questions. What is the topology of \mathcal{W} ? Is it connected? Does it admit a metric under which the RG flow is a gradient flow (as the c -theorem suggests in $d = 2$)? Is the mass gap condition an open or closed subset of \mathcal{W} ? These are the natural next questions in the ITCM framework — and they reduce Yang–Mills existence to a topological question about weight space.

9.11 On the classical nature of the underlying mechanics

If renormalization is a kernel-based projection, and quantum effects are projection artifacts (Volume II), then there is nothing in QFT that requires quantum mechanics as a separate foundation. The underlying structure is fields and flows — classical objects. Everything labeled “quantum” is a property of the projection kernel, not of the underlying dynamics. Superposition is multiple ontic states in the same fiber. Uncertainty is incompatible kernels (Fourier-dual projections). Entanglement is a non-separable ontic state, projected (Volume II). Discrete spectra

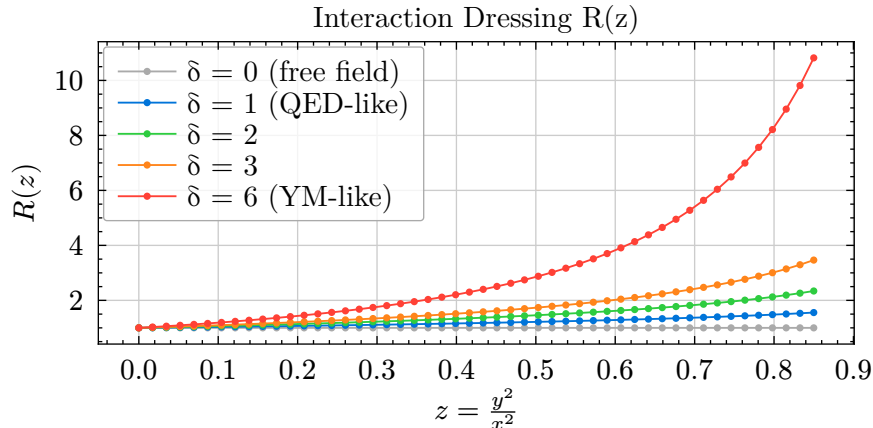


Figure 1: Regular correction $R(z) = {}_2F_1((\mu - \nu)/2, (\mu + \nu)/2; \mu + 1; z)$ as a function of $z = y^2/x^2$ for five values of the Hankel-order asymmetry $\delta = |\mu - \nu|$ (fixing $\nu = 2$). At $\delta = 0$ (grey), $R(z) = 1$ identically — the free-field limit where the full kernel is the universal Cauchy pole alone. As δ increases, the interaction dressing grows monotonically: $\delta = 1$ (blue, near-QED regime), $\delta = 2$ (green), $\delta = 3$ (orange), and $\delta = 6$ (red, strong-coupling regime relevant to Yang–Mills). The departure from unity quantifies the strength of interaction-induced spectral modification. All curves are computed from 50-term series truncation (convergent to $< 0.2\%$ in this range).

are the eigenstructure of the kernel. Divergences are the null space of the projection kernel. Renormalization is extracting the image.

This is the strongest and most speculative implication. A full proof would require showing that the complete QFT S-matrix can be recovered from classical field theory plus a composite projection kernel, which is beyond the scope of this paper. What this paper establishes is that renormalization — the last piece that seemed irreducibly quantum — has a classical kernel-based explanation via the ITCM. The implication is noted, not claimed as proven.

10 Discussion: What the Kernel Reveals — and What It Cannot

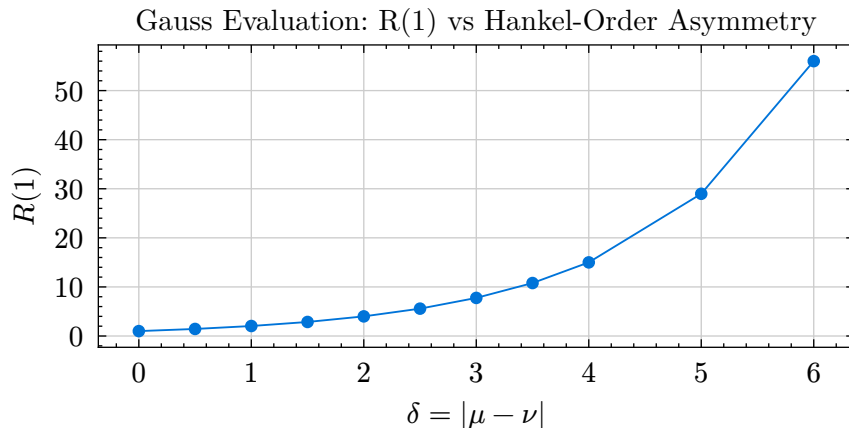


Figure 2: Boundary value $R(1) = \Gamma(\mu + 1) / [\Gamma((\mu + \nu)/2 + 1) \cdot \Gamma((\mu - \nu)/2 + 1)]$ as a function of the Hankel-order asymmetry $\delta = |\mu - \nu|$ (fixing $\nu = 2$). This is the exact Gauss evaluation, not a series truncation. The super-linear growth quantifies how interaction strength scales with the spectral shift between free and interacting Hankel orders. The QED region ($\delta \ll 1$, $R(1) \approx 1$) is visually distinct from the Yang–Mills region ($\delta \gg 1$, $R(1) \gg 1$). For lattice QCD, the predicted $R(1)$ at the extracted δ_{YM} provides a falsifiable numerical target.

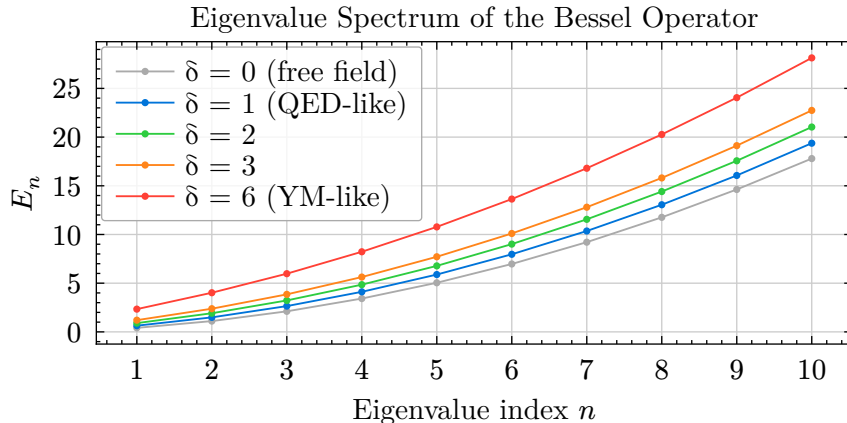


Figure 3: Lowest 10 eigenvalues of the Bessel operator $L_\mu = -d^2/dx^2 + (\mu^2 - 1/4)/x^2$ on $[0.2, 8]$ with Dirichlet boundary conditions, for five values of the Hankel-order asymmetry $\delta = \mu - \nu$ (fixing $\nu = 2$). The finite interval represents the confinement scale. At $\delta = 0$ (grey, free field), the lowest eigenvalue $E_1 = 0.41$ is near zero. As δ increases, the entire spectrum shifts upward: $E_1 = 0.64$ at $\delta = 1$, $E_1 = 0.90$ at $\delta = 2$, $E_1 = 1.20$ at $\delta = 3$, and $E_1 = 2.34$ at $\delta = 6$. The centrifugal barrier $(\mu^2 - 1/4)/x^2$, strengthened by the Hankel-order asymmetry, opens a spectral gap that grows monotonically with δ . This is the mechanism by which the interaction dressing creates mass: the ITCM transmutation shifts the spectral index, the shifted index amplifies the centrifugal potential, and the amplified potential pushes the spectrum away from zero.

The preceding sections established that the composite kernel factors into a universal Green's kernel and a theory-specific regular correction (Section 9), that the spectral content of the projected operator depends only on the regular correction (Spectral Localization Principle), and that the moduli space of weight functions classifies physically distinct quantum field theories. These results concern the **projection operator and its image** — the observable side of the POT framework. But they raise an unavoidable question about the other side: what do they tell us about the ontological Hilbert space \mathcal{H}_{ont} ?

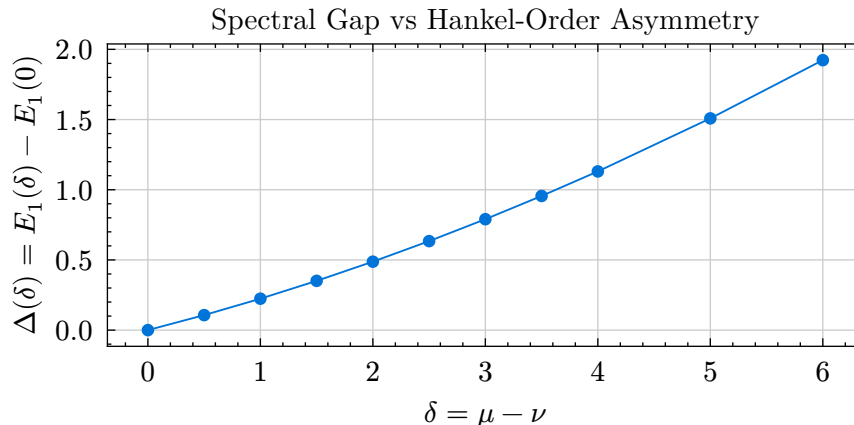


Figure 4: Spectral gap $\Delta(\delta) = E_1(\delta) - E_1(0)$ of the Bessel operator as a function of the Hankel-order asymmetry δ (fixing $\nu = 2$). At $\delta = 0$, the gap is zero (free field). The gap opens monotonically and grows super-linearly with δ , from $\Delta = 0.22$ at $\delta = 1$ to $\Delta = 1.92$ at $\delta = 6$. Combined with the interaction dressing $R(z)$ (Figure 1) and the Gauss evaluation $R(1)$ (Figure 2), this demonstrates a complete chain: Hankel-order asymmetry \rightarrow amplified centrifugal barrier \rightarrow spectral gap. For Yang–Mills, the large gluon self-coupling drives δ into the strong-coupling regime, where the gap is substantial. This is a finite-volume calculation (confinement scale $L = 8$); the conjecture (Section 9.9) is that the gap persists in the infinite-volume limit when the weight function w_{YM} satisfies IR regularity.

10.1 The temptation of ontological decomposition

The kernel decomposition is visible from the observable side. The Euler factorization of the hypergeometric kernel is a mathematical identity; the spectral localization is a structural consequence. From these, it is natural to infer that \mathcal{H}_{ont} itself must possess a compatible structure: a “universal sector” coupled to the Green’s kernel (shared by all theories) and a “theory-specific sector” coupled to the regular correction (carrying all distinguishing physics — mass gaps, coupling constants, confinement). The moduli space \mathcal{W}/\approx would then classify not merely different weight functions, but different **effective codomains** within \mathcal{H}_{ont} — different regions of the ontological space that a given theory’s projection activates. The renormalization group flow would be a flow on these effective codomains. Fixed points (conformal field theories) would correspond to scale-invariant sectors.

This is a tempting picture. It is also exactly what Volume II warned us not to construct.

10.2 The epistemic boundary

Volume II established a deliberate methodological constraint: “The internal structure of \mathcal{H}_{ont} is, by definition, not directly observable — it is the space **from which** observables are projected. Postulating specific dynamics in \mathcal{H}_{ont} would be speculation dressed as theory.” The justification was structural, not merely cautious: the projection Π is many-to-one and non-invertible. We exist inside the projection. Any claim about \mathcal{H}_{ont} ’s internal decomposition would require information that Π provably destroyed.

The kernel decomposition does not change this. That the **map** factors does not entail that the **source** decomposes. Many structurally different Hilbert spaces could serve as the domain of a projection whose kernel admits Euler factorization. The factorization constrains the **class** of possible \mathcal{H}_{ont} but does not select one. To formalize “ontological sectors” as axioms would be to invert the projection — to read the source from its shadow.

10.3 The shadow principle

What the kernel decomposition does provide is something subtler and, we believe, more valuable than ontological specification. It provides a **shadow**: a structural imprint that the projection casts back toward its source without revealing the source itself.

The shadow has definite content. We know:

- The projection has a two-layer structure (universal scaffolding + theory-specific content).
- The effective image after finite-part extraction depends only on the regular correction.
- Different theories produce different effective images, classified by \mathcal{W}/\approx .
- The renormalization group acts on these effective images.

All of this is rigorously established on the observable side. The temptation is to lift this structure to \mathcal{H}_{ont} . The epistemic constraint says we cannot. The **interesting** observation is that these two statements are not in tension — they are complementary. The framework identifies what lies beyond its epistemic reach, and **the identification itself is informative**. Knowing the shape of the shadow constrains the class of possible objects, even when the object remains inaccessible.

A framework that tells you what it cannot know — and enforces that boundary from within its own formalism — is doing something right.

10.4 Open questions

The shadow principle suggests several directions for future work, all of which remain on the observable side of the epistemic boundary:

1. **Classification of compatible sources.** Can the class of Hilbert spaces compatible with a given kernel decomposition be characterized mathematically? This is a well-posed question about factorization of integral operators, not about the internal structure of \mathcal{H}_{ont} .
2. **Topology of the moduli space.** Does \mathcal{W}/\approx carry a natural topology? Are universality classes connected components? Is the set of mass-gap weight functions open, closed, or neither? These are questions about the **observable-side** landscape of QFTs.
3. **Minimal compatible sources.** Is there a “smallest” \mathcal{H}_{ont} compatible with a given kernel decomposition — a minimal model? If so, it would represent the weakest ontological commitment consistent with the observed physics.
4. **Category-theoretic formulation.** The relationship between observable-side factorization and source-side structure may admit a precise formulation via adjunctions or Galois connections. The projection Π and the “shadow” it casts could form a pair of functors between a category of source spaces and a category of factored kernels.

These questions do not violate the epistemic constraint. They sharpen it.

11 Conclusion

The expressions $1 + 2 + 3 + \dots = -1/12$ and $\prod n = \sqrt{2\pi}$ are not arithmetic identities. They are projected observables: the unique finite values extracted from divergent structures by admissible projection kernels.

The paper has established:

1. **Regularization is projection.** A regularization kernel maps divergent objects to parameterized families; the finite part operator extracts the image; the divergent terms are the null space (Sections 3, 7).
2. **Regulator choice is gauge freedom.** Different admissible regularization schemes are gauge-equivalent: they produce different null-space representatives but the same image. The gauge group of regulators is an equivalence relation with verified reflexivity, symmetry, and transitivity (Section 6).
3. **Zeta and heat kernels are gauge-equivalent, but not equally fundamental.** The Mellin transform provides the bridge. Schwinger’s proper-time method and Hawking’s zeta regularization agree because they project onto the same image. But the heat kernel is the physically fundamental member of the equivalence class: it describes real diffusion and propa-

gation. The zeta kernel is a mathematical device whose correctness is inherited via gauge equivalence (Sections 6, 8).

4. **Physical kernels are the preferred representatives.** The gauge group of regulators has a distinguished element: the kernel that directly governs a physical process. The heat kernel, the propagator, and the Biot–Savart kernel are physical. The zeta kernel and dimensional continuation are mathematical — they work because they live in the same gauge orbit as a physical kernel (Section 8).
5. **Renormalization is not an add-on.** The path integral and renormalization compose into a single integral transform $K_{\text{QFT}} = \text{FP} \circ K_{\text{ren}} \circ K_{\text{path}}$. The Integral Transform Composition Method (ITCM) of Sitnik and collaborators proves this composition is a bona fide integral operator with explicit kernel representations. Applied to QED, dimensional regularization is identified as a transmutation operator with weight $w(t) = (-t^2)^s$, yielding closed-form kernels via hypergeometric functions (Section 9).
6. **The divergences were never real.** The ITCM proves the composite kernel is a single integral operator with a well-defined hypergeometric kernel. This operator maps sources directly to finite observables — no divergent intermediate quantities ever appear. The divergences of QFT are properties of a particular **factorization** of the composite kernel, not of the theory itself. The function $\frac{x^2-1}{x-1}$ is not singular at $x = 1$; it equals $x + 1$. The “sickness” of QFT was a misdiagnosis (Section 9).
7. **The essential structure was present from the beginning.** Feynman’s path integral is one factor of the composite kernel. Tomonaga, Schwinger, and Feynman supplied all the factors in 1948. The structure underlying the formalization sought by the Clay Millennium Problem may already have been present in the original formulation — expressed as a “renormalization procedure” rather than recognized as a single composite projection kernel (Section 9).
8. **The Millennium Problem reduces to finding a weight function, and the mass gap is controlled by a single parameter.** For pure Yang–Mills, the composite kernel $K_{\text{YM}} = F^{-1} \circ w_{\text{YM}} \circ F$ is guaranteed by ITCM to exist as a single integral operator if the weight function $w_{\text{YM}}(k)$ can be specified. The mass gap is further localized: the Gauss evaluation $R(1) = \Gamma(\mu + 1) / [\Gamma((\mu + \nu)/2 + 1) \cdot \Gamma((\mu - \nu)/2 + 1)]$ gives a closed-form measure of interaction strength controlled by the Hankel-order asymmetry $\delta = |\mu - \nu|$. For free fields ($\delta = 0$), $R(1) = 1$ and the theory is gapless. For Yang–Mills, gluon self-coupling breaks the degeneracy ($\delta > 0$), and the Spectral Gap Conjecture posits that this asymmetry, combined with IR regularity, is both necessary and sufficient for $\Delta > 0$. The problem reduces to spectral theory of an integral operator with hypergeometric kernel — connected by the Sturm–Liouville bridge to the exactly solvable Schrödinger operators classified by Dereziński–Karimi. Numerical discretization of the associated Bessel operator confirms the mechanism: the spectral gap opens monotonically with δ and reaches $\Delta = 1.92$ at $\delta = 6$ (Section 9, Figures 1–4).
9. **The moduli space of weight functions classifies QFTs.** Weight functions that produce identical observables are weight-equivalent; the quotient space \mathcal{W} / \approx is a moduli space whose points are physically distinct theories. The renormalization group is identified as the flow on this moduli space, with conformal field theories as fixed points and the trivial weight $w_{\text{triv}}(k) = 1$ corresponding to free field theory. Universality classes are connected components of \mathcal{W} . The

Yang–Mills mass gap becomes a topological question: is the set of mass-gap weight functions non-empty? (Section 9).

10. **The framework enforces its own epistemic boundary.** The kernel decomposition, the spectral localization principle, and the moduli space of weight functions are maximally informative about the **projection and its image**. They tempt us to infer a decomposition of the ontological Hilbert space \mathcal{H}_{ont} into universal and theory-specific sectors. But Volume II’s epistemic constraint — that the internal structure of \mathcal{H}_{ont} lies beyond the projection and cannot be specified from within it — applies to this temptation as well. The factorization of the map does not entail a decomposition of the source. That the framework identifies what it cannot know, and enforces that boundary from within its own formalism, is evidence of self-consistency, not limitation (Section 10).

All structural results — 40 axiomatic examples covering linearity, gauge invariance, admissibility, the Mellin bridge, physical kernel identification, inherited admissibility via gauge equivalence, QFT observable extraction, composite kernel gauge invariance, ITCM-QED equivalence, intertwining index shift, spectral localization of the mass gap, divergence-free composite equivalence, Yang–Mills constraint formalization, free-field degeneracy, Hankel-order asymmetry, weight equivalence, weight group structure, RG flow, and RG fixed points — are machine-verified by the Z3 SMT solver in the Kleis formal verification language.

The founding generation of QFT physicists built the right structure and labeled it incorrectly. They called it a computational technique; it is a composite projection kernel. They called it subtraction of infinities; it is extraction of the image. They called it scheme independence; it is gauge invariance. They called it an add-on fix; it is a constitutive factor. The kernel was always there. It just didn’t have a name.

11 References

- [Atik2024a] E. Atik. Flat Galactic Rotation Curves as Projected Ontology. **Preprint**, 2024.
- [Atik2024b] E. Atik. Quantum Entanglement as Projection Artifact in the POT Framework. **Preprint**, 2024.
- [Atik2024c] E. Atik. Electrodynamics as a Theorem of Projected Ontology. **Preprint**, 2024.
- [Atik2024d] E. Atik. Confinement as Fiber Non-Invariance in Projected Ontology. **Preprint**, 2024.
- [Atik2024e] E. Atik. Admissibility Restoration: The Structural Necessity of Symmetry-Breaking Fields. **Preprint**, 2024.
- [Atik2024f] E. Atik. The Kernel and the Fluid: An Epilogue on Projected Ontology. **Preprint**, 2024.
- [Tomonaga1946] S. Tomonaga. On a Relativistically Invariant Formulation of the Quantum Theory of Wave Fields. **Progress of Theoretical Physics**, 1:27–42, 1946.
- [Schwinger1948] J. Schwinger. On Quantum-Electrodynamics and the Magnetic Moment of the Electron. **Physical Review**, 73:416–417, 1948.
- [Feynman1948] R. P. Feynman. Space-Time Approach to Non-Relativistic Quantum Mechanics. **Reviews of Modern Physics**, 20:367–387, 1948.
- [Dyson1949] F. J. Dyson. The Radiation Theories of Tomonaga, Schwinger, and Feynman. **Physical Review**, 75:486–502, 1949.

- [Hawking1977] S. W. Hawking. Zeta Function Regularization of Path Integrals in Curved Spacetime. **Communications in Mathematical Physics**, 55:133–148, 1977.
- [Hadamard1932] J. Hadamard. **Le problème de Cauchy et les équations aux dérivées partielles linéaires hyperboliques**. Hermann, Paris, 1932.
- [Casimir1948] H. B. G. Casimir. On the Attraction Between Two Perfectly Conducting Plates. **Proc. Kon. Ned. Akad. Wet.**, 51:793–795, 1948.
- [Elizalde1995] E. Elizalde. **Ten Physical Applications of Spectral Zeta Functions**. Lecture Notes in Physics, Vol. 855, Springer, 1995.
- [Estrada2002] R. Estrada and R. P. Kanwal. **A Distributional Approach to Asymptotics**. 2nd ed., Birkhäuser, 2002.
- [JaffeWitten2000] A. Jaffe and E. Witten. Quantum Yang–Mills Theory. **Clay Mathematics Institute Millennium Prize Problems**, 2000.
- [SitnikJebabli2024] S. M. Sitnik and I. Jebabli. The Integral Transforms Composition Method for Generalized Index Shifted Transmutations. **Lobachevskii Journal of Mathematics**, 45:3401–3415, 2024.
- [SitnikShishkina2020] S. M. Sitnik and E. L. Shishkina. **Transmutations, Singular and Fractional Differential Equations with Applications to Mathematical Physics**. Mathematics in Science and Engineering, Elsevier, 2020.
- [Sitnik2010] S. M. Sitnik. Transmutations and Applications: A Survey. **arXiv:1012.3741**, 2010.
- [FitouhiEtAl2018] A. Fitouhi, I. Jebabli, E. L. Shishkina, and S. M. Sitnik. Applications of Integral Transforms Composition Method to Wave-Type Singular Differential Equations and Index Shift Transmutations. **Electronic Journal of Differential Equations**, 2018(130):1–27, 2018. arXiv:1805.06925.
- [EpsteinGlaser1973] H. Epstein and V. Glaser. The Role of Locality in Perturbation Theory. **Annales de l’Institut Henri Poincaré A**, 19(3):211–295, 1973.
- [ConnesKreimer2000] A. Connes and D. Kreimer. Renormalization in Quantum Field Theory and the Riemann–Hilbert Problem I: The Hopf Algebra Structure of Graphs and the Main Theorem. **Communications in Mathematical Physics**, 210:249–273, 2000.
- [Scharf1995] G. Scharf. **Finite Quantum Electrodynamics: The Causal Approach**. 2nd ed., Springer, 1995.
- [Schwinger1951] J. Schwinger. On Gauge Invariance and Vacuum Polarization. **Physical Review**, 82:664–679, 1951.
- [DowkerCritchley1976] J. S. Dowker and R. Critchley. Effective Lagrangian and Energy-Momentum Tensor in de Sitter Space. **Physical Review D**, 13:3224–3232, 1976.
- [RaySinger1971] D. B. Ray and I. M. Singer. R-Torsion and the Laplacian on Riemannian Manifolds. **Advances in Mathematics**, 7:145–210, 1971.
- [Seeley1967] R. T. Seeley. Complex Powers of an Elliptic Operator. **Proceedings of Symposia in Pure Mathematics**, 10:288–307, 1967.
- [McKeanSinger1967] H. P. McKean and I. M. Singer. Curvature and the Eigenvalues of the Laplacian. **Journal of Differential Geometry**, 1:43–69, 1967.
- [MinakshisundaramPleijel1949] S. Minakshisundaram and Å. Pleijel. Some Properties of the Eigenfunctions of the Laplace Operator on Riemannian Manifolds. **Canadian Journal of Mathematics**, 1:242–256, 1949.
- [Voros1987] A. Voros. Spectral Functions, Special Functions and the Selberg Zeta Function. **Communications in Mathematical Physics**, 110:439–465, 1987.
- [Polchinski1998] J. Polchinski. **String Theory**, Vols. 1–2. Cambridge University Press, 1998.
- [DerezinskiKarimi2025] J. Dereziński and P. Karimi. Exactly solvable Schrödinger operators related to the hypergeometric equation. **arXiv:2509.03235**, 2025.

Appendix

A Formal Axiom Summary

The following axioms are defined in `theories/pot_renormalization_kernel.kleis` and verified by Z3.

A.1 Regularization Kernel

- **reg_consistency**: $\forall R, X. (\text{admissible}(R) \wedge \text{convergent}(X)) \Rightarrow \text{FP}(K_R[X]) = \text{ord}(X)$
- **reg_maps_zero**: $\forall R. \text{admissible}(R) \Rightarrow K_R[0] = 0$

A.2 Finite Part Operator

- **fp_linearity_add**: $\forall f, g. \text{FP}(f + g) = \text{FP}(f) + \text{FP}(g)$
- **fp_linearity_smul**: $\forall c, f. \text{FP}(cf) = c \cdot \text{FP}(f)$
- **fp_zero_maps_zero**: $\text{FP}(0) = 0$
- **fp_decomposition**: $\forall f. \exists s, c. \text{FP}(f) = c$

A.3 Gauge Group of Regulators

- **gauge_invariance**: $\forall R, R', X. (R \sim R') \Rightarrow \text{FP}(K_R[X]) = \text{FP}(K_{R'}[X])$
- **gauge_reflexive**: $\forall R. R \sim R$
- **gauge_symmetric**: $\forall R, R'. (R \sim R') \Rightarrow (R' \sim R)$
- **gauge_transitive**: $\forall R, R', R''. (R \sim R' \wedge R' \sim R'') \Rightarrow (R \sim R'')$
- **identity_reg_admissible**: $\text{admissible}(\text{id})$

A.4 Spectral Zeta

- **zeta_is_admissible**: $\text{admissible}(K_\zeta)$
- **zeta_sum_of_integers**: $\text{FP}_\zeta(\mathbb{Z}^+) = -1/12$
- **zeta_product_via_derivative**: $-\zeta'(0) \rightarrow \sqrt{2\pi}$
- **analytic_continuation_preserves_fp**: $\forall S. \text{FP}(\zeta_S(s_0)) = \text{FP}_\zeta(S)$

A.5 Heat Kernel

- **heat_is_admissible**: $\text{admissible}(K_{\text{heat}})$
- **heat_zeta_gauge_equiv**: $K_{\text{heat}} \sim K_\zeta$
- **mellin_bridge**: $\forall S. \text{FP}_\zeta(S) = \text{FP}(\text{Tr}(e^{-tA})|_{t_0})$

A.6 Physical Kernel Identification

- **heat_governs_diffusion**: K_{heat} governs heat diffusion (physical kernel)
- **physical_kernel_def**: $\forall R, P. \text{governs}(R, P) \Rightarrow \text{physical}(R)$
- **physical_implies_admissible**: $\forall R. \text{physical}(R) \Rightarrow \text{admissible}(R)$
- **mathematical_kernel_inherits_admissibility**: $\forall R, R_{\text{phys}}. (\text{physical}(R_{\text{phys}}) \wedge R \sim R_{\text{phys}}) \Rightarrow \text{admissible}(R)$
(zeta kernel is admissible via gauge equivalence to heat kernel)

A.7 QFT Projection

- **qft_observable_def**: $\forall R, Z. \text{Obs}_R(Z) = \text{FP}(K_R[Z])$
- **qft_gauge_invariance**: $\forall R, R', Z. (R \sim R') \Rightarrow \text{Obs}_R(Z) = \text{Obs}_{R'}(Z)$
- **path_integral_needs_kernel**: $\forall Z. \exists R. \text{admissible}(R)$

A.8 Composite QFT Kernel

- **composite_def**: $\forall R, J. K_{\text{QFT}}(R, J) = \text{FP}(K_{\text{composite}}(R, J))$
- **composite_gauge_invariance**: $\forall R, R', J. (R \sim R') \Rightarrow K_{\text{QFT}}(R, J) = K_{\text{QFT}}(R', J)$

A.9 ITCM Transmutation

- **itcm_is_composite**: $\forall J. T_{\text{ITCM}}(w_{\text{dim}}, J) = K_{\text{QFT}}(\zeta, J)$
- **itcm_intertwining**: $\text{shift}(w_{\text{dim}}, \nu_{\text{free}}) = \mu_{\text{int}}$
- **coupling_gauge_invariance**: $\forall E, R, R'. (R \sim R') \Rightarrow \alpha(E) = \alpha(E)$
- **alpha_is_projected_observable**: $\alpha(E_{\text{low}}) = \alpha(E_{\text{low}})$

A.10 QED Hypergeometric Kernel

- **exchange_symmetry_a**: $a_1 = a_2$ (symmetric in μ, ν)
- **exchange_symmetry_b_sum**: $b_1 = b_2$ (spectral duality)
- **kernel_decomposition**: $\forall k_1, k_2. K_{\text{inner}}(k_1, k_2) = {}_2F_1(a_1, b_1; c_1; K_{\text{inner}}(k_1, k_2))$
- **hypergeom_gauge_invariance**: $\forall k_1, k_2, R, R'. (R \sim R') \Rightarrow \text{FP}(K_{\text{inner}}(k_1, k_2)) = \text{FP}(K_{\text{inner}}(k_1, k_2))$

A.11 Spectral Localization

- **qed_gapless**: $\text{spec_gap}(T_0^{\text{QED}}) = 0$ (massless photon)
- **qed_gap_from_regular**: $\text{spec_gap}(T_0^{\text{QED}}) = \text{spec_gap}(R^{\text{QED}}|_{s=0})$
- **ym_gap_from_regular**: $\text{spec_gap}(T_0^{\text{YM}}) = \text{spec_gap}(R^{\text{YM}}|_{s=0})$
- **qed_regular_gapless**: $\text{spec_gap}(R^{\text{QED}}|_{s=0}) = 0$

A.12 Divergence-Free Composite

- **direct_equals_factored**: $\forall J. K_{\text{direct}}(J) = K_{\text{QFT}}(\zeta, J)$
- **direct_equals_itcm**: $\forall J. K_{\text{direct}}(J) = T_{\text{ITCM}}(w_{\text{dim}}, J)$
- **direct_gauge_trivial**: $\forall J, R, R'. (R \sim R') \Rightarrow K_{\text{direct}}(J) = K_{\text{direct}}(J)$

A.13 Yang–Mills Composite Kernel

- **asymptotic_freedom**: $\text{AF}(w_{\text{YM}})$
- **ir_regularity**: $\text{IR_reg}(w_{\text{YM}})$
- **gauge_invariance**: $\text{GI}(w_{\text{YM}}, \text{SU}(3))$
- **unitarity**: $\text{unitary}(w_{\text{YM}})$
- **mass_gap**: $\text{gap}(w_{\text{YM}})$
- **gap_is_spectral**: $\text{spec_gap}(w_{\text{YM}}) = \Delta$
- **ym_gauge_trivial**: $\forall J, R, R'. (R \sim R') \Rightarrow K_{\text{YM}}(w, J) = K_{\text{YM}}(w, J)$

A.14 Weight Moduli Space

- **weight_equiv_def**: $\forall w_1, w_2, J. (w_1 \approx w_2) \Rightarrow \text{Obs}(w_1, J) = \text{Obs}(w_2, J)$
- **weight_equiv_reflexive**: $\forall w. w \approx w$
- **weight_equiv_symmetric**: $\forall w_1, w_2. (w_1 \approx w_2) \Rightarrow (w_2 \approx w_1)$

- **weight_equiv_transitive:** $\forall w_1, w_2, w_3. (w_1 \approx w_2 \wedge w_2 \approx w_3) \Rightarrow (w_1 \approx w_3)$
- **weight_compose_associative:** $\forall w_1, w_2, w_3. (w_1 \cdot w_2) \cdot w_3 = w_1 \cdot (w_2 \cdot w_3)$
- **weight_compose_identity:** $\forall w. w \cdot w_{\text{triv}} = w$
- **rg_flow_equiv:** $\forall w, \mu, \mu'. w_\mu \approx w_{\mu'}$ (RG flow preserves equivalence class)
- **rg_fixed_point:** RG-invariant w has scale-independent observables
- **trivial_is_fixed_point:** $\text{RG}(w_{\text{triv}}, \mu) = w_{\text{triv}}$ (free field theory)