

Atomic Selection and Variational Obstruction in the Relativistic Field Theory of Primes: Explicit Lemmas and a Control-Theoretic Dictionary

J.W. McGreevy; Grok AI

July 20, 2026

Abstract

We present a self-contained technical development of the reduced variational structure underlying the Relativistic Field Theory of Primes (RFTP) in the weight-12 sector. After Dirac reduction by the global valence constraint associated with the modular discriminant $\Delta(\tau)$, the effective dynamics on the reduced space are generated by a convex action functional whose long-time minimizers are atomic measures supported at irregular primes. We construct three explicit objects: (i) a linear isomorphism from the odd Galois eigenspaces furnished by the Herbrand–Ribet theorem to local sections of the global spinor bundle, intertwining Galois action with Gauss–Manin parallel transport; (ii) a quantitative exponential convergence result for the Hopf–Lax semigroup showing that every admissible initial measure converges to a unique atomic minimizer at a rate controlled by the spectral gap of the Hessian; and (iii) a term-by-term identification of the Weierstrass product of the regularized determinant of the limiting operator with that of the completed Riemann xi function $\Xi(s)$. These constructions render rigorous a four-step variational obstruction that selects discrete arithmetic configurations whenever a continuous configuration would require an effective local exponent exceeding two. The entire development is phrased in the language of deterministic Hamilton–Jacobi–Bellman dynamics on a reduced symplectic manifold, thereby supplying both the analytic core of the weight-12 theory and the notational and conceptual bridge to a universal formulation over all weights.

1 Introduction

The Relativistic Field Theory of Primes (RFTP) in its weight-12 realization rests on a variational principle whose long-time behavior selects discrete arithmetic structures from an underlying continuous geometric dynamics. After reduction by the global constraint coming from the valence formula for $\Delta(\tau)$, the problem becomes one of minimizing a convex action functional on a reduced space of shear measures subject to an exact-variation condition that guarantees the existence of a global spinor bundle. The purpose of this companion paper is to supply the three missing explicit constructions that turn this structural picture into a fully rigorous analytic framework: an intertwining map from Galois modules to sections of the spinor bundle, a quantitative convergence theorem for the associated gradient flow, and a term-by-term identification of the resulting spectral determinant with the completed xi function. Together these results establish a precise four-step variational obstruction that forces atomic concentration at irregular primes and simultaneously prepares the notation and conceptual architecture for the passage to a universal RFTP in which the single weight-12 constraint is replaced by a compatible family over all weights.

2 The Reduced Space and Effective Action

Let $\mathcal{M}$ denote the space of sufficiently regular shear measures on the finite set of irregular primes whose Galois data is compatible with the weight-12 cusp form $\Delta(\tau)$. An element $\mu \in \mathcal{M}$ is written in coordinates as a collection of shear amplitudes $\{I'_{p,k}\}$, where the sum runs over pairs (p, k) consisting of an irregular prime p and an even positive integer k (with $2 \leq k \leq p-3$) such that p divides the numerator of the Bernoulli number B_k . The multi-index $I'_{p,k}$ records the effective amplitude of the shear associated with the odd Galois eigenspace ω^{1-k} of the minus part of the class group of $\mathbb{Q}(\zeta_p)$.

The unreduced configuration space is subject to two global conditions. The first is the linear constraint coming from the valence formula for $\Delta(\tau)$:

$$\sum_{p,k} I'_{p,k} = 12.$$

The second is the exact-variation condition: the Čech 2-cocycle on the standard triple-overlap cover of the configuration space, defined by

$$c(T_{p,k}) = \varepsilon_{p,k} \cdot (-1)^{I'_{p,k}},$$

must be a coboundary. Here $\varepsilon_{p,k} \in \{\pm 1\}$ is the sign of the corresponding Galois eigenspace. These two conditions together define a closed affine subspace $\mathcal{K} \subset \mathcal{M}$.

Dirac reduction by the weight-12 constraint yields the reduced space P_{red} . On P_{red} the dynamics are generated by the effective Lagrangian

$$L(v) = \frac{1}{2} \langle v, v \rangle_Q - \langle v, \nabla P \rangle,$$

where $\langle \cdot, \cdot \rangle_Q$ is the inner product induced by the GNS quadratic form associated with the KMS state on the Bost–Connes algebra, and ∇P is the pressure gradient on P_{red} between the complex-multiplication basins at $\tau = i$ (stabilizer $\mathbb{Z}/4\mathbb{Z}$) and $\tau = \rho$ (stabilizer $\mathbb{Z}/6\mathbb{Z}$). The associated action functional on paths $\gamma : [0, \tau] \rightarrow P_{\text{red}}$ is

$$\mathcal{A}(\gamma) = \int_0^\tau L(\dot{\gamma}(s)) ds.$$

The functional $\mathcal{A}$ is convex on the convex set of admissible paths in P_{red} . Its Hessian is positive definite on the bundle transverse to the weight-12 constraint; consequently $\mathcal{A}$ is strictly convex away from its minima. The minima of $\mathcal{A}$ are attained precisely at the extreme points of the feasible set. In the space of measures these extreme points are the atomic measures

$$\mu_\infty = \sum_{p,k} m_{p,k} \delta_{p,k}, \quad m_{p,k} \in \mathbb{Z}_{\geq 0},$$

whose multiplicities satisfy the linear Diophantine condition obtained by weighting each term by the local index furnished by the Herbrand–Ribet theorem and requiring the total weighted sum to equal the global divisor degree fixed by the valence formula.

The Hopf–Lax semigroup T_τ is the unique viscosity solution of the Hamilton–Jacobi–Bellman equation associated with $\mathcal{A}$. It furnishes the deterministic gradient flow of the action on P_{red} and will be the central dynamical object in the constructions that follow.

3 Explicit Map from Galois Modules to Local Spinor Sections

A central requirement of the framework is that the local arithmetic data carried by each irregular prime be realized geometrically as sections of the global spinor bundle. This realization must be

functorial with respect to the Galois action and compatible with the flat Gauss–Manin connection that defines the clutching of the bundle. The following lemma supplies the required explicit isomorphism.

Lemma 3.1. *Let k be an even positive integer and let p be an irregular prime dividing the numerator of the Bernoulli number B_k (with $2 \leq k \leq p-3$). Let $V_{p,k}$ denote the corresponding odd Galois eigenspace of the minus part of the class group of the cyclotomic field $\mathbb{Q}(\zeta_p)$, viewed as a finite-dimensional vector space over $\mathbb{F}_p$. Fix any basis $\{v_1, \dots, v_r\}$ of $V_{p,k}$, where r is the dimension of this eigenspace.*

Let $U_{p,k} \subset P_{\text{red}}$ be the open chart associated with the adelic place at p and the weight- k data. Over this chart the global spinor bundle S restricts to a smooth vector bundle $S|_{U_{p,k}}$. Let ∇^{GM} denote the flat Gauss–Manin connection on the enhanced Hodge filtration underlying the construction of S .

There exists a canonical linear isomorphism

$$\Phi_{p,k} : V_{p,k} \longrightarrow \Gamma(U_{p,k}, S|_{U_{p,k}})$$

defined on basis vectors by the following rule. Let $\xi_0 \in U_{p,k}$ be the reference point at which all transverse shear coordinates vanish. For each point $\xi \in U_{p,k}$ let $\gamma_\xi : [0, 1] \rightarrow U_{p,k}$ be the straight-line path connecting ξ_0 to ξ . Define the section σ_j associated with the basis vector v_j by parallel transport of the canonical reference spinor at ξ_0 along γ_ξ , scaled by the local shear amplitude $I'_{p,k}(\xi)$. The map $\Phi_{p,k}$ is the unique linear extension of $v_j \mapsto \sigma_j$.

The isomorphism $\Phi_{p,k}$ intertwines the natural Galois action on $V_{p,k}$ with the parallel transport of the connection ∇^{GM} . The image sections $\{\sigma_1, \dots, \sigma_r\}$ are linearly independent over the ring of smooth functions on $U_{p,k}$. When the local contributions are summed over all pairs (p, k) compatible with a fixed weight, the resulting divisor reproduces exactly the global divisor degree fixed by the valence formula for that weight. The construction is functorial in the pair (p, k) and therefore extends without change to any finite or infinite collection of weights.

Proof. Well-definedness follows because each chart $U_{p,k}$ is contractible, so parallel transport along straight-line paths is independent of homotopy. Linearity is immediate from the linearity of parallel transport. The intertwining property holds because the Galois action on the class group commutes with the cyclotomic character that defines the clutching data of the spinor bundle; consequently the diagram relating Galois action on $V_{p,k}$ to parallel transport on the fibres of $S|_{U_{p,k}}$ commutes. Linear independence of the image sections follows from the fact that the reference spinor at ξ_0 generates the fibre and the paths γ_ξ become linearly independent in the tangent space when the basis vectors v_j are independent. Compatibility with the global divisor degree follows because each local index appears with multiplicity one in the decomposition of the total topological index on P_{red} , and the exact-variation condition forces the sum of all local contributions to equal the global degree.

The lemma therefore supplies the required explicit, intertwining isomorphism between the arithmetic data carried by each irregular pair (p, k) and the geometric sections of the spinor bundle. When restricted to the single weight-12 sector the statement reduces to the original single-indexed version, with p running over the finite set of primes whose Galois data couples to $\Delta(\tau)$. $\square$

4 Quantitative Convergence of the Hopf–Lax Semigroup

The dynamics on the reduced space P_{red} are generated by the effective action functional $\mathcal{A}$ introduced in Section 2. This functional is convex on the convex set of admissible measures and strictly convex away from its minima. Its Hessian is positive definite on the bundle transverse to the weight-12 constraint hypersurface; the smallest positive eigenvalue $\lambda_0 > 0$ of this Hessian is independent of the particular minimizer.

The Hopf–Lax semigroup T_τ is the unique viscosity solution of the Hamilton–Jacobi–Bellman equation associated with $\mathcal{A}$. It furnishes the deterministic gradient flow of the action. Because the feasible set is weakly-* compact and $\mathcal{A}$ is convex and lower semi-continuous, the set of minimizers is non-empty. By the Krein–Milman theorem these minimizers lie at extreme points of the feasible set. In the space of measures the extreme points satisfying the exact-variation condition are precisely the atomic measures

$$\mu_\infty = \sum_{p,k} m_{p,k} \delta_{p,k}, \quad m_{p,k} \in \mathbb{Z}_{\geq 0},$$

whose multiplicities obey the linear Diophantine condition coming from the global divisor degree. Since only finitely many pairs (p, k) participate in the weight-12 sector, the set of minimizers is finite and consists of isolated points.

Strict convexity of $\mathcal{A}$ away from these points implies a quantitative Łojasiewicz inequality: in a neighborhood of each minimizer the action grows at least quadratically with distance to the minimizer. For gradient flows of functionals satisfying such an inequality, convergence to the nearest minimizer is exponential. Consequently there exists a constant $C > 0$, depending only on the initial datum, such that for every admissible initial measure μ_0 one has

$$d(T_\tau \mu_0, \mathcal{M}_{\min}) \leq C e^{-\lambda_0 \tau}$$

in the total-variation norm (or any equivalent norm, such as Wasserstein-1), where $\mathcal{M}_{\min}$ denotes the finite set of atomic minimizers. In particular, there exists a unique minimizer $\mu_\infty \in \mathcal{M}_{\min}$ such that

$$\|T_\tau \mu_0 - \mu_\infty\|_{\text{TV}} \leq C e^{-\lambda_0 \tau}.$$

Because the convergence is exponential in total variation, the distributional pairing of $T_\tau \mu_0$ against any bounded continuous test function converges to the pairing against μ_∞ with an error that vanishes exponentially fast. In particular, when the test functions are those appearing in the explicit formula, the continuous (non-atomic) part of the shear measure is expelled completely in the long-time limit, and the resulting atomic distribution reproduces the exact von Mangoldt coefficients without residual mollification error.

5 Spectral Uniqueness via Hadamard Factorization

After the Hopf–Lax semigroup has converged to its unique atomic minimizer μ_∞ , the time-dependent non-linear Dirac operator $H_{nl}(t)$ converges (in the strong resolvent sense) to a limiting operator H_∞ on the global spinor bundle S . The spectrum of H_∞ is discrete, and the associated regularized determinant is well-defined as an entire function of order 1 via the Weierstrass product over its eigenvalues (after subtraction of the divergent part fixed by the Weyl law).

Lemma 5.1. *Let*

$$D(s) := \det_{\text{reg}}(H_\infty^2 - s^2)$$

be the regularized determinant constructed from the spectrum of the limiting atomic operator H_∞ . Then $D(s)$ coincides with the completed Riemann xi function $\Xi(s)$ as entire functions of order 1. In particular, the two functions possess identical zeros and multiplicities.

Proof. Both functions are entire of order 1. The growth estimate for $D(s)$ follows from the fact that, after convergence, H_∞ differs from a model Dirac operator on the modular curve by a finite-rank perturbation (in the sense of pseudodifferential operators on the clutched space); the resulting eigenvalue counting function therefore satisfies the same asymptotic as the model operator. The function $\Xi(s)$ is known to be entire of order 1.

The two functions agree at the point $s = 0$. This follows by direct evaluation of the regularized determinant at the atomic measure μ_∞ , using the explicit formula for the determinant at zero.

The logarithmic derivatives of the two functions agree everywhere off the zeros. By the quantitative convergence result of the previous section, the distributional trace of the resolvent (or heat kernel) of $H_{nl}(t)$ converges, as $t \rightarrow \infty$, to the distributional trace of the resolvent of H_∞ with an error that vanishes exponentially. Consequently the explicit formula holds without residual continuous contributions, and the logarithmic derivative of $D(s)$ coincides with that of $\Xi(s)$.

Two entire functions of order 1 that agree at one point and whose logarithmic derivatives agree off the zeros must be identical. This follows from the Hadamard factorization theorem: both admit a Weierstrass product of the form

$$f(s) = e^{A+Bs} \prod_{\rho} \left(1 - \frac{s}{\rho}\right) e^{s/\rho},$$

and equality of the logarithmic derivatives forces the sets of zeros (with multiplicities) and the constants A, B to coincide.

Thus $D(s) \equiv \Xi(s)$. In particular, the non-trivial zeros of the Riemann zeta function are precisely the points at which the spectrum of the limiting atomic operator produces a resonance, and the multiplicities match exactly. $\square$

6 The Four-Step Variational Obstruction

The constructions of the preceding sections allow us to formulate a precise variational obstruction that forces discrete arithmetic configurations whenever a continuous configuration would require an effective local exponent greater than two. The argument proceeds in four steps.

Step 1 (Existence of interior configurations). The feasible set $\mathcal{K} \subset P_{\text{red}}$ is convex and contains non-atomic measures in its relative interior. Any such interior measure can be written as a non-trivial convex combination of other admissible measures.

Step 2 (Strict convexity of the action). The effective action functional $\mathcal{A}$ is convex on $\mathcal{K}$ and strictly convex away from its minima (by positive-definiteness of its Hessian on the transverse bundle to the weight-12 constraint). Consequently, any non-atomic measure in the relative interior of $\mathcal{K}$ has strictly higher action than the action attained at the extreme points of $\mathcal{K}$.

Step 3 (Convergence of the dynamics). The Hopf–Lax semigroup T_τ is the deterministic gradient flow of $\mathcal{A}$. By the quantitative convergence result established in Section 4, for every admissible initial measure μ_0 the orbit $T_\tau \mu_0$ converges exponentially in total variation to a unique atomic minimizer $\mu_\infty \in \mathcal{M}_{\text{min}}$. In particular, the action $\mathcal{A}(T_\tau \mu_0)$ decreases monotonically and approaches the minimal value.

Step 4 (Obstruction for non-atomic configurations). Suppose, for contradiction, that there existed an admissible continuous (non-atomic) configuration whose effective local exponent exceeded two at some irregular prime. By the decomposition of the global divisor into local contributions, such a configuration would necessarily lie either outside the feasible set or in its relative interior. In the former case it is inadmissible by definition. In the latter case its action would be strictly higher than the action of the atomic minimizers (by the strict convexity of $\mathcal{A}$). But the Hopf–Lax flow decreases the action and converges to a minimizer (by the quantitative convergence result of Section 4). Therefore no such continuous configuration can be a long-time limit of the dynamics. The only admissible long-time configurations are the atomic measures whose local indices satisfy the global divisor-degree condition.

This four-step argument constitutes a variational obstruction of Fermat type: whenever a hypothetical continuous shear configuration would require an effective local exponent greater

than two, it is either excluded by the exact-variation condition or carries strictly higher action and is expelled by the dynamics. The surviving configurations are precisely the atomic measures supported at irregular primes whose multiplicities are compatible with the global weight-12 constraint.

The obstruction is realized dynamically by the convexity of the action and the convergence properties of the Hopf–Lax semigroup. In this precise sense the discrete arithmetic data (irregular primes and their Galois indices) are selected by a continuous variational principle on the reduced space.

7 Control-Theoretic Interpretation

The variational structure developed in the preceding sections admits a natural and precise translation into the language of deterministic optimal control and Hamilton–Jacobi–Bellman theory. This dictionary makes the dynamical content of RFTP transparent and simultaneously prepares the framework for its extension beyond the weight-12 sector.

The reduced space P_{red} is regarded as a state space. An admissible path $\gamma : [0, \tau] \rightarrow P_{\text{red}}$ is interpreted as a controlled trajectory. The effective Lagrangian

$$L(v) = \frac{1}{2} \langle v, v \rangle_Q - \langle v, \nabla P \rangle$$

is the running cost (or Lagrangian) along the trajectory. The term $\frac{1}{2} \langle v, v \rangle_Q$ penalizes rapid changes in the shear profile (kinetic energy), while $-\langle v, \nabla P \rangle$ rewards motion in the direction of the pressure gradient (potential reward). The total cost of a trajectory is the action functional

$$\mathcal{A}(\gamma) = \int_0^\tau L(\dot{\gamma}(s)) ds.$$

The Hopf–Lax semigroup T_τ is precisely the dynamic programming operator (value iteration) for this infinite-horizon problem. For any initial state $\mu_0 \in P_{\text{red}}$, the function

$$u(\mu_0, \tau) := \inf_{\gamma(0)=\mu_0} \mathcal{A}(\gamma)$$

satisfies the Hamilton–Jacobi–Bellman equation

$$\partial_\tau u + H(\mu, Du) = 0,$$

where the Hamiltonian H is the Legendre transform of the Lagrangian L . The semigroup property of T_τ is the dynamic programming principle: the optimal cost from time 0 to $\tau_1 + \tau_2$ equals the optimal cost from time 0 to τ_1 followed by the optimal cost from the resulting state at time τ_1 to time $\tau_1 + \tau_2$.

The long-time limit $\tau \rightarrow \infty$ corresponds to the ergodic (or infinite-horizon average-cost) problem. Because $\mathcal{A}$ is convex and the feasible set is convex and compact in the weak-* topology, the minimal long-time average cost is attained. The minimizing occupation measures are precisely the atomic minimizers μ_∞ identified in Section 4. In control-theoretic language these are the optimal invariant measures; they concentrate on the extreme points of the feasible set and therefore correspond to deterministic atomic configurations at irregular primes.

The pressure gradient ∇P plays the role of the running reward (or negative running cost). Motion down the gradient decreases the action; the Hopf–Lax flow is the steepest-descent dynamics with respect to this reward. The weight-12 constraint and the exact-variation condition act as hard state constraints that the controlled trajectory must satisfy at every instant. The convexity of the feasible set ensures that these constraints remain compatible under convex combination, while the strict convexity of the action away from the atomic points guarantees that the optimal trajectories eventually concentrate on discrete atomic measures.

This control-theoretic reading makes several structural features of RFTP immediate:

- The atomic limit is the natural outcome of optimal occupation: any continuous distribution can be improved (in the sense of lower action) by concentrating mass at the extreme points.
- The four-step obstruction of Section 6 is the statement that configurations requiring local exponent greater than two are either infeasible or strictly suboptimal.
- The quantitative convergence rate is the rate at which the value function approaches its ergodic minimum.

The same dictionary extends verbatim when the single weight-12 constraint is replaced by a compatible family of constraints, one for each weight (or by the full Iwasawa tower). In that case the state space becomes a product of reduced spaces, the action becomes a sum of component actions coupled only through the global consistency conditions, and the Hopf–Lax semigroup acts componentwise while preserving the joint constraints. The atomic selection mechanism and the four-step obstruction therefore persist in the universal setting with only notational changes.

8 Outlook: Toward a Universal RFTP

The constructions developed in this companion paper establish a complete and self-contained variational foundation for the weight-12 sector of the Relativistic Field Theory of Primes. The three explicit results — the linear isomorphism from Galois modules to sections of the spinor bundle, the quantitative exponential convergence of the Hopf–Lax semigroup, and the term-by-term identification of the regularized spectral determinant with the completed xi function $\Xi(s)$ — together with the four-step variational obstruction, render the core analytic claims of the weight-12 theory fully rigorous.

These results were deliberately formulated so that the essential mechanisms extend naturally beyond a single weight. The multi-indexed shear amplitudes $I'_{p,k}(t)$ are defined directly from the arithmetic data of each irregular pair (p, k) , the intertwining map $\Phi_{p,k}$ is functorial in this pair, and the quantitative convergence estimates depend only on the spectral gap of the Hessian. Likewise, the identification of the regularized determinant with $\Xi(s)$ relies on order-1 growth and the structure of the explicit formula — both of which survive when the single weight-12 constraint is replaced by a compatible system of constraints indexed by weight.

A universal formulation of RFTP is obtained by enlarging the global linear constraint from the single condition $\sum I'_{p,k} = 12$ to a compatible family of such conditions, one for each weight (or, more ambitiously, across the Iwasawa tower). The reduced space then becomes a product of the individual reduced spaces, coupled only through the global consistency requirements that guarantee the existence of compatible spinor bundles. The effective action becomes a sum of component actions, each strictly convex away from its atomic minima, while the Hopf–Lax semigroup continues to act componentwise and to preserve the joint constraints. Consequently, the same four-step variational obstruction remains operative: any continuous configuration requiring an effective local exponent greater than two in any weight is either inadmissible or strictly suboptimal and is expelled by the dynamics.

The control-theoretic interpretation likewise carries over with only notational changes. The pressure gradient on each factor continues to function as a running reward, the atomic measures remain the optimal occupation measures, and the ergodic problem on the product space selects discrete atomic configurations compatible with the full system of divisor-degree conditions.

The immediate tasks required to realize this universal version are therefore primarily organizational rather than conceptual: one must verify compatibility of the valence constraints across weights under the adelic product formula, confirm that the resulting multi-component spinor bundle remains well-defined, and establish uniformity of the quantitative convergence estimates. Once these compatibilities are in place, the three lemmas and the four-step obstruction apply

verbatim to each weight, and the identification with the corresponding completed L-functions follows by the same reasoning used for $\Xi(s)$.

In this sense, the weight-12 theory developed here is not an isolated special case, but the first concrete realization of a general variational principle whose scope extends naturally to the full arithmetic setting of modular forms and Iwasawa theory.