The Coercive Projection Law of Gravity: A Universal Variational Principle with Explicit Constants and Cross-Probe Falsifiers

Jonathan Washburn
Recognition Physics
jon@recognitionphysics.org

November 4, 2025

Abstract

We articulate a single, universal principle that governs gravitational inference under finite information: the *coercive projection law*. In this view, nature implements a projection from raw baryonic sources to an *effective* source through a fixed, scale– and time–aware kernel; the gravitational field is the unique minimizer of a classical energy with that effective source. We show that Information-Limited Gravity (ILG) is precisely the gravitational instantiation of this law in the *pressure* formulation, where the kernel $\mathsf{w}(k, a) = 1 + C\,(a/(k\tau_0))^\alpha$ maps observed baryons to the effective pressure p, and the potential Φ solves the classical Poisson equation $\nabla^2\Phi = 4\pi G\, a^2 p$.

Mathematically, we prove a coercivity inequality with explicit constants that (i) guarantees existence/uniqueness and stability of the projected solution, (ii) certifies positivity and monotonicity displays used in galaxies and cosmology, and (iii) yields *falsifiers* that bind probes together: rotation curves, tracer-independent E_G, the low-ℓ ISW sign, and the low-L CMB-lensing amplitude are all simultaneously constrained by the same kernel and the same coercivity constants. Operationally, this converts "fits" into *audited inferences*: each analysis ships with machine-readable *certificates* (energy values, residual norms, positivity checks, convergence diagnostics) that verify compliance with the coercive projection law.

Conceptually, the universality and the specific constants trace to Recognition Geometry: the exponent $\alpha = \frac{1}{2}(1 - \varphi^{-1})$ and prefactor $C = \varphi^{-3/2}$ are fixed by golden-ratio structure; the finite-net and rank-one projection constants match those that appear in independent CPM instantiations (e.g., protein folding), explaining cross-domain alignment. We release a minimal, dependency-light engine that implements the grid (FFT) and disk (Hankel) paths and emits certificates with each result (CPM-Cosmology-Grid-Path).

Keywords: coercive projection; information-limited gravity; variational methods; explicit constants; falsifiability; recognition geometry; rotation curves; linear growth; ISW; CMB lensing

1 Introduction

A single law. This paper advances a unifying claim: *gravitational inference in the real world is governed by a universal coercive projection.* Raw baryonic sources are first mapped through a fixed kernel into an effective source, and the observed field is the unique minimizer of a classical energy with that source. Information-Limited Gravity (ILG) is the concrete gravitational presentation of this law. In the *pressure* formulation, the kernel

$$\mathsf{w}(k,a) \;=\; 1 + C\Big(\frac{a}{k\,\tau_0}\Big)^{\alpha}, \qquad C = \varphi^{-3/2}, \quad \alpha = \tfrac{1}{2}(1-\varphi^{-1}),$$

builds an effective pressure p from baryons, and Φ solves the standard Poisson equation $\nabla^2\Phi = 4\pi G\, a^2 p$. At small scales (large k) the kernel tends to unity (laboratory gravity); at large scales (small k) it yields a mild, monotone enhancement that is the *same* in galaxies and cosmology.

From phenomenology to principle. Prior work established ILG's empirical adequacy for rotation curves under a global-only policy and recast ILG as classical gravity with a pressure source. Here we elevate ILG to a *law* by proving a *coercivity inequality with explicit constants* in a general CPM (Coercive Projection Method) framework. The inequality certifies that (i) the projection is unique and stable, (ii) positivity and monotonicity displays are structurally enforced, and (iii) *one* kernel with *one* set of constants simultaneously constrains galaxies, linear growth, lensing, and ISW. This binds probes together: if the kernel is forced globally, any failure in one regime falsifies the whole structure—no retuning.

Certificates, not just fits. The coercive projection law is operational. Each analysis can—and should—emit *certificates*:

$$\text{energy } \mathcal{E}[\Phi|p], \quad \text{residual } \|\nabla^2\Phi - 4\pi G\, a^2 p\|, \quad \text{positivity/monotonicity checks}, \quad \text{grid/Hankel convergence},$$

which together verify compliance with the law. We provide a minimal engine that implements both the 3D grid (FFT) and axisymmetric disk (Hankel) paths and writes these certificates alongside figures and tables (CPM-Cosmology-Grid-Path).

Why the constants are what they are. The specific constants are not fitted artifacts. The exponent α and prefactor C follow from Recognition Geometry's golden-ratio structure; the finite-net and rank-one projection factors that appear in the coercivity bound match those arising in independent CPM instantiations (e.g., folding as phase recognition). This cross-domain alignment explains the empirical universality: *proof optimization* (CPM) and *physical optimization* (recognition under finite information) discover the same architecture.

Contributions.

1. **Universal coercivity law.** A CPM formulation of ILG with an *explicit-constant* coercivity inequality that guarantees uniqueness, stability, and positivity/monotonicity displays.
2. **Single-kernel universality.** A "no-retuning" statement: the same kernel governs galaxies and cosmology; any per-system retuning falsifies the law.
3. **Cross-probe falsifiers.** Linked predictions for rotation-curve residuals (slope nulls), tracer-independent E_G with a monotone scale trend, negative low-ℓ ISW sign, and mild low-L CMB-lensing enhancement—*all* from the same kernel and constants.
4. **Certificates as data.** A practical audit layer (energy, residuals, positivity, convergence, kernel checks) that ships with results and is machine-verifiable.

5. **Constant structure explained.** Alignment of CPM constants and ILG exponents via Recognition Geometry (golden-ratio rigidity), clarifying why universality holds across domains.

Roadmap. Section 2 formalizes the CPM structure (structured set, projection, energy, defect, nets, aggregation). Section 3 instantiates ILG as the pressure formulation and proves existence/uniqueness. Section 4 records the explicit constants (K_{net}, C_{proj}, C_{eng}) that yield $c = 49/162$. Section 5 derives cross-probe falsifiers (rotation curves, E_G, ISW, lensing). Section 6 specifies the certificate schema. Section 7 aligns CPM constants across domains. Section 8 gives new, sign/slope-level predictions. Section 9 sketches the relativistic program. Appendices provide technical details, algorithms, and reproducibility notes.

2 The Coercive Projection Law (Abstract CPM → Physics)

We state the coercive projection law abstractly and specialize it to gravity. The ingredients are: (i) an admissible *structured set* of potentials, (ii) a *projection* map to the unique energy minimizer, (iii) *energy* and *defect* quantifying distance to structure, (iv) *finite nets* and *dispersion* control, and (v) an *aggregation* principle elevating local positivity to global guarantees with *explicit constants.*

Structured set and projection

Let $\Omega \subset \mathbb{R}^3$ be either (a) a bounded Lipschitz domain with Dirichlet boundary data (isolated systems, with the convention $\Phi \to 0$ at infinity), or (b) a periodic box $\mathbb{T}^3$ with meanzero potentials. Define the admissible class

$$\mathcal{V} \;=\; \begin{cases} H_0^1(\Omega), & \text{Dirichlet (isolated)},\\ \{\Phi \in H^1(\mathbb{T}^3) :\ \textstyle\fint_{\mathbb{T}^3} \Phi\, dx = 0\}, & \text{periodic (meanzero).} \end{cases}$$

The *structured set* for the potential is the admissible set itself, endowed with the classical Dirichlet metric (Section 1). The *projection* Π is the solution operator that maps any effective source p to the unique minimizer $\Phi^* = \Pi(p) \in \mathcal{V}$ solving Poisson's equation

$$\nabla^2 \Phi \;=\; 4\pi G\, a^2\, p, \qquad p \;=\; \big[\mathsf{w}(\nabla, a)\, s\big], \tag{1}$$

with the *same* kernel w used in all probes. Here s is the raw baryonic source: $s = \rho_b$ for galaxies (at $a = 1$), and $s = \bar{\rho}_b(a)\, \delta_b$ for cosmology. The operator $\mathsf{w}(\nabla, a)$ is the isotropic convolution with Fourier symbol

$$\mathsf{w}(k, a) \;=\; 1 \;+\; C \left(\frac{a}{k\, \tau_0}\right)^{\alpha}, \qquad C = \varphi^{-3/2}, \quad \alpha = \tfrac{1}{2}(1 - \varphi^{-1}), \tag{2}$$

so that $\mathsf{w} \to 1$ in the laboratory limit (large k) and increases monotonically at long wavelengths.

Energy and defect

For fixed a and source p, define the classical energy functional

$$\mathcal{E}[\Phi \,|\, p] \;=\; \frac{1}{8\pi G}\int_{\Omega} |\nabla\Phi|^2\,dx \;+\; \int_{\Omega} a^2\,p\,\Phi\,dx, \qquad \Phi \in \mathcal{V}. \tag{3}$$

Standard firstvariation gives the Euler-Lagrange equation (1), so the projection Π returns the unique minimizer Φ^*. We measure *defect* in two equivalent ways:

$$\mathsf{D}_{H^1}(\Phi) \;=\; \int_{\Omega} |\nabla(\Phi - \Phi^*)|^2\,dx, \tag{4}$$

$$\mathsf{D}_{\mathrm{res}}(\Phi) \;=\; \left\|\nabla^2\Phi - 4\pi G\,a^2 p\right\|_{H^{-1}}^2, \tag{5}$$

and record the *energy-defect control* (a consequence of Poincaré and elliptic regularity): there exists $C_{\mathrm{eng}} > 0$ depending only on boundary conditions and domain geometry such that

$$\mathsf{D}_{H^1}(\Phi) \;\le\; C_{\mathrm{eng}}\left(\mathcal{E}[\Phi \,|\, p] - \mathcal{E}[\Phi^* \,|\, p]\right), \qquad \mathcal{E}[\Phi \,|\, p] - \mathcal{E}[\Phi^* \,|\, p] \;\gtrsim\; \mathsf{D}_{\mathrm{res}}(\Phi). \tag{6}$$

In periodic boxes one may take $C_{\mathrm{eng}} = 1$ by construction; under Dirichlet data, C_{eng} is an $O(1)$ constant fixed by the Poincaré constant of Ω.

Finite nets and dispersion control

The CPM template localizes distancetostructure by covering admissible modes with a finite εnet and controlling the orthogonal projection error with an explicit constant. In the present setting:

- *Projection constant.* A rankone/Hermitian estimate yields $C_{\mathrm{proj}} \le 2$ for the fiberwise projection that removes components orthogonal to the structured set.
- *Net constant.* For a unit εnet on spectral shells (FFT) or Hankel bands (disks), one records $K_{\mathrm{net}} = ((1+\varepsilon)/(1-\varepsilon))^2$. An *eighttick aligned* choice $\varepsilon = 1/8$ gives $K_{\mathrm{net}} = (9/7)^2$.
- *Dispersion hygiene.* CIC/TSC assignment windows and a 2/3 spectral cutoff suppress aliasing; in Hankel space, logarithmic sampling and Besselkernel quadrature control leakage. These rules ensure that the discrete projection respects the continuous positivity of w.

Aggregation of local positivity

Positivity and monotonicity of w ($\mathsf{w} \ge 1,\ \partial_k \mathsf{w} < 0,\ \partial_a \mathsf{w} > 0$) imply local window tests cannot manufacture sign flips in the effective source or nonmonotone displays in derived quantities. Let $\{T_W\}$ denote a boundedoverlap family of local tests (e.g., residual norms in radial windows for galaxies, bandpowers for cosmology). A standard CPM aggregation yields the global bound

$$\mathsf{D}_{H^1}(\Phi) \;\le\; M\,K_{\mathrm{net}}\,C_{\mathrm{proj}}\,\sup_{W}\,T_W[\Phi], \tag{7}$$

where M is the window overlap constant (fixed by the analysis design). If the righthand side is below a *critical threshold* determined by (6), the energy gap forces small global defect and, hence, proximity to the structured solution.

Explicit constants and the role of φ and τ_0

The coercivity constant that appears in the inequality

$$\mathcal{E}[\Phi \,|\, p]\mathcal{E}[\Phi^* \,|\, p] \;\geq\; c\,\mathsf{D}_{H^1}(\Phi), \qquad c \;=\; \frac{1}{K_{\rm net}\,C_{\rm proj}\,C_{\rm eng}}, \tag{8}$$

is *explicit.* With the eighttick choice $\varepsilon = 1/8$, $K_{\rm net} = (9/7)^2$, the Hermitian bound $C_{\rm proj} \leq 2$, and periodic energy normalization $C_{\rm eng} = 1$, one finds

$$c \;=\; \frac{1}{(9/7)^2 \cdot 2 \cdot 1} \;=\; \frac{49}{162} \;\approx\; 0.302\,.$$

The *kernel constants* derive from Recognition Geometry: the exponent $\alpha = \frac{1}{2}(1\varphi^1)$ and prefactor $C = \varphi^{3/2}$ fix the longwavelength slope and amplitude of w, while the fundamental tick τ_0 sets the (dimensionless) gate between laboratory and cosmic regimes through the ratio $a/(k\,\tau_0)$. These constants explain why the same projection law governs galaxies, growth, and optics without persystem tuning.

3 ILG as the Gravitational Instantiation

Pressure source. In InformationLimited Gravity (ILG), the effective source is the *pressure* field obtained by filtering the raw baryonic source s through the universal kernel w:

$$p(\boldsymbol{x}, a) \;=\; \big[\mathsf{w}(\nabla, a)\, s\big](\boldsymbol{x}), \qquad \widehat{p}(\boldsymbol{k}, a) \;=\; \mathsf{w}(k, a)\, \widehat{s}(\boldsymbol{k}, a), \tag{9}$$

with

$$\mathsf{w}(k, a) \;=\; 1 \;+\; C\left(\tfrac{a}{k\,\tau_0}\right)^{\alpha}, \qquad C = \varphi^{-3/2}, \quad \alpha = \tfrac{1}{2}(1 - \varphi^{-1}). \tag{10}$$

For galaxies (present epoch, $a = 1$) one takes $s = \rho_b$; for cosmology, $s = \bar{\rho}_b(a)\,\delta_b$ in comoving coordinates. The potential Φ solves the *classical* Poisson equation with this source,

$$\nabla^2\Phi(\boldsymbol{x}, a) \;=\; 4\pi G\, a^2\, p(\boldsymbol{x}, a), \tag{11}$$

under Dirichlet decay at infinity (isolated) or periodic meanzero (cosmology) boundary conditions.

Variational statement; existence and uniqueness. For fixed scale factor a and source p, consider the energy $\mathcal{E}[\Phi \,|\, p]$ in (3). The first variation yields (11), so the admissible minimizer $\Phi^* \in \mathcal{V}$ is the *unique* weak solution. Coercivity of the Dirichlet form and the Poincaré inequality on $\mathcal{V}$ imply *existence and uniqueness* by the Lax-Milgram theorem when $p \in H^1(\Omega)$ (e.g., $p \in L^2$ suffices). In periodic boxes, fixing the zero mode of $\widehat{\Phi}$ yields a unique meanzero solution; in isolated domains with $p \in L^1 \cap L^{6/5}$, the Greenś representation

$$\Phi(\boldsymbol{x}, a) \;=\; -\,G\,a^2 \int_{\mathbb{R}^3} \frac{p(\boldsymbol{y}, a)}{|\boldsymbol{x}\boldsymbol{y}|}\, d^3y \tag{12}$$

solves (11) in the distributional sense and decays as required. In either case, the energy gap controls the H^1 distance to the solution by (6), establishing stability.

Positivity, monotonicity, and the laboratory limit. Because $\mathsf{w}(k,a) \geq 1$ for all $k > 0$ and $a \in (0,1]$, the Fouriermultiplier operator $\mathsf{w}(\nabla, a)$ is *positive* in the operator sense:

$$\langle f,\ \mathsf{w}(\nabla,a)\, f\rangle \;=\; \int |\widehat{f}(\boldsymbol{k})|^2\, \mathsf{w}(k,a)\, \frac{d^3k}{(2\pi)^3} \;\geq\; \int |\widehat{f}(\boldsymbol{k})|^2\, \frac{d^3k}{(2\pi)^3} \;=\; \|f\|_2^2. \tag{13}$$

Moreover,

$$\partial_k\, \mathsf{w}(k,a) \;<\; 0, \qquad \partial_a\, \mathsf{w}(k,a) \;>\; 0, \qquad \lim_{k\to\infty} \mathsf{w}(k,a) \;=\; 1, \tag{14}$$

so increasing scale (smaller k) or later times a monotonically enhance the effective source, while the *laboratory limit* is recovered exactly as $k \to \infty$. These properties propagate to displaylevel quantities: (i) effective surface profiles built from p inherit nonpathological signs; (ii) ratios such as $w(R) = v^2(R)/v^2_{\text{baryon}}(R)$ are monotone in regimes where the baryonic Hankel power is concentrated at low k; and (iii) smallscale predictions reduce to standard gravity because $\mathsf{w} \to 1$.

4 Coercivity with Explicit Constants

We now record the explicit constants that certify stability of the coercive projection and bind all probes under a single kernel. The bound reads

$$\mathcal{E}[\Phi\,|\,p] - \mathcal{E}[\Phi^*\,|\,p] \;\geq\; c\, \mathsf{D}_{H^1}(\Phi), \qquad c \;=\; \frac{1}{K_{\text{net}}\, C_{\text{proj}}\, C_{\text{eng}}}. \tag{15}$$

Projection constant (rankone/Hermitian). The fiberwise projection step admits the Hermitian rankone estimate $\min_{\lambda\geq 0,\ \|v\|=1} \|H - \lambda\, v\otimes v^*\|^2_{\text{HS}} \leq 2\,\|H - \frac{\operatorname{tr} H}{d} I\|^2_{\text{HS}}$, hence one may take $C_{\text{proj}} \leq 2$. This constant is domainagnostic and matches independent CPM instantiations.

Net constant (eighttick nets or 2/3 spectral cutoff). For a unit εnet on spectral shells (FFT) or Hankel bands (disks), the conevsnet bound records $K_{\text{net}} = ((1+\varepsilon)/(1-\varepsilon))^2$. The *eighttick* alignment $\varepsilon = 1/8$ gives $K_{\text{net}} = (9/7)^2$. On periodic grids with a 2/3 spectral cutoff, the effective ε induced by shell spacing/window overlap yields a comparable constant; we retain the eighttick value for analysis invariance.

Energy-control (periodic/Dirichlet classes). With the energy normalization in (3), the linear source term cancels at the minimizer, giving the identity $\mathcal{E}[\Phi\,|\,p] - \mathcal{E}[\Phi^*\,|\,p] = \frac{1}{8\pi G}\int |\nabla(\Phi - \Phi^*)|^2$, so $C_{\text{eng}} = 1$ on periodic and Dirichlet classes (additive constant fixed by meanzero/decay).

Coercivity constant and RS alignment. Combining the three constants yields the *explicit* coercivity bound:

$$\boxed{c \;=\; \frac{1}{K_{\text{net}}\, C_{\text{proj}}\, C_{\text{eng}}} \;=\; \frac{1}{(9/7)^2 \cdot 2 \cdot 1} \;=\; \frac{49}{162} \;\approx\; 0.302} \tag{16}$$

Universality and crossdomain structure. The same eighttick net and Hermitian projection constant appear in other CPM domains (e.g., folding), yielding the *same* numerical c. This is strong evidence that projection geometry—not problemspecific tuning—governs stability. Meanwhile, the kernel's exponent and prefactor ($\alpha = \frac{1}{2}(1-\varphi^{-1})$, $C = \varphi^{-3/2}$) and the gate τ_0

fix the longwavelength behavior of w via Recognition Geometry. Together these explain why a *single* coercivity constant and *single* kernel constrain galaxies, growth, and optics without persystem retuning.

5 Aggregation to Falsifiers Across Probes

The CPM aggregation bound (7) elevates *local* window tests to a *global* defect control with explicit constants. Because the same kernel w and the same coercivity constant c govern all modalities, a single family of falsifiers binds galaxies and cosmology together: if any probe fails under the declared windows and hygiene, the universal law is falsified (*no retuning*).

Galaxies: windows → global defect

Let $\{W\}$ be radial windows on each rotation curve after fairness masks (inclination, inner beam, outer reliability, bar/warp excision). With tests T_W (e.g., windowed residual norms) and bounded overlap M, (7) gives

$$\mathsf{D}_{H^1}(\Phi) \;\le\; M\, K_{\mathrm{net}}\, C_{\mathrm{proj}} \sup_W T_W[\Phi].$$

Two immediate consequences become falsifiers:

- **No Retuning Theorem.** With kernel w and constants fixed globally, acceptable residuals *cannot* require pergalaxy changes to w. Any such retuning implies that a single projection does not minimize a single energy across systems, violating the law.
- **Residualslope nulls.** Across windows spanning $\sim$ 2-4 disc scale lengths, the slope of residuals $\Delta v(R)$ should be unbiased and *uncorrelated* with basic baryonic observables (surface brightness Σ_*, gas fraction f_{gas}, morphology) under the globalonly policy. Statistically significant correlations constitute a failure of dispersion hygiene or positivity and thus falsify the law.

Positivity/monotonicity of w propagates to displaylevel checks:

- **Sign/monotone displays.** The effective surface profile $P(R) = \int p(R,z)\,dz$ must be non-pathological when $\Sigma_b \ge 0$ is smooth; the derived display $w(R) = v^2(R)/v^2_{\mathrm{baryon}}(R)$ should be monotone in regimes where Hankel power concentrates at low k. Persistent sign flips or nonmonotone behavior in clean systems falsify the positivity/monotonicity inheritance from w.

Cosmology: linked predictions from the same kernel

In linear theory, the growth equation with the pressure source is $\ddot{\delta}_b + 2\mathcal{H}\dot{\delta}_b - 4\pi G\, a^2\, \bar{\rho}_b(a)\, \mathsf{w}(k,a)\, \delta_b = 0$, so all cosmological predictions inherit the same w.

- **Tracerindependent E_G factorization.** Define $E_G(a,k) = \dfrac{a\,k^2\,\widehat{\Phi}(a,k)}{H(a)\,f(a,k)\,\widehat{\delta}_b(a,k)}$. Using $k^2\widehat{\Phi} = 4\pi G\, a^2\, \bar{\rho}_b\, \mathsf{w}\, \widehat{\delta}_b$, one gets

$$E_G(a,k) \;=\; \left[\tfrac{4\pi G\, a^3\, \bar{\rho}_b(a)}{H(a)}\right] \frac{\mathsf{w}(k,a)}{f(a,k)}, \tag{17}$$

which is *tracerindependent*. Falsifier: significant tracerdependent splits or a residual scale trend opposite to the monotone w (after controlling for f).

- **Mild, monotone scale dependence in $f(a,k)$.** Since w mildly enhances long wavelengths, the growth rate $f(a,k) = \partial \ln D/\partial \ln a$ acquires a controlled, monotone $k$dependence at late times, reverting to the standard limit at early times/small scales. Falsifier: strong, nonmonotone $k$dependence inconsistent with $\partial_k \mathsf{w} < 0$.
- **ISW sign (low ℓ).** The growth of $\mathsf{w}(k,a)$ with a slows the decay of Φ and can make $\dot{\Phi} < 0$ on the largest scales, predicting a *negative* lowℓ ISW crosscorrelation. Falsifier: a robust, maskstable *positive* lowℓ signal.
- **CMB lensing amplitude (low L).** The lineofsight average of w mildly increases the lensing amplitude at low multipoles L, with a smooth return to GR ($\mathsf{w} \to 1$) at high L. Falsifier: a significant *decrease* toward low L or nonsmooth trends incompatible with the monotone kernel.

All four predictions are locked to the *same* constants and *same* kernel w. A pass/fail in one regime cannot be repaired by retuning another: the coercive projection law binds galaxies, growth, ISW, and lensing as a *single* auditable structure.

6 Certificates as FirstClass Outputs

The coercive projection law converts fits into *audited inferences*. Each plot, table, or quantitative claim should ship with a compact, machinereadable *certificate* that verifies compliance with existence/uniqueness, positivity/monotonicity, and discretization hygiene. This section specifies the minimum fields, schemas, and default thresholds.

What to publish with every figure

For each result (galaxy, growth bandpower, lensing bin), attach:

- **Energy:** $\mathcal{E}[\Phi \,|\, p]$ and gap $\mathcal{E}[\Phi \,|\, p] - \mathcal{E}[\Phi^* \,|\, p]$.
- **Residual norms:** $\|\nabla^2\Phi - 4\pi G\, a^2 p\|_{L^2}$ (or H^1) and windowed residuals T_W.
- **Positivity/monotonicity:** pass/fail and metrics for (i) operator positivity checks (nonnegative quadratic form), (ii) displaylevel sign (e.g., $P(R) \geq 0$ where applicable), (iii) monotone display trends (e.g., slope of $w(R)$).
- **Grid/Hankel convergence:** relative change under resolution doubling/padding (e.g., $\|\Phi_{2N} - \Phi_N\|/\|\Phi_{2N}\|$, $|v_{2M}^2 - v_M^2|/v_{2M}^2$).
- **Kernel checks:** (i) $\mathsf{w} \to 1$ at large k (laboratory limit), (ii) longwavelength slope sign $\partial_k \mathsf{w} < 0$, time monotonicity $\partial_a \mathsf{w} > 0$, (iii) numerical stability (logevaluation for enhancement term), (iv) kernel checksum for provenance.

Schema (JSON)

We recommend a single JSON object per artifact (figure/table). The schema includes:

- `artifact_id`: unique identifier (e.g., `"fig_3_panel_b"`)
- `context`: probe type, dataset, object/system, constants ($\varphi, \alpha, C, \tau_0, G$)

- `energy`: $\mathcal{E}[\Phi|p]$ and gap
- `residuals`: global norms (L^2, H^{-1}) and windowed T_W
- `positivity_monotonicity`: operator quadratic form min, display signs/slopes, pass/fail
- `convergence`: grid/Hankel relative errors
- `kernel`: high$k$ deviation, slope signs, checksum (SHA256)
- `provenance`: commit hash, environment, seeds, timestamp

A minimal example (`fig_3_panel_b.json`):


```
{"artifact_id": "fig_3_panel_b",
 "context": {"probe": "galaxy", "dataset": "SPARC", "object": "NGC3198",
             "constants": {"phi": 1.618034, "alpha": 0.190983,
                           "C": 0.485868, "tau0": 1.0, "G": 4.3e-6}},
 "energy": {"E": 1.234e5, "gap": 2.1e3},
 "residuals": {"norm_L2": 3.5e-3, "windows": [...]},
 "positivity_monotonicity": {"passes": true},
 "convergence": {"hankel": {"rel_error": 0.008}},
 "kernel": {"checksum_sha256": "abc123..."},
 "provenance": {"code_commit": "abd3eba", "timestamp": "2025-11-04T..."}}
```


Default thresholds

Thresholds should be set globally and predeclared (per analysis commit):

- **Convergence:** $\|\Phi_{2N} - \Phi_N\|/\|\Phi_{2N}\| < 1\%$; $|v_{2M}^2 - v_M^2|/v_{2M}^2 < 1\%$ over reported radii.
- **Kernel (hi$k$):** $|\mathsf{w}(k_{\max}) - 1| < 5\%$ on the mesh; slope signs $\partial_k \mathsf{w} < 0$, $\partial_a \mathsf{w} > 0$.
- **Positivity/monotonicity:** operator quadratic form ≥ 1 to numerical tolerance; display sign nonnegative where $\Sigma_b \geq 0$; monotone trend consistent with low$k$ dominance.
- **Windows:** bounded overlap M recorded; $\sup_W T_W$ below the critical threshold implied by (6) and (15).

"Green checkmark" reproducibility and audit

We recommend a short, perartifact box that lists pass/fail for each certificate class. A result is marked with a *green checkmark* only if all items pass under the frozen configuration:

- **Reproducibility:** code commit hash; kernel checksum; constants file $(\varphi, \alpha, C, \tau_0, G)$; environment and version pins; random seeds.
- **Energy/residuals:** gap and norms reported; thresholds met.
- **Positivity/monotonicity:** operator and displays pass stated tests.
- **Convergence:** grid/Hankel tolerances met; padding documented for isolated systems.
- **Windows:** masks and overlap constant M recorded; $\sup_W T_W$ below threshold.

Certificates should be archived alongside figures/tables (e.g., as `.json` sidecars) and referenced in captions. This enables third parties to audit compliance with the coercive projection law without rerunning the full pipeline.

7 CrossDomain Constant Structure

The same CPM constants that control stability in gravity appear in independent recognition problems (e.g., folding as phase recognition). Table 1 aligns the numerical values and their provenance, and highlights how the goldenratio structure fixes the kernel exponent α and prefactor C in gravity.

Quantity	Symbol	Folding (CPM)	Gravity (ILG)	Provenance
Net constant	K_{net}	$\left(\frac{9}{7}\right)^2$	$\left(\frac{9}{7}\right)^2$	Eighttick ($\varepsilon = 1/8$) finite net
Projection constant	C_{proj}	≤ 2	≤ 2	Hermitian rankone bound
Energy control	C_{eng}	1	1	Dirichlet/periodic normalization
Coercivity	c	$\frac{49}{162} \approx 0.302$	$\frac{49}{162} \approx 0.302$	$1/(K_{\text{net}} C_{\text{proj}} C_{\text{eng}})$
Kernel exponent	α	—	$\frac{1}{2}(1-\varphi^{-1})$	Recognition geometry (golden ratio)
Kernel prefactor	C	—	$\varphi^{-3/2}$	Recognition geometry (golden ratio)
Gate parameter	τ_0	—	global (fixed)	Finite refresh; lab $\to$ cosmic gate

Table 1: Alignment of CPM constants across domains. The same net and projection constants yield the same coercivity c. In gravity, the kernel exponent α and prefactor C follow from golden-ratio structure, while τ_0 sets the dimensionless gate between laboratory and cosmic regimes.

Interpretation: why proof and physics discover the same constants. Recognition Science (RS) explains why *proof optimization* (CPM) and *physical optimization* (inference under finite refresh) converge to the same architecture:

- The eighttick alignment ($\varepsilon = 1/8$) that optimizes covering nets in CPM coincides with the timing structure (eightbeat ledger cycles) that optimizes recognition capacity.
- The Hermitian rankone bound ($C_{\text{proj}} \leq 2$) reflects the same minimal projection geometry across convex cones and recognition modes.
- In gravity, goldenratio rigidity fixes the kernel's longwavelength behavior: $\alpha = \frac{1}{2}(1 - \varphi^{-1})$ and $C = \varphi^{-3/2}$ are *not* tunable; they follow from the recognition cost functional and the golden ratio's unique fixedpoint property.

The result: a single set of constants that stably governs galaxies, growth, and optics without persystem dials. The constant alignment ($c = 49/162$ in both folding and gravity) is not numerology—it is RS architecture discovered from both directions.

8 New Predictions (from the Law, not a Model)

The coercive projection law yields *linked, sign and slopelevel predictions* that do not depend on auxiliary modeling choices. They arise from positivity and monotonicity of w, the explicit coercivity constant, and the singlekernel universality that binds probes together.

The rotation–lensing–growth triangle

- **Linked signs and slopes.** A monotone $\mathsf{w}(k, a)$ ($\partial_k \mathsf{w} < 0$, $\partial_a \mathsf{w} > 0$) demands: (i) outer rotationcurve displays $w(R)$ nondecreasing over radii where Hankel power concentrates at low k; (ii) a mild, monotone $k$dependence in the latetime growth rate $f(a, k)$; (iii) a negative lowℓ ISW sign; and (iv) a mild low$L$ enhancement in CMB lensing amplitude with a smooth return to GR at high L.

- **Nogo behaviors.** The following cannot occur under the law: (i) persistent nonmonotone $w(R)$ in clean disks; (ii) strong, oscillatory $k$dependence in $f(a,k)$ after controlling for background; (iii) a robust *positive* lowℓ ISW crosscorrelation; (iv) a decreasing lensing amplitude toward low L; or (v) perprobe retuning of w to reconcile inconsistent signs/slopes. Any one is a structural falsifier.

Nearfield slope and nanogravity trend

At high wavenumber, $\mathsf{w}(k,a) = 1 + C(a/(k\tau_0))^\alpha$ with $\alpha > 0$ implies a *negative* logarithmic slope $d\ln \mathsf{w}/d\ln k = -\alpha \frac{C(a/(k\tau_0))^\alpha}{1+C(a/(k\tau_0))^\alpha} < 0$. Thus, nearfield deviations must be small, negativeslope corrections approaching unity from above as $k \to \infty$. The *nanogravity* regime therefore exhibits a gentle, monotone approach to GR with no oscillatory or positiveslope features. Controlled experiments that recover the *opposite* sign or a nonmonotone behavior would falsify the kernel form.

Crossprobe amplitude − band constraints

The lineofsight average of w imposes consistent amplitude bounds across probes and bands: rotationcurve displays, tracerindependent E_G, lowℓ ISW, and low$L$ lensing form a *single amplitude budget*. A scale/time band that demands enhancement in one probe must produce a commensurate response in the others. Conversely, a band that appears enhanced in one probe but suppressed in another (after identical hygiene) violates singlekernel universality.

9 Outlook: Relativistic Completion via Coercive Projection

The nonrelativistic law invites a relativistic presentation that preserves coercivity and universality.

Effective stress–energy and route identities

Define an effective, divergencecontrolled stress–energy (or pressure 2form) built from the filtered source $p = \mathsf{w}(\nabla, a)\, s$, and choose a gauge in which the coercivity identity (15) holds at the level of metric potentials. *Route identities* (Kgates) then lock normalizations by equating independent constructions (e.g., time − to − length routes), eliminating ambiguity in gauge presentations and preventing hidden degrees of freedom. The same constants $(\varphi, \alpha, C, \tau_0)$ fix the longwavelength sector.

Program for Nbody and multiprobe synthesis

- **Nbody with prefiltered sources.** Incorporate the prefilter step ($p = \mathsf{w} * s$) at each time slice, then solve standard Poisson and advance particles. Emit certificates (energy, residuals, convergence, kernel checks) per snapshot.
- **Multiprobe joins.** Enforce singlekernel universality by sharing kernel arrays and constants across rotationcurve, growth, lensing, and ISW pipelines; attach perprobe certificates and a

crossprobe consistency summary.

- **Release practice.** Archive certificate sidecars (JSON), kernel checksums, and environment pins with each public figure/table to enable independent audit without reruns.

This program extends the present nonrelativistic law to surveyscale inference while preserving its decisive feature: *a single, auditable structure* that ties galaxies, growth, and optics together by coercivity and explicit constants.

10 Conclusion

We have articulated and instantiated a *coercive projection law of gravity*: nature projects raw baryonic sources through a fixed, scale and timeaware kernel to construct an effective pressure source, and the gravitational field is the unique minimizer of a classical energy with that source. In this presentation, InformationLimited Gravity (ILG) is not a tunable phenomenology but the gravitational face of a universal projection principle. A CPM (Coercive Projection Method) framework with *explicit constants* certifies existence/uniqueness, positivity, and stability, and binds galaxies, growth, ISW, and lensing as a *single* auditable structure.

Three distinguishing features.

1. **Universality.** One kernel and one set of coercivity constants apply across probes; persystem retuning falsifies the law.
2. **Falsifiability.** Positivity and monotonicity of the kernel generate linked, sign and slopelevel predictions (rotation − growth − lensing triangle; nearfield slope sign; crossprobe amplitude bands) that cannot be violated without breaking the principle.
3. **Auditability.** Every result ships with compact *certificates*—energy values, residual norms, positivity/monotonicity checks, convergence diagnostics, kernel sanity and checksums—so that independent parties can verify compliance without reruns.

Why the constants align. The coercivity constants (net, projection, energy) match those arising in independent CPM domains (e.g., folding), while the kernel's exponent and prefactor follow from Recognition Geometry's goldenratio rigidity; the fundamental tick τ_0 sets the gate between laboratory and cosmic regimes. This alignment explains why *proof optimization* (CPM) and *physical optimization* (recognition under finite refresh) converge: the same architecture is discovered from both directions.

The engine and next steps. We release a minimal engine (CPMCosmologyGridPath) that implements the grid (FFT) and disk (Hankel) paths, emits certificates by default, and enables multiprobe synthesis under a single kernel. Nearterm directions: (i) relativistic completion preserving coercivity via effective stress–energy and Kgates; (ii) Nbody pipelines with prefiltered sources and persnapshot certificates; (iii) surveyscale audits reporting pass/fail against predeclared thresholds. The decisive test is not a better fit but a *better law*: one kernel, explicit constants, green checkmarks, and crossprobe predictions that stand or fall together.

Appendix A: Functional Setting and Existence Details

Function spaces and boundary conditions. For an isolated domain $\Omega \subset \mathbb{R}^3$ with $\Phi \to 0$ at infinity, take $\mathcal{V} = H_0^1(\Omega)$. For periodic $\mathbb{T}^3$, take $\mathcal{V} = \{\Phi \in H^1(\mathbb{T}^3) : \int \Phi = 0\}$. Assume $p \in H^{-1}(\Omega)$ ($p \in L^2$ suffices). The bilinear form $a(\Phi, \Psi) = (8\pi G)^{-1} \int \nabla\Phi \cdot \nabla\Psi$ is continuous and coercive on $\mathcal{V}$.

Lax-Milgram and energy identity. By Lax-Milgram, for each $p \in H^{-1}$ there exists a unique $\Phi^* \in \mathcal{V}$ solving $\nabla^2\Phi = 4\pi G\, a^2 p$ with the stated boundary conditions. Moreover,

$$\mathcal{E}[\Phi \,|\, p] - \mathcal{E}[\Phi^* \,|\, p] \;=\; \frac{1}{8\pi G} \int_\Omega |\nabla(\Phi - \Phi^*)|^2 \, dx,$$

giving $C_{\rm eng} = 1$ in the energy-defect control.

Green's representation (isolated). If $p \in L^1(\mathbb{R}^3) \cap L^{6/5}(\mathbb{R}^3)$, the potential $\Phi(\boldsymbol{x}, a) = -\,G\, a^2 \int p(\boldsymbol{y}, a)/|\boldsymbol{x} - \boldsymbol{y}|\, d^3y$ solves the Poisson equation in the distributional sense and decays at infinity; the Dirichlet energy is finite.

Operator positivity. For real $f \in L^2$, $\langle f, \mathsf{w}(\nabla, a) f\rangle = \int \mathsf{w}(k, a) |\widehat{f}|^2\, d^3k/(2\pi)^3 \geq \|f\|_2^2$ since $\mathsf{w} \geq 1$. This is an operatorlevel statement and does not require a pointwise $W(r, a) \geq 0$.

Appendix B: Discretization Hygiene and Algorithms

FFT grid path (periodic).

- Deposit s on an $N_x \times N_y \times N_z$ mesh; FFT to $\widehat{s}$.
- Multiply by $\mathsf{w}(k, a)$ in Fourier space (evaluate enhancement in logs).
- Solve $\widehat{\Phi} = -4\pi G\, a^2 \widehat{p}/k^2$ for $\boldsymbol{k} \neq 0$; set $\widehat{\Phi}(0) = 0$.
- Inverse FFT; differentiate spectrally for forces; apply a 2/3 spectral cutoff.

Hankel disk path (axisymmetric).

- Compute $\tilde{\Sigma}_b(k) = \int R\, \Sigma_b(R)\, J_0(kR)\, dR$ on a log grid (FFTLog).
- Form $\tilde{P}(k) = \mathsf{w}(k, 1)\, \tilde{\Sigma}_b(k)$ and $v^2(R) = 2\pi G\, R \int k\, J_1(kR)\, \tilde{P}(k)\, dk$.
- Use thickness corrections via k_z quadrature or standard kernels as needed.

Convergence and padding. Double resolution (grid) or sample count (Hankel) and require relative changes $< 1\%$. For isolated boxes, zeropad by ≥ 2 per dimension and verify stability against padding.

Appendix C: Certificate Fields and Thresholds

Minimum fields.

- Energy $\mathcal{E}$, gap; residual norms (L^2, H^1); windowed residuals T_W.
- Positivity/monotonicity: operator quadratic form min; $P(R)$ sign; $w(R)$ slope sign.
- Convergence: grid/Hankel relative errors; padding factor for isolated systems.
- Kernel: high$k$ deviation from unity; slope signs; kernel checksum.
- Provenance: commit, constants $(\varphi, \alpha, C, \tau_0, G)$, environment pins, seeds, timestamp.

Default thresholds. Convergence $< 1\%$; high$k$ $|\mathsf{w} - 1| < 5\%$; slope signs $\partial_k \mathsf{w} < 0,\ \partial_a \mathsf{w} > 0$; positivity within numerical tolerance; window overlap M recorded; $\sup_W T_W$ below the critical threshold implied by coercivity.

Appendix D: Constants and Provenance

Golden ratio $\varphi = (1+\sqrt{5})/2$; exponent $\alpha = \frac{1}{2}(1-\varphi^{-1})$; prefactor $C = \varphi^{-3/2}$; tick τ_0 catalogglobal; Newton's constant G in consistent units (kpc/Mpc conventions noted in artifacts). Release kernel checksum (SHA256) and constants JSON with each analysis.

Appendix E: Reproducibility Notes

Repository: CPMCosmologyGridPath. Dependencies: Python 3.11, NumPy $\geq$ 1.26, SciPy $\geq$ 1.11. Examples: grid path (growth, lensing, ISW) and disk path (rotation curves). Provide a onecommand script to regenerate figures with certificate sidecars; pin environments and record seeds. Artifacts should include perartifact JSON certificates and a run manifest.

Appendix F: Growth and E_G Details

Growth ODE. Integrate $\ddot{\delta}_b + 2\mathcal{H}\dot{\delta}_b - 4\pi G\, a^2\, \bar{\rho}_b\, \mathsf{w}\, \delta_b = 0$ in a using $\delta \propto a$ initial conditions at early times; report $D(a,k) = \delta/\delta_0$ and $f(a,k) = a\,\dot{\delta}/\delta$.

E_G estimator. Use (17) with surveyspecific geometry factors; report tracerindependent values and consistency across tracers as a certificate item.

Appendix G: NoRetuning Theorem (Sketch)

Assume a single kernel w, fixed constants, and bounded window overlap M. Suppose acceptable residuals require pergalaxy changes to w. Then the projection Π cannot be a unique minimizer of a single energy across the survey, contradicting the coercive projection law (existence/uniqueness with explicit constants). Equivalently, $\sup_W T_W$ computed under the global kernel exceeds the critical threshold implied by coercivity for some systems; replacing w galaxybygalaxy constitutes a change of law rather than a parameter choice. Therefore pergalaxy retuning *falsifies* the universal law.

Appendix H: Notation and Symbols

Symbol	Meaning	Notes
$\Phi(\boldsymbol{x}, a)$	Gravitational potential	Admissible $\Phi \in \mathcal{V}$ (Dirichlet or periodic)
$p(\boldsymbol{x}, a)$	Effective pressure source	$p = \mathsf{w}(\nabla, a)\, s$
s	Raw baryonic source	ρ_b (galaxies), $\bar{\rho}_b\, \delta_b$ (cosmology)
$\mathsf{w}(k, a)$	ILG kernel (Fourier symbol)	$1 + C(a/(k\tau_0))^\alpha$
C, α	Kernel constants	$C = \varphi^{-3/2},\ \alpha = \frac{1}{2}(1 - \varphi^{-1})$ (RSderived)
τ_0	Fundamental tick	Lab $\to$ cosmic gate (global, fixed)
$\mathcal{E}[\Phi\|p]$	Energy functional	$\frac{1}{8\pi G} \int \|\nabla\Phi\|^2 + \int a^2 p\, \Phi$
D_{H^1}	Dirichlet defect	$\int \|\nabla(\Phi - \Phi^*)\|^2$
$\mathsf{D}_{\mathrm{res}}$	Residual defect	$\\|\nabla^2\Phi - 4\pi G a^2 p\\|^2_{H^{-1}}$
Π	Projection	$\Phi^* = \Pi(p)$ unique Poisson solution
K_{net}	Net constant	$((1+\varepsilon)/(1-\varepsilon))^2$; $\varepsilon = 1/8$ (eighttick)
C_{proj}	Projection constant	≤ 2 (rankone/Hermitian bound)
C_{eng}	Energy-control	1 (normalization)
c	Coercivity constant	$49/162 \approx 0.302$ (universal across CPM domains)
$f(a, k)$	Growth rate	$\partial \ln D / \partial \ln a$
$E_G(a, k)$	Tracerindependent ratio	$[4\pi G a^3 \bar{\rho}_b / H]\, \mathsf{w}/f$

Repository and reproducibility.
Code, tests, examples, and certificate templates: `github.com/jonwashburn/CPM-Cosmology-Grid-Path`.