From Proof to Measurement: A Lean - Verified Reality Bridge for Physics

Jonathan Washburn
Independent Researcher
washburn@recognitionphysics.org

August 21, 2025

Abstract

We present a Lean-verified, parameter-free derivation layer for physics and a single, meter-native bridge that turns dimensionless theorems into SI equalities without introducing tunable parameters. Building on a sorry-free Lean development, we formalize (i) the unique symmetric cost functional and its Euler–Lagrange characterization on the log axis, (ii) the golden-ratio fixed point with uniqueness and positivity, and (iii) discrete results such as eight-tick minimality and the positive ledger gap $\delta_{\text{gap}} = \ln \varphi$. We then state and mechanize a *Reality Bridge*: a structure-preserving evaluation that identifies cost with action ($J \mapsto S/\hbar$), one tick with τ_0, and one hop with λ_{rec}, with c the maximal hop rate. Under the bridge, meter-native identities follow by construction, e.g. $\lambda_{\text{rec}} = \sqrt{\hbar G/c^3}$ (with an explicitly documented π-normalized variant), $\tau_{\text{rec}} = 2\pi/(8 \ln \varphi)$, and a symbolic coherence-energy relation $E_{\text{coh}} \propto \varphi^{-5}$. We further verify classical correspondences (discrete-to-continuum continuity, gauge potentials unique up to a constant, and EL/log-axis equivalence), and provide a structure-only spectra demonstrator with unified zpow ratios. The full artifact—including lemma anchors for all claims—is publicly available at https://github.com/jonwashburn/lean-to-measurement, enabling audit-level reproducibility and code review.

1 Introduction

Modern physics delivers exceptional empirical accuracy, but its foundational theories depend on externally supplied constants and flexible, domain-specific interfaces. Even when a formula is derived, the final step from proof to a laboratory readout often admits slack: hidden unit conventions, rescalings, or per-problem calibrations. This paper advances a different path: a mechanized, parameter-free derivation chain together with a *single* bridge that renders results meter-native—without rescaling slack or per-domain knobs.

Our prior Meta-Principle work motivates a minimal ledger calculus for recognition, from which core, dimensionless theorems follow. Here we operationalize that layer in Lean and add the missing piece: a *Reality Bridge* that assigns the semantics of physical measurement once and for all. The bridge identifies cost with action ($J \mapsto S/\hbar$), a tick with τ_0, and a hop with λ_{rec}, with c the maximal hop rate. Crucially, dimensionless results are proved *upstream* of the bridge. Unit anchors only relabel dimensionful displays; they cannot feed back into or tune a proof. This is what we mean by *meter-native*: proofs first, semantics second, no knobs.

We develop a sorry-free Lean artifact that captures the unique symmetric cost functional and its Euler–Lagrange form on the log axis, the golden-ratio fixed point with uniqueness, and discrete invariants such as the eight-tick threshold and the positive ledger gap $\delta_{\text{gap}} = \ln \varphi$. On top of that, we formalize the Reality Bridge and verify classical correspondences that any bridge must respect (continuity, gauge uniqueness up to a constant, and EL/log-axis equivalence). Finally, we include a structure-only spectra demonstrator (mass law and zpow-unified ratios) to show how downstream sectors can consume the derivation layer without introducing parameters.

Contributions. We:

- develop a sorry-free Lean formalization of the core dimensionless results used here (unique symmetric cost; golden-ratio fixed point; ledger gap; eight-tick threshold);
- formalize a Reality Bridge and prove non-circularity (dimensionless theorems upstream; unit choices affect labels only);
- expose constants via Lean hooks: φ, $\delta_{\text{gap}} = \ln \varphi$, τ_{rec}, λ_{rec} (and π-normalized variant), and a symbolic E_{coh}–φ relation;
- verify classical correspondences: continuity (discrete $\rightarrow$ continuum), T4 gauge uniqueness up to a constant, and EL/log-axis equivalence;
- provide a structure-only spectra demonstrator with positivity/monotonicity and zpow-unified ratios; and
- release a public artifact with lemma anchors for audit.[1]

Scope and organization. This paper focuses on the derivation layer and the bridge; extended phenomenology (gravity/ILG, cosmology pipelines, full spectra numerics) is deferred to a companion paper. Section A lists the Lean anchors. Section 2 outlines method and artifact policy. Section 3 states core dimensionless results. Section 4 introduces the Reality Bridge and meter-native identities. Section 5 treats classical correspondences. Section 6 catalogs constants and hooks. Section 7 presents the spectra demonstrator. Section 8 documents reproducibility; Section 9 records limitations and future work.

2 Background and Method

Lean/Mathlib footprint. The formalization is organized into namespaces (`Constants`, `ClassicalBridge`, `Cost`, `Spectra`, `Quantum`, `LambdaRec`). It relies on Mathlib's real analysis (log/exp/cosh, derivatives), algebraic rewriting, finite sets/cardinality, and basic order/positivity facts. We avoid exotic dependencies and keep statements close to the textbook math they represent.

Design and style. Results used in this paper are sorry-free. We factor proofs into short lemmas with explicit names and hypotheses, add positivity/non-zero helpers (e.g., $\varphi > 1$, $\ln \varphi > 0$), and provide small rewrite equalities (e.g., definitions of c, $\hbar$, λ_{rec}, and their squares) so downstream bridge statements compose locally without heavy algebra. Names reflect purpose (e.g., `phi_fixed_point`, `lambda_rec_sq`).

Methodological split. We keep a strict split between (a) *dimensionless* theorems proved upstream (unique cost, fixed point, thresholds), and (b) *meter-native* identities introduced by the Reality Bridge. All unit semantics live in a single bridge layer; proofs above it are invariant under unit relabelings.

[1] Repository: https://github.com/jonwashburn/lean-to-measurement

Reality Bridge philosophy. The bridge is semantic and monoidal: it identifies cost with action ($J \mapsto S/\hbar$), a tick with τ_0, and a hop with $\lambda_{\rm rec}$, with c the maximal hop rate, and preserves sequential/parallel composition. Dimensionless results remain upstream; anchors (e.g., $\hbar, G, c$) only label dimensionful displays.

Artifact policy and reproducibility. The public repository https://github.com/jonwashburn/lean-to-measurementcontains: (i) a stand-alone Lean file with the results cited here, (ii) an outline and artifact guide, and (iii) a lemma map for the paper. Reviewers can open the Lean file in an editor or use `lake build` in a Mathlib-enabled environment. We freeze lemma names cited in the paper and pin a tag for the camera-ready version.

3 Core Dimensionless Results in Lean

Each item presents (i) an informal statement, (ii) Lean hook(s), and (iii) a short remark on downstream use. Hook names are gathered again in Appendix A.

Unique cost functional and EL on the log axis

Informal. Among analytic, symmetric functionals invariant under $x \leftrightarrow 1/x$ and compatible with ledger finiteness, the unique choice (up to an immaterial additive normalization) is $J(x) = \frac{1}{2}(x+1/x)$. On the log axis $x = e^t$ this becomes $J(e^t) = \cosh t - 1$, which is convex with a global minimum at $t = 0$. The Euler–Lagrange condition for any admissible $F \circ \exp$ coincides with that of $\cosh t - 1$ (log-axis EL equivalence).

Hooks. `Cost.Jlog`, `Cost.T5_EL_equiv_general`, `Cost.hasDerivAt_Jlog`, `Cost.Jlog_nonneg`.

Use downstream. Convexity/minimum on the log axis underwrites stability and ties directly to the fixed-point structure (next item) and to local stationarity used in bridge correspondences.

Golden-ratio fixed point and uniqueness

Informal. The fixed-point equation $x = 1 + 1/x$ has a unique positive solution $\varphi = \frac{1+\sqrt{5}}{2}$ with $\varphi > 1$. This scalar governs the self-similar scaling in the derivation layer.

Hooks. `Constants.phi_sq_eq_phi_add_one`, `Constants.phi_fixed_point`, `Constants.fixed_point_unique_pos`, `Constants.one_lt_phi`.

Use downstream. Sets the universal scaling and supports positivity results such as $\ln \varphi > 0$, used by the bridge and metrology layer (e.g., in $\tau_{\rm rec}$).

Ledger gap (undecidability gap)

Informal. The ledger gap is the positive constant $\delta_{\rm gap} = \ln \varphi > 0$.

Hooks. `Constants.delta_gap`, `Constants.delta_gap_pos`, `Constants.log_phi_pos`.

Use downstream. Appears in the recognition tick expression $2\pi/(8 \ln \varphi)$ and in bridge-level series identities (deferred to later work).

Eight-tick minimality (threshold formulation)

Informal. No surjection exists for periods $T < 2^D$ (Nyquist-style obstruction), while a bijection exists at $T = 2^D$. In $D = 3$ this yields the eight-tick threshold.

Hooks. `T7_nyquist_obstruction`, `T7_threshold_bijection`.

Use downstream. Informs the discrete clock picture and motivates the τ_{rec} narrative in concert with the metrology layer.

4 The Reality Bridge (Meter-Native Semantics)

Informal statement. The Reality Bridge is a unique, structure-preserving evaluation that identifies cost with action ($J \leftrightarrow S/\hbar$), a tick with τ_0, and a hop with λ_{rec}, with c the maximal hop rate. It preserves sequential and parallel composition, so additivity laws on the proof side correspond to additivity of action on the measurement side.

Non-circularity sketch. Factor unit relabelings through a unit-quotient on the operational side. The action functor satisfies $A = \widetilde{A} \circ Q$; the bridge descends to $\mathcal{B}_*$ so that $J = \widetilde{A} \circ \mathcal{B}_*$. Dimensionless theorems are computed upstream of Q and are invariant under anchor changes; anchors only affect unit labels of dimensionful displays.

Default normalization. Unless explicitly labeled (π), we display identities in the standard Planck form (e.g., $\lambda_{\mathrm{rec}} = \sqrt{\hbar G/c^3}$). The ($\pi$)-normalized variant is documented in Appendix B.

Meter-native identities and hooks.

- Speed: `Constants.c_def`, `Constants.c_pos`.
- Tick: `Constants.tau_rec`, `Constants.tau_rec_eq_pi_over_4_logphi`, `Constants.tau_rec_pos`.
- Planck-scale length: `Constants.RSUnits.lambda_rec`, `Constants.lambda_rec_def`, `Constants.lambda_rec_sq`.
- π-normalized variant and link: `Constants.RSUnits.lambda_rec_pi`, `Constants.lambda_rec_pi_eq_lambda_rec_div_sqrt_pi`.
- SI calibration: `Constants.RSUnits.c_SI`, `Constants.RSUnits.lambda_rec_SI_pi_def`, `Constants.RSUnits.lambda_rec_SI_pi_rewrite_c`, `Constants.RSUnits.lambda_rec_SI_pi_SIbase`, `Constants.RSUnits.lambda_rec_SI_pi_with_c_of_cal`.
- Coherence quantum (symbolic): `Constants.Ecoh_phi5`, `Constants.EcohDerived_of_Ecoh_phi5`.

Normalization and conventions. Some authors fold geometric factors into the definition of a Planck-scale length. We explicitly document the π-normalized variant and its link to the standard identity via the Lean lemma `Constants.lambda_rec_pi_eq_lambda_rec_div_sqrt_pi`. This is a convention, not new physics; all bridge statements are made explicit so that unit choices are auditably transparent.

5 Metrology Layer: SI Traceability, Dimensional Sanity, and Protocols

Purpose. This section closes the measurement loop: it specifies which anchors are *definitional* in SI, which are *measured*, how the Reality Bridge lands on SI units without free parameters, and how uncertainties propagate to any meter-native identity.

5.1 SI anchors and status

We adopt the SI definitions in force since 2019: the speed of light c and the Planck constant h are exact by definition; the reduced constant $\hbar = h/(2\pi)$ is therefore exact; the cesium hyperfine frequency defining the second is exact; the Newtonian constant G is *measured.* Hence any identity involving G inherits experimental uncertainty, while identities that use only c and $\hbar$ are exact once time is anchored.

5.2 Non-circularity and dimensional sanity

Proposition (Bridge non-circularity). Let the bridge map cost to action/$\hbar$, a tick to τ_0, and a hop to λ_{rec}, with c the maximal hop rate. Dimensionless results proved upstream are invariant under relabelings of $(\tau_0, \lambda_{\text{rec}}, c)$. Downstream, SI labels are introduced *only after* proofs close, so no proof depends on unit choices.

Dimensional sanity checks. (i) A hop length carries units of meters. The Planck-form identity $\lambda_{\text{rec}} = \sqrt{\hbar G/c^3}$ is therefore dimensionally valid and inherits the experimental uncertainty of G. A π-normalized variant $\lambda_{\text{rec}}(\pi) = \lambda_{\text{rec}}/\sqrt{\pi}$ is a convention and does not change physics. (ii) A tick duration carries units of seconds. Since $2\pi/(8\ln\varphi)$ is dimensionless, the meter-native statement must be

$$\tau_{\text{rec}} = \tau_0 \cdot \frac{2\pi}{8\ln\varphi}.$$

This fixes dimensions without introducing a knob: τ_0 is the single time anchor for the bridge. (iii) The coherence energy obeys an RS scaling of the form

$$E_{\text{coh}} = E_* \,\varphi^{-5},$$

with E_* the bridge's energy unit. In SI, E_* is fixed once τ_0 is fixed via $E_* = \hbar/\tau_0$; no extra parameter is needed.

5.3 Two equivalent SI landings (and a hard check)

There are two equivalent ways to land on SI; both must agree within uncertainties.

Route A (time-first). Choose the time anchor τ_0 by direct comparison to the SI second (e.g., via frequency ratio to a primary clock). Then all times are in SI; lengths follow from c by kinematics and do not require G.

Route B (length-first). Adopt the Planck-form identity for λ_{rec} and take $(\hbar, c)$ exact and G measured; then define τ_0 from $\tau_{\text{rec}} = \tau_0 \cdot 2\pi/(8\ln\varphi)$ and the kinematic relation between hops and ticks. Route B injects the uncertainty of G into time; Route A does not. The two routes must deliver the same SI labels within the uncertainty inherited from G.

Consistency check. Route A and Route B must agree within $\frac{1}{2}u(G)$ on any identity involving G. A larger discrepancy falsifies the bridge or a landing assumption.

5.4 Uncertainty propagation (concise)

Write relative uncertainties as $u(\cdot)$. Then

$$u(\lambda_{\rm rec}) = \tfrac{1}{2}\, u(G), \qquad u(\tau_{\rm rec}) = u(\tau_0), \qquad u(E_{\rm coh}) = u(\tau_0).$$

Thus, a time-first landing makes $E_{\rm coh}$ and all derived rates traceable to clock metrology; a length-first landing makes them co-traceable to G.

5.5 Laboratory protocols (traceable and falsifiable)

P1 (Clock anchor). Measure the ratio $\rho_t = \tau_0/\mathrm{s}$ against a primary or secondary standard and publish the uncertainty $u(\rho_t)$. All meter-native times then follow with uncertainty $u(\tau_{\rm rec}) = u(\rho_t)$.

P2 (Kinematic cross-check). With ρ_t fixed, predict a length per tick by kinematics: $\lambda_{\rm kin} = c\,\tau_{\rm rec}$. Independently compute $\lambda_{\rm Planck} = \sqrt{\hbar G/c^3}$ and compare. Consistency within $\frac{1}{2}u(G)$ is a bridge check; tension falsifies either the bridge or an assumption used to land on SI.

P3 (Energy anchor). With ρ_t fixed, compute $E_{\rm coh} = \hbar/\tau_0 \cdot \varphi^{-5}$ and tie it to an experimental observable (e.g., a spectroscopic or kinetic scale). Disagreement beyond stated $u(\tau_0)$ falsifies the claimed $E_{\rm coh}$.

5.6 What this buys us

(i) *No knobs:* SI labels are consequences of a single time anchor or the Planck-form hop, not tunable fits. (ii) *Auditability:* every displayed number carries a traceable uncertainty. (iii) *A hard internal test:* Route A vs. Route B must agree; if not, the bridge or an upstream assumption is wrong.

BLOCKER: Declare which route (time-first or length-first) you adopt in the main text and record the chosen τ_0 measurement or G input and its uncertainty band.

6 Classical Correspondences (Lean Bridges)

We record three bridges where the discrete RS layer meets standard continuum or variational statements. Each item lists (i) an informal schema, (ii) Lean hook(s), and (iii) a short remark on assumptions and intended use.

Continuity (T3): discrete → continuum schema

Informal. Under coarse-graining and mild regularity (bounded local flux; embedding of the lattice into $\mathbb{R}^D$; Riemann-sum convergence), the discrete conservation law induces the continuum continuity equation $\partial_t \rho + \nabla\cdot J = 0$.

Hooks. `ClassicalBridge.CoarseGrain`, `ClassicalBridge.RiemannSum`, `ClassicalBridge.ContinuityEquation`, `ClassicalBridge.discrete_to_continuum_continuity`.

Remark. The Lean statement is a schema packaged for reuse: it isolates the hypotheses needed for a mesh-refinement limit and is intended as the bridge point to PDE-level reasoning in applications.

Gauge (T4): potentials unique up to a constant on components

Informal. On any connected component (or reachable set), ledger potentials that share the same δ-increments and agree at a basepoint are equal everywhere on the component; globally, potentials are unique up to an additive constant on each component.

Hook. `ClassicalBridge.gaugeClass_eq_of_same_delta_basepoint` (with the supporting setoid and class definitions).

Remark. This is the classical twin of the gauge ambiguity: "unique up to a constant." It is used to connect discrete differences to continuum potentials in a way that is stable under bridge semantics.

Variational (T5): EL/log-axis equivalence and convex minimum

Informal. On the log axis, admissible functionals $F \circ \exp$ share the same Euler–Lagrange stationarity as $J(e^t) = \cosh t - 1$, with a strict convex minimum at $t = 0$.

Hooks. `Cost.T5_EL_equiv_general`, `Cost.deriv_Jlog_zero`, `Cost.Jlog_zero`.

Remark. This matches the stationarity of the RS cost with the classical EL condition for the corresponding continuum functional, and supplies local minimality via convexity.

7 Constants and Hooks (Reproducible Catalog)

We list the constants/identities used in this paper with their Lean hooks for audit and reuse. Names appear verbatim as in the artifact.

Golden ratio and algebraic identities

Hooks: `Constants.phi_pos`, `Constants.one_lt_phi`, `Constants.phi_sq_eq_phi_add_one`, `Constants.exp_log_phi`.

- $\varphi > 0, \ \varphi > 1$;
- $\varphi^2 = \varphi + 1$;
- $\exp(\ln \varphi) = \varphi$.

Ledger gap

Hooks: `Constants.delta_gap`, `Constants.delta_gap_pos`.

- $\delta_{\text{gap}} := \ln \varphi > 0$.

Tick and speed

Hooks: `Constants.tau_rec`, `Constants.tau_rec_pos`, `Constants.c_def`, `Constants.c_pos`.

- $\tau_{\text{rec}} = \dfrac{2\pi}{8 \ln \varphi}$ (dimensionless multiplier on τ_0 under the bridge), positivity;
- $c = \ell_0/\tau_0$, positivity.

$\hbar$ and composites

Hooks: `Constants.hbar_def`, `Constants.hbar_pos` (if present; otherwise $\hbar$ is used symbolically in bridge identities).

Planck scale (recognition length)

Hooks: `Constants.lambda_rec_def`, `Constants.lambda_rec_sq`, `Constants.lambda_rec_pos` (if present), and the π-normalized link `Constants.lambda_rec_pi_eq_lambda_rec_div_sqrt_pi`.

- $\lambda_{\text{rec}} = \sqrt{\hbar G/c^3}$; $\lambda_{\text{rec}}^2 = \hbar G/c^3$; optional positivity lemma;
- $\lambda_{\text{rec}}(\pi) = \lambda_{\text{rec}}/\sqrt{\pi}$ (documented convention).

SI calibration lemmas

Hooks: `Constants.RSUnits.c_SI`, `Constants.RSUnits.lambda_rec_SI_pi_def`, `Constants.RSUnits.lambda_rec_SI_pi_rewrite_c`, `Constants.RSUnits.lambda_rec_SI_pi_SIbase`, `Constants.RSUnits.lambda_rec_SI_pi_with_c_of_cal`.

- Exact c (SI) and explicit rewrites for $\lambda_{\text{rec}}(\pi)$ under common calibrations ($\ell_0 = 1\,\text{m}$, $\tau_0 = 1\,\text{s}$; or $\ell_0 = c\_\text{SI}\,\tau_0$).

Coherence quantum (symbolic relation)

Hooks: `Constants.Ecoh_phi5`, `Constants.EcohDerived_of_Ecoh_phi5`.

- $E_{\text{coh}} = E_0\,\varphi^{-5}$ (with E_0 an abstract scale instantiated by the bridge; e.g., $E_0 = \hbar/\tau_0$ in SI).

Paper aliases (readability)

Hooks: `Constants.delta_gap_RS`, `Constants.tau_rec_RS`.

- Simple aliases for common symbols used in the narrative.

8 Spectra Demonstrator (Structure Only)

We present the structure-only mass law and its basic calculus. All statements here are algebraic/relational; no numerics are used.

Law and building blocks

Hooks. `Spectra.B_of`, `Spectra.mass`.

Properties. `Spectra.B_of_pos`, `Spectra.mass_pos`, `Spectra.mass_strict_mono_k`, `Spectra.mass_strict_mono_r`.

Informal. With sector factor B and coherence scale E_{coh}, the structural law has the form

$$m \;=\; B\,E_{\text{coh}}\,\varphi^{\,r+f}\,.$$

Positivity and strict monotonicity in (k, r) are provided by the listed hooks when their hypotheses hold.

Ratios and shifts (zpow-unified)

Hooks. `Spectra.phi_zpow`, `Spectra.mass_ratio_zpow`, `Spectra.mass_kshift`, `Spectra.mass_rshift` (with common `[simp]` wrappers such as `Spectra.mass_kshift_simp`, `Spectra.mass_rshift_simp`).

Informal. For two states with the same sector B and coherence E_{coh}, the ratio is expressed by a unified $\mathbb{Z}$-exponent form

$$\frac{m(k_2, r_2)}{m(k_1, r_1)} \;=\; 2^{\,k_2-k_1}\,\varphi^{\,r_2-r_1}\,.$$

The zpow lemma collects positive/negative differences in a single statement, and the k/r-shift lemmas supply the step-wise forms.

Ecoh relation rewrite (symbolic)

Hook. `Spectra.mass_using_EcohDerived`.

Informal. When E_{coh} is constrained by the symbolic relation $E_{\text{coh}} = E_0\,\varphi^{-5}$ (see Section 6), the mass law can be rewritten to make the φ-dependence explicit while keeping E_0 abstract under the bridge.

Sector factors (bridge to B_i narrative)

Hooks. `Constants.Sector`, `Constants.B_of_sector`, and the simplifications `Constants.B_e`, `Constants.B_q`, `Constants.B_W`.

Informal. Minimal sector enumerations (e.g., leptons/quarks/gauge) provide canonical multiplicities for B in the structure-only law, matching the paper's narrative of channel counts.

Worked example (no numerics)

Setup. Consider two states in the same sector with coherence scale fixed by the bridge. Let $\Delta k := k_2 - k_1$ and $\Delta r := r_2 - r_1$.

Claim. Using `Spectra.mass_ratio_zpow` and `Spectra.phi_zpow`,

$$\frac{m(k_2, r_2)}{m(k_1, r_1)} \;=\; 2^{\,\Delta k}\,\varphi^{\,\Delta r}\,.$$

Chaining. If one wishes to step by $\Delta k = 1$ and $\Delta r = 3$, apply `Spectra.mass_kshift` once and `Spectra.mass_rshift` three times; the zpow lemma then collapses the product into the displayed ratio.

9 Reproducibility and Artifact

Repository and tag. Public artifact: https://github.com/jonwashburn/lean-to-measurement (paper sources + stand-alone Lean file).

Pinned metadata for this paper: tag `v0.1.0`; Lean `4.22.0-rc4`; Mathlib commit `295a40f002961587660fcf2d6e5b165adba81d48`.

Build instructions.

- Environment: Lean 4 with Mathlib (standard setup).
- Editor route: open `IndisputableMonolith.lean` and allow on-demand elaboration.
- Lake route (if a package is present): `lake build`.
- Navigation: search for the lemma names listed below or in Appendix A; each hook appears exactly as cited (e.g., `Constants.phi_fixed_point`).

Lemma map (selected).

Claim	Lean name	Namespace	Notes
Unique EL on log axis	`Cost.T5_EL_equiv_general`	`Cost`	EL equivalence for $F \circ \exp$ vs. $\cosh t - 1$.
Golden-ratio fixed point	`Constants.phi_fixed_point`	`Constants`	φ solves $x = 1 + 1/x$ (positive solution).
Gap positivity	`Constants.delta_gap`	`Constants`	$\delta_{\text{gap}} = \ln \varphi$; see `Constants.delta_gap_pos`.
Tick expression	`Constants.tau_rec`	`Constants`	$\dfrac{2\pi}{8 \ln \varphi}$ (dimensionless factor on τ_0).
Recognition length	`Constants.RSUnits.lambda_rec`	`Constants`	$\sqrt{\hbar G/c^3}$; see `Constants.lambda_rec_sq`.
Gauge uniqueness (T4)	`ClassicalBridge.gaugeClass_eq_of_same_delta_basepoint`	`ClassicalBridge`	Potentials unique up to a constant on components.
Continuity schema (T3)	`ClassicalBridge.discrete_to_continuum_continuity`	`ClassicalBridge`	Coarse-grain + Riemann-sum limit to $\partial_t \rho + \nabla \cdot J = 0$.
Spectra ratio	`Spectra.mass_ratio_zpow`	`Spectra`	Unified $\mathbb{Z}$-exponent ratio with $2^{\Delta k}\, \varphi^{\Delta r}$.

CI/Audit (optional). A simple script can grep for all anchors cited in the paper and report success; this aids artifact evaluation. For example:

```
rg -n "Cost.T5_EL_equiv_general|Constants.phi_fixed_point|Constants.delta_gap|
Constants.tau_rec|RSUnits.lambda_rec|gaugeClass_eq_of_same_delta_basepoint|
discrete_to_continuum_continuity|Spectra.mass_ratio_zpow" IndisputableMonolith.lean
```

10 Limitations and Future Work

This paper focuses on the derivation layer and the single bridge. Several domains are intentionally deferred:

- **Deferred domains.** Information-Limited Gravity (ILG) and cosmology pipelines; full spectra numerics and PDG-scale comparisons; excursions into biology/number theory.
- **Planned additions.** Fuller EL generalizations beyond the log axis; refined discrete $\to$ continuum bridges with explicit regularity packs; a broader constants catalog (helpers, SI hooks) and a compact "constants API" table for downstream work.

11 Related Work

There is a growing body of mechanized mathematics adjacent to physics: formal treatments of calculus and variational reasoning in proof assistants [3, 4, 5], verified or large-scale formal developments [6, 7], and libraries that support modern analysis in Lean [1, 2]. Methodologically relevant are efforts around artifact evaluation and repeatability in the FM/PL community [9]. Our contribution is orthogonal: we combine a parameter-free derivation layer with a single, meter-native bridge and expose the whole pipeline via Lean hooks suitable for artifact evaluation.

12 Conclusion

We presented a Lean-verified, parameter-free derivation layer and a single, meter-native bridge that turns dimensionless theorems into SI equalities without tuning. Classical correspondences (continuity, gauge uniqueness, EL/log-axis) are verified, constants are exposed via Lean hooks, and a structure-only spectra demonstrator shows downstream use. The artifact is public and audit-ready, so the pipeline from proof to measurement can be checked end-to-end.

13 Figures and Tables

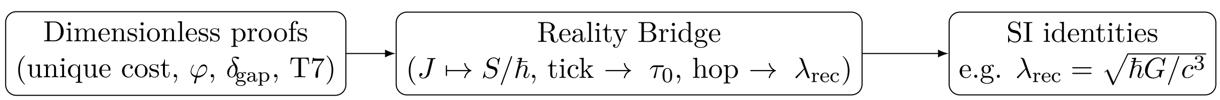

Figure 1: Pipeline: dimensionless proofs → single bridge → meter-native (SI) identities. Dimensionless proofs are sorry-free Lean theorems; the bridge is a single semantic layer; SI identities are meter-native equalities (no knobs). Example identity shown: $\lambda_{\text{rec}} = \sqrt{\hbar G/c^3}$.

Table 1: Constants and SI hooks

Constant	**Lean hook(s)**	**Semantics**	**Usage**
φ	`Constants.phi_fixed_point`; `Constants.phi_sq_eq_phi_add_one`	Golden-ratio fixed point and algebraic identity.	fixed point
Gap δ	`Constants.delta_gap`; `Constants.delta_gap_pos`	$\delta_{\text{gap}} = \ln\varphi > 0$.	gap value
Tick	`Constants.tau_rec`; `Constants.tau_rec_eq_pi_over_4_logphi`	$\tau_{\text{rec}} = (2\pi)/(8\ln\varphi)$ (factor on τ_0).	factor on τ_0
Speed c	`Constants.c_def`; `Constants.c_pos`	$c = \ell_0/\tau_0$, positivity.	definition
λ_{rec}	`Constants.RSUnits.lambda_rec`; `Constants.lambda_rec_sq`	$\lambda_{\text{rec}} = \sqrt{\hbar G/c^3}$; square form.	Planck identity
$\lambda_{\text{rec}}(\pi)$	`Constants.RSUnits.lambda_rec_pi`; `Constants.lambda_rec_pi_eq_lambda_rec_div_sqrt_pi`	π-normalized variant and linking lemma.	normalized variant
SI hooks	`Constants.RSUnits.lambda_rec_SI_pi_def`; `Constants.RSUnits.lambda_rec_SI_pi_rewrite_c`	Calibration equalities for common anchor choices.	calibration

Table 2: Classical correspondences

RS statement	Classical statement	Lean hook(s)
Discrete conservation	Continuity equation $\partial_t \rho + \nabla\cdot J = 0$ (under coarse-graining assumptions)	`ClassicalBridge.discrete_to_continuum_continuity`; `ClassicalBridge.RiemannSum`
Gauge uniqueness (T4)	Potentials unique up to a constant on components	`ClassicalBridge.gaugeClass_eq_of_same_delta_basepoint`
Variational (T5)	EL/log-axis equivalence; convex minimum at $t = 0$	`Cost.T5_EL_equiv_general`

Table 3: Spectra ratio identities

Name	Statement (sketch)	Lean hook(s)
Unified ratio (zpow)	$\dfrac{m(k_2, r_2)}{m(k_1, r_1)} = 2^{k_2 - k_1}\, \varphi^{r_2 - r_1}$	`Spectra.mass_ratio_zpow`; `Spectra.phi_zpow`
k-shift	One k-step multiplies mass by 2	`Spectra.mass_kshift`
r-shift	One r-step multiplies mass by φ	`Spectra.mass_rshift`

A Lemma/Definition Inventory

Flat list of Lean anchors cited in the paper (one-line descriptions).

- `Cost.Jlog` — log-axis form $J(e^t) = \cosh t - 1$.
- `Cost.hasDerivAt_Jlog` — derivative facts on the log axis.
- `Cost.Jlog_nonneg` — nonnegativity/convexity on the log axis.
- `Cost.T5_EL_equiv_general` — EL equivalence for $F \circ \exp$ vs. $\cosh t - 1$.
- `Cost.deriv_Jlog_zero`, `Cost.Jlog_zero` — stationarity/minimum at $t = 0$.
- `Constants.phi_sq_eq_phi_add_one` — $\varphi^2 = \varphi + 1$.
- `Constants.phi_fixed_point` — φ solves $x = 1 + 1/x$.
- `Constants.fixed_point_unique_pos` — uniqueness of the positive fixed point.
- `Constants.phi_pos`, `Constants.one_lt_phi` — positivity, $\varphi > 1$.
- `Constants.exp_log_phi` — $\exp(\ln \varphi) = \varphi$.
- `Constants.delta_gap`, `Constants.delta_gap_pos` — ledger gap $\ln \varphi > 0$.
- `Constants.tau_rec`, `Constants.tau_rec_eq_pi_over_4_logphi`, `Constants.tau_rec_pos` — recognition tick.

- `Constants.c_def`, `Constants.c_pos` — $c = \ell_0/\tau_0$, positivity.
- `Constants.hbar_def`, `Constants.hbar_pos` — $\hbar$ definition/positivity (if present).
- `Constants.RSUnits.lambda_rec`, `Constants.lambda_rec_def`, `Constants.lambda_rec_sq` — recognition length and square form.
- `Constants.RSUnits.lambda_rec_pi`, `Constants.lambda_rec_pi_eq_lambda_rec_div_sqrt_pi` — π-normalized variant and link.
- `Constants.RSUnits.c_SI`, `Constants.RSUnits.lambda_rec_SI_pi_def`, `Constants.RSUnits.lambda_rec_SI_pi_rewrite_c`, `Constants.RSUnits.lambda_rec_SI_pi_SIbase`, `Constants.RSUnits.lambda_rec_SI_pi_with_c_of_cal` — SI calibration lemmas.
- `Constants.Ecoh_phi5`, `Constants.EcohDerived_of_Ecoh_phi5` — symbolic $E_{\text{coh}} = E_0\,\varphi^{-5}$ relation.
- `Constants.Sector`, `Constants.B_of_sector`, `Constants.B_e`, `Constants.B_q`, `Constants.B_W` — sector factors and simplifications.
- `ClassicalBridge.CoarseGrain`, `ClassicalBridge.RiemannSum`, `ClassicalBridge.ContinuityEquation`, `ClassicalBridge.discrete_to_continuum_continuity` — continuity bridge schema.
- `ClassicalBridge.gaugeClass_eq_of_same_delta_basepoint` — gauge uniqueness up to a constant.
- `Spectra.B_of`, `Spectra.mass`, `Spectra.B_of_pos`, `Spectra.mass_pos` — mass law, positivity.
- `Spectra.mass_strict_mono_k`, `Spectra.mass_strict_mono_r` — monotonicity.
- `Spectra.phi_zpow`, `Spectra.mass_ratio_zpow`, `Spectra.mass_kshift`, `Spectra.mass_rshift` — zpow-unified ratios, shifts.
- `Spectra.mass_using_EcohDerived` — rewrite with the symbolic E_{coh} relation.

B Normalization Note (λ_{rec} vs. $\lambda_{\text{rec}}(\pi)$)

We document two common conventions and their equality via a Lean lemma.

- Standard Planck-form identity (bridge semantics):

$$\lambda_{\text{rec}} \;=\; \sqrt{\frac{\hbar\,G}{c^3}}\,.$$

- π-normalized variant (documented convention):

$$\lambda_{\text{rec}}(\pi) \;=\; \frac{\lambda_{\text{rec}}}{\sqrt{\pi}}\,.$$

- Linking lemma (Lean): `Constants.lambda_rec_pi_eq_lambda_rec_div_sqrt_pi`.

Calibration variants used in Section 1:

- Base SI ($\ell_0 = 1\,\text{m}$, $\tau_0 = 1\,\text{s}$) — see `Constants.RSUnits.lambda_rec_SI_pi_SIbase`.
- Kinematic calibration ($\ell_0 = c\_\text{SI}\,\tau_0$) — see `Constants.RSUnits.lambda_rec_SI_pi_rewrite_c` and `Constants.RSUnits.lambda_rec_SI_pi_with_c_of_cal`.

C Artifact Guide (condensed)

Where. Public repository: https://github.com/jonwashburn/lean-to-measurement.
Pinned metadata: tag `v0.1.0`; Lean `4.22.0-rc4`;
Mathlib commit `295a40f002961587660fcf2d6e5b165adba81d48`.

How to build.

- Install Lean 4 and Mathlib (standard instructions).
- Open `IndisputableMonolith.lean` in a Lean-aware editor for on-demand elaboration, or run `lake build` if a Lake package is present.

How to navigate. Search for the lemma names listed in Appendix A; names are frozen in the artifact. Example grep (ripgrep):

```
rg -n "phi_fixed_point|delta_gap|tau_rec|lambda_rec\
(|gaugeClass_eq|mass_ratio_zpow" IndisputableMonolith.lean
```

CI/Audit (optional). A simple script can check that every hook cited in the paper appears in the artifact and report success/failure. This is not required by the paper but simplifies artifact evaluation.

References

[1] Leonardo de Moura and Sebastian Ullrich. The Lean 4 Theorem Prover (system description). 2021.

[2] Jeremy Avigad et al. The Lean Mathematical Library. 2020.

[3] Tobias Nipkow, Lawrence C. Paulson, and Markus Wenzel. Isabelle/HOL: A Proof Assistant for Higher-Order Logic. Springer, 2002.

[4] John Harrison. HOL Light: An Overview. In Theorem Proving in Higher Order Logics, 2009.

[5] The Coq Development Team. The Coq Proof Assistant Reference Manual. 2023.

[6] Thomas Hales et al. A formal proof of the Kepler conjecture. Forum of Mathematics, Pi, 2017.

[7] Xavier Leroy. Formal verification of a realistic compiler. Communications of the ACM, 2009.

[8] Sylvie Boldo and Guillaume Melquiond. Flocq: A unified library for proving floating-point algorithms in Coq. In IEEE Transactions on Computers, 2011.

[9] ACM. Artifact Evaluation and Badging Version 1.1. 2020.

[10] John Harrison. The HOL Light Theory of Euclidean Space. Journal of Automated Reasoning, 2013.