

# Operational Law Closure from Finite Distinguishability

## Why Effective Dynamics Lives on the Stable Record Algebra

---

### For the General Reader

#### What this paper is about

Physics is full of laws — Newton's laws, Maxwell's equations, Einstein's field equations. But there is a prior question that physics rarely asks explicitly: **what makes something qualify as a macroscopic law in the first place?**

This paper gives a precise answer grounded in the structure of physical records.

Any real law must be storable, testable, and comparable between observers. This means the variables appearing in a law must correspond to distinctions that can be stably written, reliably read, and meaningfully exchanged between different observers. A distinction that evaporates before it can be read, or that two observers cannot compare, cannot serve as a primitive variable in a physical law.

From this observation — together with the fact that all records require finite resources, and that certain distinctions once made cannot be undone — we derive a structural theorem: **macroscopic laws must be defined on the algebra of stable, committed, observer-comparable distinctions.** No degree of freedom lacking these properties can appear as an independent primitive variable in any admissible law.

A key distinction runs throughout: we are not claiming that physics in any direction *disappears*. Microscopic fields still exist and can still exert forces. What we claim is structural: if a direction cannot support independent stable records, it cannot contribute an *independent state variable* to any macroscopic law. Its influence persists, absorbed into the effective couplings of the reduced theory rather than constituting a separate degree of freedom. This is the standard logic of effective field theory — integrating out a direction reduces the arena of independent variables, not the arena of physical influence.

---

#### What the two papers prove together

This paper is one half of a two-paper programme. The companion paper (*Finite Distinguishability and Dimensional Reduction*) handles the geometry: it proves that thin physical

regions — specifically, galactic disk slabs thinner than a characteristic scale  $\ell_{\text{eff}}$  — cannot support independent three-dimensional stable record structure. The present paper handles the law theory: it proves that wherever stable records cannot be maintained along a direction, the effective macroscopic laws governing that region must reduce dimensionally. Each paper closes the other's logical gap, and taken together they form a complete derivation.

**The complete argument in plain language runs as follows:**

Everything in physics that we call a law must ultimately be defined on things that can be stably recorded, reliably read, and compared between observers. In a region so thin that its vertical dimension falls below the minimum scale at which stable records can be maintained — which is set by the physics of record-keeping itself, not by any external assumption — the vertical direction loses the ability to support independent stable records. By the law-closure theorem proved here, it therefore loses the ability to contribute an independent state variable to any macroscopic law. The effective laws of that thin region must be two-dimensional. The only local, isotropic, second-order field equation in two dimensions is the 2D Laplace equation, whose gravitational potential falls off logarithmically rather than as  $1/r$ . This logarithmic falloff produces flatter rotation curves than three-dimensional Newtonian gravity predicts — without any new force law, dark matter halo, or modified gravity theory.

**The physical predictions**

When the two papers are combined, the derivation produces concrete, testable predictions about galactic rotation curves:

**A universal surface density threshold — derived, not fitted.** There is a characteristic surface density  $\Sigma^*$  at which the transition from three-dimensional to two-dimensional effective gravity occurs in galactic disks. Under vertical hydrostatic equilibrium, imposing the dimensional transition trigger directly yields the prediction  $\Sigma^* \sim \sigma_z^2 / (\pi G \ell_{\text{eff}})$ . Two of the three inputs come from atomic physics:  $\sigma_z \sim 10 \text{ km s}^{-1}$  is the sound speed of thermally stable neutral hydrogen at  $T \sim 8000 \text{ K}$  (set by the balance between photoelectric heating and atomic fine-structure cooling, the same in every outer disk); the atomic cooling function  $\Lambda(T)$  at 8000 K is fixed by atomic transition rates. The third input — the outer-disk HI density  $n \sim 0.5\text{--}2 \text{ cm}^{-3}$  — is a galaxy-scale parameter that varies across the galaxy population. It is not fixed by atomic physics. It is constrained to this range observationally: the galaxies that exhibit MOND-like rotation curves are those with outer-disk HI conditions in this regime, and this range is consistent with (though not predicted by) the condition  $h < \ell_{\text{eff}}$  that triggers the dimensional transition. The output is  $\Sigma^* \approx 148 \text{ M}\odot \text{ pc}^{-2}$ , consistent with the observed value of  $\sim 137 \text{ M}\odot \text{ pc}^{-2}$ .

**A characteristic acceleration scale — derived relation, predicted band.** The 2D Laplace equation implies a characteristic acceleration  $a_0 = 2\pi G \Sigma^* \sim 2\sigma_z^2 / \ell_{\text{eff}}$ . This is a derived relationship: once  $\sigma_z$  and  $\ell_{\text{eff}}$  are constrained by outer-disk microphysics,  $a_0$  follows. The claim is not that  $a_0$  is parameter-free; it is that the two inputs are themselves constrained by atomic physics to a narrow band —  $\sigma_z \sim 8\text{--}12 \text{ km s}^{-1}$  from the WNM thermal floor,  $\ell_{\text{eff}} \sim$

20–100 pc from the atomic cooling relaxation length — and that this band predicts  $a_0$  must lie in the MOND decade ( $\sim 6 \times 10^{-11}$  to  $\sim 3 \times 10^{-10}$  m s<sup>-2</sup>). The physical meaning of  $a_0$  is now precise: it is the acceleration at which a baryonic slab becomes too thin to maintain independent three-dimensional record channels. The reason it is nearly universal across galaxies is the same reason the boiling point of water is universal: both inputs are set by atomic constants, not by galaxy-scale properties.

**A falsifiable scaling law:  $a_0 \propto n \Lambda(T) / T^{\{1/2\}}$ .** Because both  $\sigma_z$  and  $\ell_{\text{eff}}$  are anchored in ISM conditions, the programme predicts that  $a_0$  should vary in a structured way across galaxies with different outer-disk environments: denser HI disks should have slightly larger  $a_0$ ; lower-metallicity galaxies (less efficient cooling) should have slightly smaller  $a_0$ . These are quantitative, testable predictions that MOND and dark matter models cannot make, accessible with current SPARC and HI survey data.

**A directly testable formula.** At the dimensional transition,  $\ell_{\text{eff}} \approx h(r_h)$ , which simplifies the prediction to  $a_0 \sim 2\sigma_z^2/h(r_h)$ . For any galaxy where the vertical velocity dispersion and disk scale height at the transition radius can be measured, this formula provides a clean test without any dark matter modelling.

**Thickness dependence — a prediction MOND cannot make.** Because the transition is triggered by disk thickness, the radius at which rotation curves depart from Newtonian predictions should be correlated with independently measurable disk scale heights. Thicker disks should show the transition at larger radii. No dark matter model or MOND-type theory predicts this correlation, because neither connects the transition to physical disk geometry. It is the sharpest single observational test of the mechanism.

**A kinematic lag.** The stable records that define the transition propagate at the speed of the baryonic matter, not at the orbital velocity. The geometric transition occurs where  $h(r) = \ell_{\text{eff}}$ ; the kinematically observed break in the rotation curve occurs where the baryonic record dynamics has had time to establish a stable committed partition at the new scale, which lags the geometric condition by one propagation time — roughly  $\sigma_z/v_{\text{orb}}$  times the transition circumference, of order 0.5–1 disk scale lengths. This means the kinematically observed transition radius lags the geometric transition radius by roughly 1.5–3 kiloparsecs in Milky Way-scale galaxies (derived in the companion paper, §5.3). This lag is measurable in resolved rotation curve data and absent in all MOND and dark matter predictions.

---

## How this differs from MOND and dark matter

MOND (Modified Newtonian Dynamics) introduces a new, phenomenological force law that interpolates between standard Newtonian gravity at high accelerations and a modified regime at low accelerations, with  $a_0$  as a fitted parameter whose physical origin is unknown. The present framework does not modify the force law. It derives the same phenomenology — the flattening of rotation curves, the value of  $a_0$  — from the structure of physical records and the geometry of thin disks. The two frameworks are not equivalent: the present framework predicts thickness

dependence and kinematic lag that MOND does not, and it explains *why*  $a_0$  has the value it does rather than treating it as a given.

Dark matter models add an unseen mass component whose density profile is fitted to match observed rotation curves. The present framework requires no additional matter. The rotation curve modification comes entirely from the change in the effective dimensionality of gravity in thin regions, driven by the baryonic geometry alone.

---

## What remains assumed

The derivation is transparent about what it takes as input and what it derives. Three inputs are adopted as axioms rather than derived: the requirement that law statements close under record exchange between observers — Axiom LC, from which observer comparability is derived as Proposition 2.1 (motivated by, and consistent with, the gauge-invariance structure of all known fundamental physics); the standard effective- field-theory result that integrated-out modes produce local corrections at long distances; and the scalar, isotropic, second-order character of the thin-disk gravitational effective theory. Everything else — the dimensional reduction, the logarithmic Green's function, the value of  $a_0$ , the thickness dependence, the kinematic lag — follows from these inputs and the geometry of the disk.

---

## The paper's structure

The derivation proceeds in two explicitly separated stages:

**Stage 1 — Law closure.** From three axioms (finite distinguishability, irreversible commitment, and law closure under record exchange — from which observer comparability is derived), we prove that the admissible macroscopic state space has the same dimension as the number of independent stable partition channels. When a spatial direction loses stable channels, it loses independent macroscopic state variables.

**Stage 2 — Operator selection.** Given that the effective theory lives on the reduced two-dimensional base, symmetry arguments — locality, isotropy, and second-order derivatives — select the unique field operator. The logarithmic Green's function follows. This stage requires additional physical assumptions beyond the three axioms, stated explicitly.

The result: the Bridge Principle assumed in the companion paper — that effective dynamics lives on the stable record algebra — is not a free assumption. It is a consequence of a coherent account of what macroscopic laws are. When that result is combined with the companion paper's geometric theorem, the derivation runs from first principles to observable galactic rotation curve phenomenology without postulating new forces, new particles, or new fundamental fields.

---

# Table of Contents

1. Motivation and Scope
2. Axiomatic Framework — Principle PCP: Prediction Composition Principle (*definitional of prediction; grounds LI*) — Axiom FD: Finite Distinguishability — Axiom IC: Irreversible Commitment — Definition 2.1: Admissible Representation Change — Axiom LI: Law Invariance Under Admissible Representation Change (*corollary of PCP*) — Theorem 2.LC: Law Closure derived from LI (*representation theorem; LC as invariant-subalgebra result*) — Proposition 2.0: Contextual Collapse under LI Failure (*structural necessity*) — Operational Metric  $\mathcal{D}_\mathcal{O}$  (*introduced here alongside IC*) — Proposition 2.1: Observer comparability as consequence of LI + FD + IC
3. Stable Record Algebra — Definition 3.1: Stable Record Partition — Definition 3.2: Stable Record  $\sigma$ -Algebra ( $\mathcal{A}_{\text{stable}}$ ) — Definition 3.3: Stable Equivalence and Macroscopic State Space — Definition 3.4: Record Coherence Length  $\ell_{\text{eff}}$  (*defined in record-theoretic terms*) — Definition 3.5: Commitment Time  $\tau_{\text{commit}}$  (*intrinsic  $\mathcal{D}_\mathcal{O}$  definition; cooling time as empirical realization*) — Theorem 3.1: Mixing-Commitment Competition Determines  $\ell_{\text{eff}}$  (*ballistic, diffusive, turbulent channels*) — Lemma 3.2: Dominant Mixing Channel Determines  $\ell_{\text{eff}}$  (*fastest channel governs; H2 as channel-shift*)
4. Admissible Macroscopic Observables — Definition 4.1: Macroscopic Observable (*four conditions including law-eligibility*) — Definition 4.2: Counterfactual Stability Condition
5. The Representation Theorem — Theorem 5.1: Representation of law-eligible observables (*nontrivial content in Lemmas 5.2c, 5.2*) — Definition 5.2a:  $\mathcal{D}_\mathcal{O}$  restated for reference (*introduced in §2*) — Lemma 5.2b: Monotonicity of  $\mathcal{D}_\mathcal{O}$  under admissible channels (*data-processing inequality*) — Lemma 5.2c: No Resurrection by Aggregation (*stated before Lemma 5.2; closes aggregation loophole*) — Lemma 5.2: Non-Accumulation of Uncommitted Fluctuations (*direct pathway; cites 5.2c for indirect*)
6. Law-Admissibility — Definition 6.1: Law-Admissibility — Definition 6.2: Independent Macroscopic State Variable ( *$\sigma$ -algebra measurability criterion*) — Lemma 6.1:  $\sigma$ -Algebra Non-Enlargement (*proper  $\sigma$ -algebra proof of independence exclusion*)
7. Operational Law Closure Theorem — Theorem 7.1: Admissible laws are functionals on  $\mathcal{A}_{\text{stable}}$  — Remark: Anticipated referee objection (observability vs. state variables)
8. Dimensional Consequence — Proposition 8.1: Dimensional reduction of the admissible state space when  $h < \ell_{\text{eff}}$  — Corollary 8.2: Microphysical anchoring of  $\Sigma^*$  and  $a_0$  from vertical equilibrium
9. Two-Stage Structure of the Derivation — Proposition 9.1: 2D Laplacian and logarithmic Green's function from symmetry
10. Relation to Assumption BP — Corollary 10.1: Assumption BP as a derived consequence
11. What This Paper Does Not Claim
12. Conclusion
13. Combined Implications of the Two Papers — §13.1: The complete logical chain (axioms  $\rightarrow a_0$ ) — §13.2: Microphysical anchoring of the acceleration scale — §13.3: What the combined result actually asserts — §13.4: What remains genuinely assumed (LC, EFT locality, symmetry) — §13.5: Relationship to MOND, modified gravity, dark matter, emergent gravity — §13.6: Falsifiability — the decision tree across SPARC tests —

§13.7: Open questions and future directions — §13.8: Structural dependency on companion paper's Lemma II (*three conditions; gap analysis*)

- Appendix — Relationship to Standard Physics — A.1: Gauge-invariant observables and LC — A.2: IC and the erasure of uncommitted fluctuations — A.3: EFT locality and suppression of non-local terms

---

## Abstract

We prove that any physically meaningful macroscopic law must be defined on the  $\sigma$ -algebra generated by stable record partitions under finite distinguishability, irreversible commitment, and law closure under record exchange (LC). The central result is a **representation theorem**: every admissible macroscopic observable is measurable with respect to the stable record  $\sigma$ -algebra  $\mathcal{A}_{\text{stable}}$ . Observer comparability — previously assumed as an independent axiom — is here derived as a consequence of LC + FD + IC (Proposition 2.1), reducing the axiomatic base from four independent postulates to three. From the representation theorem and its supporting lemmas (including a No-Resurrection lemma establishing that uncommitted fluctuations cannot accumulate into law-eligible channels even through coherent aggregation), we derive that the number of independent macroscopic state variables equals the number of independent stable partition channels. When geometric or dynamical constraints eliminate stable channels along a spatial direction, that direction cannot contribute a primitive macroscopic state variable, and its residual influence on remaining variables is representable as a local operator on the reduced algebra under standard EFT locality. This establishes **operational law closure** as a structural consequence of the axioms, not a free assumption. As a corollary, the Bridge Principle (Assumption BP) of the companion dimensional reduction paper follows from this framework. Combined with a microphysical anchoring of the record coherence length to the WNM thermal relaxation scale in outer galactic disks, the framework predicts the MOND acceleration scale  $a_0$  lies in the decade  $\sim 6 \times 10^{-11}$  to  $\sim 3 \times 10^{-10}$  m s<sup>-2</sup> — a band prediction derived from atomic physics inputs, not a calibration to observed MOND phenomenology.

---

## 1. Motivation and Scope

The companion paper (*Finite Distinguishability and Dimensional Reduction*) proves that a slab of physical thickness  $h < \ell_{\text{eff}}$  cannot support independent three-dimensional stable record partitions, and derives from this a surface-density threshold  $\Sigma^*$  and characteristic acceleration scale  $a_0$ . That result is geometric and structural. Its connection to effective field equations requires an additional assumption — **Assumption BP (Operational Closure)** — which states that the effective action governing macroscopic dynamics is defined on the algebra of admissible stable records, and that integrating out non-representable degrees of freedom yields a reduced effective theory.

The purpose of the present paper is to show that Assumption BP is not an independent postulate. It follows from a coherent and independently motivated account of **macroscopic lawhood** — specifically, from the requirement that laws support counterfactual reuse and observer comparability under finite distinguishability and irreversible commitment.

**Scope clarification.** This paper does not claim that unstable degrees of freedom cannot *influence* macroscopic dynamics. A 3D field can exist microscopically and mediate forces even if vertical structure is not independently recordable. What is claimed is structural: unstable degrees of freedom cannot contribute *independent* macroscopic state variables. Their influence on stable variables is not eliminated; it is absorbed into the effective couplings of the reduced theory, in exact analogy with standard Wilsonian integration of heavy modes. The claim is about independence, not influence. This distinction is the paper's central contribution.

---

## 2. Axiomatic Framework

We work with the same minimal admissibility structure as the companion paper, with three axioms:

---

**Axiom FD — Finite Distinguishability.** There exists a minimum resolvable partition scale  $v_0 > 0$  such that no physically realisable record partition resolves cells of measure smaller than  $v_0$ .

---

**Axiom IC — Irreversible Commitment.** Let  $\mathcal{O}(R, \varepsilon)$  denote the class of admissible physical processes — those implementable with resources  $R$  and failure probability less than  $\varepsilon$ . A macroscopic distinction is *committed* if it is non-recombinable within  $\mathcal{O}(R, \varepsilon)$ : once written, it cannot be erased by any admissible process. Committed distinctions accumulate across readout intervals; uncommitted distinctions do not.

*Remark.* IC is strictly stronger than FD. FD requires that distinctions cost a minimum resource to write. IC additionally requires that certain distinctions, once written, cannot be undone within  $\mathcal{O}(R, \varepsilon)$ . The two axioms are independent: a system can satisfy FD without any commitment (if all partitions are reversible) and can satisfy IC without FD (if irreversible distinctions exist but have no minimum scale). Together they define the class of *stable committed records* — the only distinctions that can accumulate into macroscopic differences over operational timescales.

---

**Principle PCP (Prediction Composition Principle).** If two admissible experimental procedures  $P_1, P_2 \in \mathcal{O}(R, \varepsilon)$  are *operationally equivalent* — i.e., they produce the same ensemble of outcomes on all admissible inputs within  $\mathcal{O}(R, \varepsilon)$  — then any law statement evaluated on the outputs of  $P_1$  must yield the same truth value as when evaluated on the outputs of  $P_2$ .

*Remark.* PCP is not a claim about laws. It is a consistency condition on what "prediction" means within  $\mathcal{O}(\mathcal{R}, \varepsilon)$ : operationally equivalent procedures are, by definition, the same procedure at the operational level, so they cannot yield different predictions without rendering the prediction ill-defined. Any framework that accepts that experimental procedures have definite predictions — not merely procedure-relative ones — accepts PCP. It is prior to and more primitive than any claim about representation invariance.

**Definition 2.1 (Admissible Representation Change).** An *admissible representation change* is any invertible map  $U \in \mathcal{O}(\mathcal{R}, \varepsilon)$  acting on the record description of an observer — change of measurement basis, pointer relabelling, coarse-graining dictionary change — that preserves admissibility:  $U$  maps  $\mathcal{O}(\mathcal{R}, \varepsilon)$ -implementable procedures to  $\mathcal{O}(\mathcal{R}, \varepsilon)$ -implementable procedures, and  $U^{-1} \in \mathcal{O}(\mathcal{R}, \varepsilon)$ .

*Remark.* Admissible representation changes are operationally equivalent to the identity on physical outcomes: the physical ensemble that reaches an observer is the same before and after  $U$ ; only the record description changes. By PCP, therefore, any prediction computed from the  $U$ -transformed description must agree with the prediction computed from the original description.

**Axiom LI — Law Invariance Under Admissible Representation Change.** (*Corollary of PCP + closure of  $\mathcal{O}(\mathcal{R}, \varepsilon)$  under invertible transformations.*)

For any law statement  $L$  and any admissible representation change  $U \in \mathcal{O}(\mathcal{R}, \varepsilon)$ :

$$L(\rho) = L(U\rho) \quad \text{for all admissible ensembles } \rho,$$

whenever the records being compared are compatible (i.e., the observers' stable partition states agree on  $\mathcal{A}_1 \cap \mathcal{A}_2$ ).

*Derivation from PCP.* An admissible representation change  $U$  is a procedure in  $\mathcal{O}(\mathcal{R}, \varepsilon)$  that acts only on the record description, leaving the physical ensemble unchanged. The procedures "measure with description  $\rho$ " and "measure with description  $U\rho$ " are therefore operationally equivalent within  $\mathcal{O}(\mathcal{R}, \varepsilon)$ : they produce the same physical outcomes. By PCP,  $L$  evaluated on  $\rho$  and  $L$  evaluated on  $U\rho$  must agree. This is exactly  $L(\rho) = L(U\rho)$ . LI is therefore not an additional axiom — it is the statement that law predictions respect PCP over the symmetry group generated by admissible representation changes.  $\square$

*Remark.* The logical order is now: PCP (definitional of prediction)  $\rightarrow$  LI (symmetry group of  $\mathcal{O}(\mathcal{R}, \varepsilon)$  on record descriptions)  $\rightarrow$  Theorem 2.LC (LC as invariant-subalgebra result)  $\rightarrow$  Proposition 2.1. LI is not a strategic postulate encoding the desired conclusion. It is what PCP forces on law statements when the symmetry group is the invertible maps in  $\mathcal{O}(\mathcal{R}, \varepsilon)$ .

**Theorem 2.LC (LI implies Law Closure — representation theorem).** *If LI holds, then every law statement  $L$ , being invariant under all admissible representation changes  $U \in \mathcal{O}(\mathcal{R}, \varepsilon)$ , must be evaluable on the invariant subalgebra under those transformations. For compatible observers  $O_1, O_2$ , this invariant subalgebra is the intersection  $\mathcal{A}_1 \cap \mathcal{A}_2$ , and the consistent refinement  $\pi$  required by LC exists within it. That is, LC holds as a derived representation theorem under LI.\**

**Proof.** Let  $L$  be a law statement and  $U$  an admissible representation change. By LI,  $L(\rho) = L(U\rho)$  for all admissible  $\rho$ . This means  $L$  is constant on  $U$ -orbits in the space of admissible ensembles —  $L$  depends only on  $U$ -invariant features of  $\rho$ . The class of all admissible representation changes generates an equivalence on record descriptions; the invariant subalgebra is the algebra of events whose membership is unchanged by any  $U \in \mathcal{O}(\mathbb{R}, \varepsilon)$ .

For compatible observers  $O_1, O_2$  with algebras  $\mathcal{A}_1, \mathcal{A}_2$ : any event in  $\mathcal{A}_1 \setminus \mathcal{A}_2$  can be relabelled by an admissible  $U$  (acting on  $O_1$ 's record description) to produce a different event description for the same underlying microstate, while leaving  $O_2$ 's description unchanged. By LI,  $L$  must be invariant under this relabelling. Therefore  $L$  cannot depend on events in  $\mathcal{A}_1 \setminus \mathcal{A}_2$  — its domain is restricted to events that survive the  $U$ -symmetry, which is  $\mathcal{A}_1 \cap \mathcal{A}_2$ . A partition  $\pi^* \in \mathcal{A}_1 \cap \mathcal{A}_2$  that refines both  $\pi_1$  and  $\pi_2$  and supports joint evaluation of  $L$  exists precisely because  $L$ 's domain has already been restricted to the intersection. The two failure modes of the earlier argument (no common domain; contextual truth) are both excluded: no-common-domain fails because  $L$  is defined on  $\mathcal{A}_1 \cap \mathcal{A}_2$  by the symmetry argument; contextual truth fails because LI forces  $L(\rho) = L(U\rho)$ , so the truth value cannot depend on which coupling generated the records.  $\square$

*Remark.* The logical structure is: LI is a symmetry principle  $\rightarrow$  law statements are  $U$ -invariant  $\rightarrow$  they factor through the invariant subalgebra  $\mathcal{A}_1 \cap \mathcal{A}_2 \rightarrow$  LC (consistent refinement in the intersection) is the required domain. This is the exact pattern of gauge invariance: gauge symmetry  $\rightarrow$  physical observables are gauge-invariant  $\rightarrow$  they live on the gauge-invariant algebra. LC is the event-algebraic analogue of "gauge-invariant observables are the primitive objects of law statements."

*Definition (Consistent Refinement).* For the statement of Theorem 2.LC and subsequent proofs, a *consistent refinement* of partitions  $\pi_1 \in \mathcal{A}_1$  and  $\pi_2 \in \mathcal{A}_2$  is a partition  $\pi^*$  satisfying:

(a)  $\sigma(\pi^*) \supseteq \sigma(\pi_1)$  and  $\sigma(\pi^*) \supseteq \sigma(\pi_2)$  ( $\pi$  refines both)\*

(b)  $\pi^* \in \mathcal{A}_1 \cap \mathcal{A}_2$  ( $\pi$  lies in the invariant subalgebra)\*

(c) For any law statement  $L$ , the truth value of  $L$  on  $\pi^*$ -equivalence classes is the same whether derived from  $O_1$ 's records or  $O_2$ 's.

Condition (b) is the load-bearing clause, forced by the symmetry argument of Theorem 2.LC: the invariant subalgebra is exactly  $\mathcal{A}_1 \cap \mathcal{A}_2$ , and  $\pi^*$  must live there.

*Remark on LC as derived rather than assumed.* In previous versions of this programme, Axiom LC was assumed directly. Theorem 2.LC shows it is a representation theorem: given that law statements are invariant under admissible representation changes (LI), they must be evaluable on the invariant subalgebra, which is  $\mathcal{A}_1 \cap \mathcal{A}_2$ . LC is what the symmetry requirement forces on the event domain of law statements. The programme now rests on three axioms (FD, IC, LI) rather than the earlier LC, and LC appears as a structural consequence of the symmetry group of  $\mathcal{O}(\mathbb{R}, \varepsilon)$ .

**Why LI is not optional.** If LI fails, there exists an admissible representation change  $U$  such that  $L(\rho) \neq L(U\rho)$  for some law statement  $L$  and admissible  $\rho$ . Then two observers who differ only in their representation choice — both within  $\mathcal{O}(R, \varepsilon)$  — disagree on the truth value of  $L$ . "The law" is not a single object but a family indexed by admissible measurement couplings  $M$ . LI is the minimal symmetry requirement that prevents this: it is what distinguishes a physical law from a coupling-dependent regularisation artifact.

**Proposition 2.0 (Contextual Collapse under LI Failure).** *If LI fails, then there exist admissible observers  $O_1, O_2$  and a law statement  $L$  such that the truth value of  $L$  is not well-defined under record exchange.*

*Proof.* If LI fails, there exists  $U \in \mathcal{O}(R, \varepsilon)$  such that  $L(\rho) \neq L(U\rho)$  for some  $\rho$ . Take  $O_1$  using description  $\rho$  and  $O_2$  using description  $U\rho$  — both admissible. Then  $L$  has different truth values under the two representations. The standard contextual- pointer construction realises this: take  $\mathcal{A}_1$  and  $\mathcal{A}_2$  as the stable algebras of two complementary von Neumann projectors; their intersection is trivial, so any non-trivial law statement has no  $U$ -invariant domain on which to be jointly evaluated.  $\square$

*Remark.* The paper's logical spine is therefore:  $FD + IC + LI \rightarrow LC$  (Theorem 2.LC, representation theorem)  $\rightarrow$  Proposition 2.1  $\rightarrow$  Theorem 5.1  $\rightarrow$  Lemma 6.1  $\rightarrow$  Theorem 7.1.

**Operational Metric  $\mathcal{D}_{\mathcal{O}}$ .** To make LC precise and to formalise "stable committed distinction" in a metric compatible with the channel structure of  $\mathcal{O}(R, \varepsilon)$ , we introduce the *admissible distinguishability* at this point, since it is used both in the statement of LC and in later proofs.

For two ensembles  $\rho, \rho'$  of microstates over  $\Omega$ , the *admissible distinguishability at tolerance  $\varepsilon$*  is:

$$\mathcal{D}_{\mathcal{O}}(\rho, \rho') \equiv \sup_{\{\Pi \in \mathcal{O}(R, \varepsilon)\}} \text{TV}(P_{\Pi}(\rho), P_{\Pi}(\rho'))$$

where the supremum is over all measurement procedures  $\Pi$  implementable within  $\mathcal{O}(R, \varepsilon)$ , and TV denotes total variation distance. A distinction is *committed* over  $[0, \tau]$  if  $\mathcal{D}_{\mathcal{O}}(\rho_t, \rho'_t) > \varepsilon$  for all  $t \in [0, \tau]$ ; it is *uncommitted* if  $\mathcal{D}_{\mathcal{O}}(\rho_t, \rho'_t) \leq \varepsilon$  for some  $t \leq \tau$ . Properties of  $\mathcal{D}_{\mathcal{O}}$ , including its monotonicity under admissible channels, are established in §5.

**Proposition 2.1 (LC + FD + IC imply Observer Comparability).** Under Axioms FD, IC, and LI (via Theorem 2.LC), any primitive macroscopic state variable that can appear in a law statement satisfying LC must correspond to a distinction that is (i) stably committed under all admissible observers' measurement couplings and (ii) invariant under admissible record exchange. Equivalently:

Primitive law variables must lie in  $\mathcal{A}_1 \cap \mathcal{A}_2$  for every pair of admissible observers  $O_1, O_2$  — i.e., they must be measurable with respect to the intersection  $\bigcap_i \mathcal{A}_i$  over all admissible observers, which equals  $\mathcal{A}_{\text{stable}}$ .

**Proof.** Let  $X$  be a primitive variable in some admissible law  $L$ . We show  $X \in \mathcal{A}_{\text{stable}}$ .

*Step 1 — X must lie in both observers' stable algebras.* Let  $O_1, O_2$  be any two admissible observers who have exchanged records and whose aggregate stable partition states are compatible (i.e., their records agree on all events in  $\mathcal{A}_1 \cap \mathcal{A}_2$ ). By LC, there exists a consistent refinement  $\pi^* \in \mathcal{A}_1 \cap \mathcal{A}_2$  such that the truth value of  $L$  — and hence the value of  $X$  — is representation-independent on  $\pi^*$ -equivalence classes. For  $X$  to have a well-defined value on  $\pi^*$ -equivalence classes,  $X$  must be measurable with respect to  $\sigma(\pi^*)$ . Since  $\pi^* \in \mathcal{A}_1 \cap \mathcal{A}_2$ , we have  $\sigma(\pi^*) \subseteq \mathcal{A}_1 \cap \mathcal{A}_2$ , so  $X$  must lie in  $\mathcal{A}_1 \cap \mathcal{A}_2$  for this pair.

*Step 2 — Asymmetric stability does not permit X to escape.* Suppose  $X$  is stable under  $O_1$ 's coupling but unstable under  $O_2$ 's. Then  $X \in \mathcal{A}_1$  but  $X \notin \mathcal{A}_2$ . The consistent refinement  $\pi^*$  requires  $\pi^* \in \mathcal{A}_1 \cap \mathcal{A}_2$  (condition (b) of the consistent refinement definition). Since  $X \notin \mathcal{A}_2$ ,  $X$  cannot be measurable with respect to any event in  $\mathcal{A}_2$ , and hence not with respect to any  $\pi^* \in \mathcal{A}_1 \cap \mathcal{A}_2$ . Therefore  $X$  cannot appear as a primitive variable in any law statement satisfying LC: there is no consistent refinement on which  $X$  is representation-independent for this observer pair.

*Step 3 — IC eliminates uncommitted distinctions.* Suppose  $X$  is stable under all admissible observers' couplings over some subinterval of  $[0, \tau]$  but becomes uncommitted — that is,  $\mathcal{D}_O(\rho_{t^*}, \rho_{t^*}') \leq \varepsilon$  at some  $t^* \leq \tau$ . By IC, committed distinctions are exactly those that remain above  $\varepsilon$  in  $\mathcal{D}_O$  throughout  $[0, \tau]$ ; any distinction that falls to  $\leq \varepsilon$  at  $t^*$  is by definition uncommitted and cannot accumulate as an independent constraint channel. (This is the content of IC stated in metric form via  $\mathcal{D}_O$ ; the formal derivation requiring only FD + IC appears as Lemma 5.2 in §5 and does not depend on this proposition.) Therefore  $X$  cannot support a stably committed event at all  $t \in [0, \tau]$ , and  $X \notin \mathcal{A}_{\text{stable}}$ .

*Step 4 — The intersection is  $\mathcal{A}_{\text{stable}}$ .* From Steps 1–3: any primitive variable  $X$  in a law satisfying LC must lie in  $\mathcal{A}_i$  for every admissible observer  $O_i$  and must correspond to a stably committed distinction over  $[0, \tau]$ . The intersection  $\bigcap_i \mathcal{A}_i$  over all admissible observers consists exactly of the events that are stably committed under every admissible coupling — which is by definition  $\mathcal{A}_{\text{stable}}$ . Therefore  $X \in \mathcal{A}_{\text{stable}}$ .  $\square$

*Remark.* Proposition 2.1 is what was previously called Axiom OC. It is now a *derived* consequence of LC + FD + IC rather than an independent postulate. The proof handles the asymmetric case (Step 2) and the IC-closure case (Step 3) that the earlier sketch did not address. The key technical move is that "consistent refinement" requires the witnessing partition to lie in  $\mathcal{A}_1 \cap \mathcal{A}_2$  (condition (b) of the LC definition), which forces  $X$  into both observers' algebras simultaneously. All subsequent references to "OC" in this paper refer to the content of Proposition 2.1, which is now a theorem rather than an axiom. LC allows auxiliary, gauge-dependent, or unobservable variables in intermediate calculations; it restricts only primitive law statements to representation-independent distinctions.

---

### 3. Stable Record Algebra

Let  $(\Omega, \mathcal{S}, \mu)$  be a measurable microstate space with reversible dynamics  $\Phi_t : \Omega \rightarrow \Omega$  preserving  $\mu$  (Liouville property). Fix a readout duration  $\tau > 0$  and failure threshold  $\varepsilon \in (0, 1/2)$ .

**Definition 3.1 (Stable Record Partition).** A partition  $\pi$  of  $\Omega$  into cells  $\{C_i\}$  is a *stable record partition* if:

- (i)  $\mu(C_i) \geq v_0$  for all  $i$ , [Axiom FD]
- (ii) *Stability over  $\tau$* :  $\text{Prob}[\Phi_t(\omega) \notin C_i \mid \omega \in C_i] < \varepsilon$  for all  $t \in [0, \tau]$  and  $\mu$ -almost all  $\omega$ . [Non-mixing condition]

A stable record partition is *committed* if, in addition, no admissible process in  $\mathcal{O}(R, \varepsilon)$  can invert the readout equivalence class over  $\tau$  — that is, no  $\mathcal{E} \in \mathcal{O}(R, \varepsilon)$  maps distinguishable cell-states back to indistinguishable ones within  $[0, \tau]$ . Commitment is formalised via  $\mathcal{D}_\tau$  (introduced in §2); stability (condition ii) is necessary but not sufficient for commitment.

*Choice of  $\tau$  (readout horizon).* The stability horizon  $\tau$  is not a free parameter; it is fixed by the intended law domain. A macroscopic law is a reusable intervention–outcome relation, so  $\tau$  must exceed the minimum time between independent admissible readouts required to test that relation (counterfactual reuse), while remaining short compared to the timescale on which the background state evolves. In the galactic disk application,  $\tau$  lies between the local thermal or phase relaxation time of the record-bearing medium (below which distinct record basins cannot be stably distinguished) and the orbital or structural evolution time (above which the slab geometry changes). Specifically:  $\tau$  must be long enough for WNM thermal basins to re-establish after perturbation — of order  $\tau_{\text{cool}} \sim 5 \times 10^6$  yr — yet short compared to secular disk evolution timescales of order Gyr. Any  $\tau$  in this window gives the same class of stable record partitions; the results of §5–7 are  $\tau$ -independent within this window.

**Definition 3.2 (Stable Record  $\sigma$ -Algebra).** The *stable record  $\sigma$ -algebra* is:

$$\mathcal{A}_{\text{stable}} = \sigma(\{C_i\} : \pi \text{ is a stable record partition})$$

This is the smallest  $\sigma$ -algebra containing all stable partition cells. It is the algebra of macroscopic distinctions admissible under FD + IC.

**Definition 3.3 (Stable Equivalence).** Two microstates  $\omega, \omega' \in \Omega$  are *stably equivalent* ( $\omega \sim_{\text{stable}} \omega'$ ) if they belong to the same cell of every stable record partition. The *macroscopic state space* is the quotient:

$$\mathbb{M}_{\text{macro}} = \Omega / \sim_{\text{stable}}$$

**Definition 3.4 (Record Coherence Length  $\ell_{\text{eff}}$ ).** The *record coherence length*  $\ell_{\text{eff}}$  is defined in terms of the record-theoretic objects already introduced, without reference to any specific physical model:

$\ell_{\text{eff}} := \inf\{r > 0 : \exists \text{ two spatial regions } R_1, R_2 \text{ with } |R_1 - R_2| = r \text{ such that the restriction of } \mathcal{A}_{\text{stable}} \text{ to } (R_1, R_2) \text{ jointly supports a committed distinction with } \mathcal{D}_{\mathcal{O}} > \varepsilon \text{ over } [0, \tau]\}$

Equivalently,  $\ell_{\text{eff}}$  is the smallest separation at which two regions can independently sustain committed distinguishable record states over the readout horizon  $\tau$ . Below  $\ell_{\text{eff}}$ , any distinction between the two regions fails commitment and falls below  $\varepsilon$  in  $\mathcal{D}_{\mathcal{O}}$  before  $\tau$ .

**Definition 3.5 (Commitment Time  $\tau_{\text{commit}}$ ).** The *commitment time* of a local dynamical process is defined intrinsically in terms of  $\mathcal{D}_{\mathcal{O}}$ :

$$\tau_{\text{commit}} := \inf\{t \geq 0 : \mathcal{D}_{\mathcal{O}}(\rho_t, \rho'_t) > \varepsilon \text{ and } \mathcal{D}_{\mathcal{O}}(\rho_s, \rho'_s) > \varepsilon \text{ for all } s \in [t, t + \tau]\}$$

That is,  $\tau_{\text{commit}}$  is the first time at which  $\mathcal{D}_{\mathcal{O}}$  crosses above  $\varepsilon$  and *remains* above  $\varepsilon$  for a full readout horizon  $\tau$  afterward. Before  $\tau_{\text{commit}}$ , the distinction between  $\rho$  and  $\rho'$  has not yet been committed in the sense of Definition 3.1. After  $\tau_{\text{commit}}$ , it is committed and, by Lemma 5.2b (monotonicity), cannot be erased by any admissible operation.

*Remark.*  $\tau_{\text{commit}}$  is a property of the dynamics, not a physical clock. It depends on the pair  $(\rho, \rho')$  and the channel structure of  $\mathcal{O}(R, \varepsilon)$  — not on any specific microphysical timescale such as cooling time. The identification  $\tau_{\text{commit}} \approx \tau_{\text{cool}}$  in the WNM regime is then an empirical claim: in that physical context, the dominant process driving  $\mathcal{D}_{\mathcal{O}}$  from below  $\varepsilon$  to stably above  $\varepsilon$  is thermal relaxation, and it does so on the cooling timescale. Cooling time is the physical mechanism that *realises*  $\tau_{\text{commit}}$  in the WNM regime — it is not  $\tau_{\text{commit}}$  by definition. Alternative mechanisms (turbulent cascade, phase transition, GMC formation) realize different  $\tau_{\text{commit}}$  values and accordingly different  $\ell_{\text{eff}}$ , which is the content of the H2 weak universality scenario (§13.4).

**Theorem 3.1 (Mixing-Commitment Competition Determines  $\ell_{\text{eff}}$ ).** Let  $\tau_{\text{mix}}(\ell)$  denote the time for coupling or mixing processes to erase a distinction of spatial extent  $\ell$  within  $\mathcal{O}(R, \varepsilon)$ , and let  $\tau_{\text{commit}}$  denote the time for the local dynamics to commit a distinction to a stable basin. Then:

$\ell_{\text{eff}}$  is the solution of  $\tau_{\text{mix}}(\ell_{\text{eff}}) = \tau_{\text{commit}}$

*Below this scale, mixing outruns commitment and no stable distinct basins persist. Above this scale, commitment outruns mixing and distinct stable record states can coexist independently.*

**Proof sketch.** Let  $\ell < \tau_{\text{mix}}^{-1}(\tau_{\text{commit}})$  — that is,  $\tau_{\text{mix}}(\ell) < \tau_{\text{commit}}$ . At this separation, mixing equilibrates two candidate regions before their local dynamics can commit each to a stable basin. Operationally: mixing drives  $\mathcal{D}_{\mathcal{O}}(\rho_t, \rho'_t)$  below  $\varepsilon$  on timescale  $\tau_{\text{mix}}(\ell) < \tau_{\text{commit}}$ . Once  $\mathcal{D}_{\mathcal{O}} \leq \varepsilon$ , the distinction is uncommitted by Definition 3.1, and by the monotonicity of  $\mathcal{D}_{\mathcal{O}}$  under admissible channels — which follows from closure of  $\mathcal{O}(R, \varepsilon)$  under sequential composition (established formally as Lemma 5.2b in §5, where the full proof is given; the property is used here only as a premise) — no subsequent admissible operation can restore  $\mathcal{D}_{\mathcal{O}}$  above  $\varepsilon$ . The distinction therefore cannot contribute to  $\mathcal{A}_{\text{stable}}$  at scale  $\ell$ . Conversely, for  $\ell > \tau_{\text{mix}}^{-1}(\tau_{\text{commit}})$ , mixing requires more than  $\tau_{\text{commit}}$  to equilibrate the two regions, so

the local dynamics has time to commit each region independently before mixing intervenes. Both regions then sustain  $\mathcal{D}_O > \varepsilon$  over  $[0, \tau]$  and constitute independent committed record states. The infimum of the second regime is  $\ell_{\text{eff}}$ .  $\square$

*Remark (derivation for different mixing channels).* Theorem 3.1 gives  $\ell_{\text{eff}}$  as a derived quantity in all cases once the mixing channel is specified:

- **Ballistic mixing** (signal propagation at speed  $v_{\text{mix}}$ ):  $\tau_{\text{mix}}(\ell) \sim \ell/v_{\text{mix}}$ , so  $\ell_{\text{eff}} \sim v_{\text{mix}} \cdot \tau_{\text{commit}}$ . For WNM with  $v_{\text{mix}} = c_s$  and  $\tau_{\text{commit}} = \tau_{\text{cool}}$ :  $\ell_{\text{eff}} \sim c_s \cdot \tau_{\text{cool}} \sim 20\text{--}100$  pc. (*Primary channel: §8 Corollary 8.2.*)
- **Diffusive mixing** (thermal conduction with diffusivity  $\chi$ ):  $\tau_{\text{mix}}(\ell) \sim \ell^2/\chi$ , so  $\ell_{\text{eff}} \sim \sqrt{\chi \cdot \tau_{\text{commit}}}$ . For WNM at  $T \sim 8000$  K this recovers the Field length  $\lambda_F \sim \sqrt{(\kappa T)/(n^2 \Lambda(T))}$  — the scale below which conduction prevents distinct thermal phases from coexisting. This is a more conservative (smaller) lower bound on  $\ell_{\text{eff}}$  than the ballistic estimate, appropriate when conduction is efficient.
- **Turbulent mixing**:  $\tau_{\text{mix}}(\ell) \sim \ell/v_{\text{turb}}(\ell)$  (using the turbulent velocity at scale  $\ell$ ), giving an implicit equation for  $\ell_{\text{eff}}$  depending on the turbulent power spectrum.

In all cases  $\ell_{\text{eff}}$  is computed from physical inputs — not postulated. The choice of primary channel (ballistic WNM thermal relaxation) is stated as an explicit assumption in §13.4 and testable via weak universality predictions (H2 scenario).

**Lemma 3.2 (Dominant Mixing Channel Determines  $\ell_{\text{eff}}$ ).** *Let  $\{\tau_{\text{mix}}^{\{k\}}(\ell)\}$  be the set of mixing timescales for all admissible mixing channels  $k$  at scale  $\ell$  (ballistic, diffusive, turbulent, magnetic, etc.). Define:*

$$\tau_{\text{mix}}(\ell) := \min_k \tau_{\text{mix}}^{\{k\}}(\ell)$$

*Then  $\ell_{\text{eff}}$  is determined by the fastest channel alone:*

$$\ell_{\text{eff}} = \inf\{\ell : \min_k \tau_{\text{mix}}^{\{k\}}(\ell) \geq \tau_{\text{commit}}\}$$

**Proof.** A distinction at scale  $\ell$  is committed if and only if it survives all admissible mixing processes over  $[0, \tau]$ . Survival requires  $\tau_{\text{mix}}^{\{k\}}(\ell) \geq \tau_{\text{commit}}$  for every channel  $k$  simultaneously — any single channel with  $\tau_{\text{mix}}^{\{k\}}(\ell) < \tau_{\text{commit}}$  is sufficient to drive  $\mathcal{D}_O$  below  $\varepsilon$  before commitment. Therefore the binding constraint is the minimum over all channels.  $\ell_{\text{eff}}$  is the smallest scale at which even the fastest admissible mixing channel fails to erase the distinction within  $\tau_{\text{commit}}$ .  $\square$

*Remark.* Lemma 3.2 removes the apparent "choice" of mixing channel from the framework. The physical world does not choose ballistic vs diffusive — it runs all mixing channels simultaneously, and the fastest one governs.  $\ell_{\text{eff}}$  is determined by whichever process erases distinctions most rapidly at the relevant scale. In WNM outer disks, ballistic (sound-speed) propagation dominates over conduction and turbulence at 20–100 pc scales (turbulent velocities are subsonic at these scales; conduction is slower than sound-speed propagation in the weakly-ionised WNM). When a different channel dominates — as in turbulence-dominated inner disks

or magnetised environments —  $\ell_{\text{eff}}$  shifts accordingly, and the prediction for  $a_0$  shifts with it. This is the physical content of the H2 weak universality scenario: different dominant channels  $\rightarrow$  different  $\ell_{\text{eff}} \rightarrow$  different  $a_0$ , with  $a_0 \propto n \Lambda(T) / T^{1/2}$  replaced by a channel-dependent scaling.

## 4. Admissible Macroscopic Observables

**Definition 4.1 (Macroscopic Observable).** A function  $f : \Omega \rightarrow \mathbb{R}$  is a *macroscopic observable* if it is:

(i) *Dynamically stable*:  $|f(\Phi_t(\omega)) - f(\omega)| < \delta$  for all  $t \in [0, \tau]$  and  $\mu$ -almost all  $\omega$ , for some  $\delta > 0$ .

(ii) *Partition-resolvable*: there exists a stable record partition  $\pi$  such that  $f$  is constant on cells of  $\pi$  up to fluctuations of order  $\delta$ .

(iii) *Intervention-responsive*: for some admissible intervention class  $\text{do}(X = x)$ ,  $\text{do}(X = x')$ , the counterfactual condition of Definition 4.2 is satisfied.

(iv) *Law-eligible*:  $f$  is intended to appear as a primitive state variable or outcome in an admissible law statement. Equivalently, its counterfactual dependence claims must be testable by stable partitions within  $\mathcal{O}(\mathbb{R}, \varepsilon)$ .

*Remark.* Condition (iv) is what separates macroscopic observables from arbitrary stable averages. A quantity such as an ergodic time-average or a hydrodynamic expectation value may be dynamically stable (condition i) and even partition-resolvable (condition ii) without being law-eligible: if its variation cannot be linked to admissible interventions testable by stable partitions, it cannot appear as a primitive state variable in a law satisfying OC. Theorem 5.1 is a representation theorem for the class defined by all four conditions — not for all conceivable stable functions on  $\Omega$ .

**Definition 4.2 (Counterfactual Stability Condition).** A macroscopic observable  $f$  satisfies the *counterfactual stability condition* with respect to variables  $X, Y$  if: for any admissible intervention  $\text{do}(X = x)$  implementable within  $\mathcal{O}(\mathbb{R}, \varepsilon)$ , there exists a stable partition  $\pi$  such that

$$\text{TV}(\mathbb{P}_{\pi}(f(Y) | \text{do}(X = x)), \mathbb{P}_{\pi}(f(Y) | \text{do}(X = x'))) > \eta$$

for some  $\eta > 0$ , and this inequality is preserved under admissible coarse-graining of  $\pi$ .

Here TV denotes total variation distance, and the interventions  $\text{do}(X = x)$  must themselves be implementable within  $\mathcal{O}(\mathbb{R}, \varepsilon)$  — interventions on unstable variables are not admissible, since they cannot be reliably implemented or compared across observers.

*Remark.* The requirement that interventions lie within  $\mathcal{O}(\mathbb{R}, \varepsilon)$  is essential. Without it, counterfactuals on unstable variables would be formally definable but operationally meaningless

— the intervention could not be reproducibly realised or compared between observers, violating OC.

---

## 5. The Representation Theorem

**Theorem 5.1 (Representation of Law-Eligible Observables).** *Under Axioms FD, IC, and LI (via Theorem 2.LC) (equivalently, FD + IC + the observer-comparability consequence Proposition 2.1), every observable that is law-eligible in the sense of Definition 4.1(iv) is measurable with respect to  $\mathcal{A}$ -stable.*

*Remark on scope.* This is a representation theorem for the class defined by all four conditions of Definition 4.1, not for all conceivable stable functions on  $\Omega$ . An arbitrary stable average (e.g., an ergodic time-average or hydrodynamic expectation) may satisfy conditions (i)–(ii) without satisfying (iii)–(iv). Theorem 5.1 applies only to the law-eligible subclass — observables whose counterfactual dependence claims are testable by stable partitions within  $\mathcal{O}(\mathbb{R}, \varepsilon)$ . The nontrivial content is not the measurability statement itself (which follows almost by definition for condition-ii observables) but the combination of Lemma 5.2 and Lemma 5.2c, which establish that no uncommitted fluctuation can accumulate into a law-eligible distinction even through aggregation or coherent superposition.

**Proof.** Let  $\omega \sim_{\text{stable}} \omega'$ . Then  $\omega$  and  $\omega'$  belong to the same cell of every stable record partition. By Definition 4.1(ii),  $f$  is constant on cells of some stable partition  $\pi$  up to fluctuations  $\delta$ . Since  $\omega$  and  $\omega'$  are in the same cell of  $\pi$ ,  $|f(\omega) - f(\omega')| < \delta$ .

It remains to show that any residual intra-cell variation is law-irrelevant: it cannot support counterfactual reuse under admissible interventions and readouts within  $\mathcal{O}(\mathbb{R}, \varepsilon)$ , and therefore cannot define a primitive macroscopic distinction in the sense of Definition 4.1(iv). By Lemma 5.2, any fluctuation uncommitted under IC falls below  $\varepsilon$  in  $\mathcal{D}_{\mathcal{O}}$  at some  $t^* \leq \tau$  and remains below  $\varepsilon$  under admissible processing (Lemma 5.2b). By Lemma 5.2c, no later aggregation within  $\mathcal{O}(\mathbb{R}, \varepsilon)$  can resurrect it as an independent channel — including via non-law-eligible aggregates such as ergodic averages or hydrodynamic moments that might feed back into law-eligible variables. Therefore the intra-cell variation cannot constitute a law-eligible distinction under Definition 4.1(iv), and  $f$  is  $\mathcal{A}$ -stable-measurable.  $\square$

*Remark.* The role of IC is essential here. FD alone would give measurability up to coarse-graining, but could leave room for fine-scale fluctuations to persist as silent influences. IC closes this: uncommitted fluctuations are operationally erased within  $\tau$  and cannot accumulate into law-relevant distinctions. This is what distinguishes "cannot be resolved" (epistemic) from "cannot accumulate as an independent constraint channel" (structural).

---

**Definition 5.2a (Admissible Distinguishability — restated for reference).** As introduced in §2,  $\mathcal{D}_\mathcal{O}(\rho, \rho') \equiv \sup_{\Pi \in \mathcal{O}(R, \varepsilon)} \text{TV}(P_{\Pi}(\rho), P_{\Pi}(\rho'))$ . A distinction is committed over  $[0, \tau]$  if  $\mathcal{D}_\mathcal{O}(\rho_t, \rho'_t) > \varepsilon$  for all  $t \in [0, \tau]$ ; uncommitted if  $\mathcal{D}_\mathcal{O}(\rho_t, \rho'_t) \leq \varepsilon$  for some  $t \leq \tau$ .

**Lemma 5.2b (Monotonicity of  $\mathcal{D}_\mathcal{O}$  under Admissible Channels).** For any admissible channel  $\mathcal{E} \in \mathcal{O}(R, \varepsilon)$ :

$$\mathcal{D}_\mathcal{O}(\mathcal{E}\rho, \mathcal{E}\rho') \leq \mathcal{D}_\mathcal{O}(\rho, \rho')$$

**Proof.** For any measurement  $\Pi \in \mathcal{O}(R, \varepsilon)$ , the composed operation  $\Pi \circ \mathcal{E}$  is also in  $\mathcal{O}(R, \varepsilon)$  (since  $\mathcal{O}$  is closed under sequential composition within resource bounds). Therefore:

$$\text{TV}(P_{\Pi}(\mathcal{E}\rho), P_{\Pi}(\mathcal{E}\rho')) = \text{TV}(P_{\{\Pi \circ \mathcal{E}\}}(\rho), P_{\{\Pi \circ \mathcal{E}\}}(\rho')) \leq \sup_{\Pi' \in \mathcal{O}(R, \varepsilon)} \text{TV}(P_{\{\Pi'\}}(\rho), P_{\{\Pi'\}}(\rho')) = \mathcal{D}_\mathcal{O}(\rho, \rho')$$

Taking the supremum over  $\Pi$  on the left gives  $\mathcal{D}_\mathcal{O}(\mathcal{E}\rho, \mathcal{E}\rho') \leq \mathcal{D}_\mathcal{O}(\rho, \rho')$ .  $\square$

*Remark.* This is the data-processing inequality stated in the language of  $\mathcal{O}(R, \varepsilon)$ . It ensures that "committed" is a stable property under admissible processing: if a distinction is committed ( $\mathcal{D}_\mathcal{O} > \varepsilon$ ) at time  $t$ , no admissible channel applied subsequently can raise  $\mathcal{D}_\mathcal{O}$  above a value it has already fallen below, so uncommitted distinctions remain operationally inaccessible.

**Lemma 5.2c (No Resurrection by Aggregation).** Let  $\rho, \rho'$  be ensembles satisfying  $\mathcal{D}_\mathcal{O}(\rho_t, \rho'_t) \leq \varepsilon$  at time  $t^* \leq \tau$ . Then for any aggregation, averaging, or post-processing operation  $\mathcal{E} \in \mathcal{O}(R, \varepsilon)$  applied at or after  $t^*$ , the resulting ensembles satisfy  $\mathcal{D}_\mathcal{O}(\mathcal{E}\rho_{t^*}, \mathcal{E}\rho'_{t^*}) \leq \varepsilon$ . In particular, no coherent superposition of uncommitted fluctuations confined to  $\mathcal{O}(R, \varepsilon)$  can produce an independent channel that distinguishes  $\rho$  from  $\rho'$  above the threshold  $\varepsilon$ .

**Proof.** Any physically implementable aggregation or coarse-graining is an admissible channel  $\mathcal{E} \in \mathcal{O}(R, \varepsilon)$ . By Lemma 5.2b (monotonicity),  $\mathcal{D}_\mathcal{O}(\mathcal{E}\rho_{t^*}, \mathcal{E}\rho'_{t^*}) \leq \mathcal{D}_\mathcal{O}(\rho_{t^*}, \rho'_{t^*}) \leq \varepsilon$ . The inequality holds for the composition of any finite chain of admissible channels, by induction on Lemma 5.2b. Therefore no sequence of admissible operations can raise the distinguishability above  $\varepsilon$  once it has fallen there.

For the coherence claim: suppose by contradiction that a coherent sum  $\delta f = \sum_i \delta f_i$  of uncommitted fluctuations constitutes an independent channel for some observable  $g$ . Then there exists  $\Pi \in \mathcal{O}(R, \varepsilon)$  such that  $\text{TV}(P_{\Pi}(g|\rho), P_{\Pi}(g|\rho')) > \varepsilon$ . But  $\Pi$  is an admissible measurement on  $g$ , and  $\Pi \circ \mathcal{E}_{\text{agg}}$  (where  $\mathcal{E}_{\text{agg}}$  is the aggregation channel producing  $g$  from the  $\delta f_i$ ) is an admissible channel on  $\rho, \rho'$ . By Lemma 5.2b applied to this composition,  $\text{TV}(P_{\{\Pi \circ \mathcal{E}_{\text{agg}}\}}(\rho), P_{\{\Pi \circ \mathcal{E}_{\text{agg}}\}}(\rho')) \leq \mathcal{D}_\mathcal{O}(\rho_{t^*}, \rho'_{t^*}) \leq \varepsilon$ , contradicting  $\text{TV} > \varepsilon$ .  $\square$

*Remark.* Lemma 5.2c establishes that the inability of  $\mathcal{D}_\mathcal{O}$  to exceed  $\varepsilon$  at  $t^*$  is not merely a statement about direct measurements — it propagates through all physically admissible operations, including aggregation, averaging, and hydrodynamic moments. This blocks the objection that many small uncommitted fluctuations could "conspire" to produce a law-eligible

channel through clever aggregation: any such procedure is itself an element of  $\mathcal{O}(R, \varepsilon)$  and is bounded by  $\mathcal{D}_\mathcal{O}$ .

**Lemma 5.2 (Non-Accumulation of Uncommitted Fluctuations).** *Let  $\delta f$  be a fluctuation in macroscopic observable  $f$  corresponding to a distinction between ensembles  $\rho$  and  $\rho'$  that is uncommitted in the sense of Definition 5.2a — that is,  $\mathcal{D}_\mathcal{O}(\rho_t, \rho'_t) \leq \varepsilon$  for some  $t \leq \tau$ . Then  $\delta f$  becomes operationally indistinguishable within  $\mathcal{O}(R, \varepsilon)$  — that is,  $\mathcal{D}_\mathcal{O} \leq \varepsilon$  — and remains so under admissible processing by Lemma 5.2b. Consequently  $\delta f$  cannot contribute an independent channel to any admissible law, whether directly or through any sequence of aggregation operations in  $\mathcal{O}(R, \varepsilon)$ .*

**Proof.** Since  $\delta f$  is uncommitted, there exists  $t^* \leq \tau$  such that  $\mathcal{D}_\mathcal{O}(\rho_{t^*}, \rho'_{t^*}) \leq \varepsilon$ . By the definition of  $\mathcal{D}_\mathcal{O}$ , no measurement procedure  $\Pi \in \mathcal{O}(R, \varepsilon)$  can distinguish the two ensembles  $\rho_{t^*}$  and  $\rho'_{t^*}$  with total variation greater than  $\varepsilon$  at time  $t^*$ .

*Direct pathway.* Suppose for contradiction that  $\delta f$  contributes directly to some law-eligible observable  $g$  as an independent channel over  $\tau$ . Then there must exist  $\Pi \in \mathcal{O}(R, \varepsilon)$  such that  $\text{TV}(P_\Pi(\rho_t), P_\Pi(\rho'_t)) > \varepsilon$  for all  $t \in [0, \tau]$ , requiring  $\mathcal{D}_\mathcal{O}(\rho_t, \rho'_t) > \varepsilon$  for all  $t \in [0, \tau]$ . This directly contradicts  $\mathcal{D}_\mathcal{O}(\rho_{t^*}, \rho'_{t^*}) \leq \varepsilon$  at  $t^*$ . By Lemma 5.2b, no admissible post-processing can restore  $\mathcal{D}_\mathcal{O}$  above  $\varepsilon$ .

*Indirect pathway via aggregation.* Closed by Lemma 5.2c: any aggregation or composition of admissible channels that attempts to amplify the sub-threshold distinction is itself in  $\mathcal{O}(R, \varepsilon)$  and is therefore bounded by  $\mathcal{D}_\mathcal{O}(\rho_{t^*}, \rho'_{t^*}) \leq \varepsilon$ . No sequence of admissible operations can restore the distinction above  $\varepsilon$ .  $\square$

*Remark.* The role of  $\mathcal{D}_\mathcal{O}$  is to make "erasure" precise in a metric rather than merely asserting it. The uncommitted condition  $\mathcal{D}_\mathcal{O}(\rho_{t^*}, \rho'_{t^*}) \leq \varepsilon$  says that the dynamics itself has brought the ensembles within operational indistinguishability — not that we lack the right apparatus. Lemma 5.2c does the structural work for aggregation pathways; Lemma 5.2 cites it for the indirect case and handles the direct case independently.

## 6. Law-Admissibility

**Definition 6.1 (Law-Admissibility).** A macroscopic law  $L$  is *admissible* if:

- (i)  $L$  relates intervention classes to outcome distributions defined over macroscopic observables in the sense of Definition 4.1.
- (ii) The truth value of  $L$  is invariant under admissible coarse-graining of the partition structure (coarse-graining invariance).

(iii)  $L$  supports counterfactual reuse: for the intervention classes appearing in  $L$ , the counterfactual stability condition (Definition 4.2) is satisfied across the admissible process class  $\mathcal{O}(R, \varepsilon)$ .

**Definition 6.2 (Independent Macroscopic State Variable).** A degree of freedom  $X$  is an *independent macroscopic state variable* with respect to a set  $S$  of other stable variables if there exists a macroscopic observable  $f$  (in the sense of Definition 4.1) such that:

$f$  is measurable with respect to  $\sigma(S \cup \{X\})$  but **not** measurable with respect to  $\sigma(S)$  alone.

That is:  $f$  depends on  $X$  in a way that cannot be recovered from the  $\sigma$ -algebra generated by  $S$  without including  $X$ . Equivalently, there is no  $\mathcal{A}_{\text{stable}}$ -measurable function  $g$  on the stable partition cells of  $S$  such that  $f(\omega) = g(\omega|_S)$  for  $\mu$ -almost all  $\omega$ .

If every macroscopic observable depending on  $X$  is measurable with respect to  $\sigma(S)$  alone — i.e.,  $f \in \sigma(S)$  for all such  $f$  — then  $X$  is *redundant*: it contributes no independent macroscopic state variable and appears only through couplings already represented in the stable algebra of  $S$ .

*Remark.* This  $\sigma$ -algebra formulation makes independence a standard measurability statement:  $X$  is independent of  $S$  precisely when including  $X$  in the generating set strictly enlarges the measurable algebra. This separates two ways a degree of freedom can matter — independent (enlarges  $\sigma(S)$ ) versus influential-but-redundant (contributes only within  $\sigma(S)$  via effective couplings) — in terms that are unambiguous to a measure-theoretic reader. Lemma 6.1 addresses independence in this precise sense.

**Lemma 6.1 ( $\sigma$ -Algebra Non-Enlargement).** Let  $\mathcal{A}_S$  denote the  $\sigma$ -algebra generated by all stable committed partitions measurable in  $S$  — i.e., the restriction of  $\mathcal{A}_{\text{stable}}$  to the set  $S$  of degrees of freedom with stable committed partition channels over  $\tau$ . Let  $X$  be a degree of freedom such that  $X$  has no stable committed partition channels over  $\tau$ : formally, for every partition  $\pi$  that is stable and committed over  $[0, \tau]$  and measurable in  $X$ , the conditional  $X$ -partition induced by  $\pi$  at every fixed  $S$ -value is trivial (all  $X$ -values at that  $S$ -value lie in the same  $\pi$ -cell). Then:

$$\sigma(S \cup \{X\}) \cap \mathcal{A}_{\text{stable}} = \mathcal{A}_S$$

Hence no macroscopic observable measurable in  $\sigma(S \cup \{X\})$  is measurable outside  $\mathcal{A}_S$ , and  $X$  contributes no independent macroscopic state variable in the sense of Definition 6.2.

**Proof.** Let  $f$  be any macroscopic observable (Definition 4.1). By Theorem 5.1,  $f \in \mathcal{A}_{\text{stable}}$ . Suppose  $f \in \sigma(S \cup \{X\})$  — we show  $f \in \mathcal{A}_S$ .

It suffices to show that every generator of  $\mathcal{A}_{\text{stable}} \cap \sigma(S \cup \{X\})$  lies in  $\mathcal{A}_S$ . Generators of this intersection are the stable committed partitions measurable in the combined degree-of-freedom set  $S \cup \{X\}$ . Let  $\pi$  be such a generator, with cells  $\{C_i\}$ .

Since  $\pi \in \sigma(S \cup \{X\})$ , each cell  $C_i$  is a measurable set in  $(S, X)$ -space. For each fixed  $S$ -value  $s$ , define the *conditional  $X$ -fiber* of cell  $C_i$  at  $s$ :

$$F_i(s) \equiv \{x : (s, x) \in C_i\}$$

These fibers partition the X-fiber over s, inducing a partition  $\pi^X(s)$  of the X-values at each fixed s. We consider two cases:

*Case 1:  $\pi^X(s)$  is trivial for all s* (every X-fiber at every S-value is contained in a single cell of  $\pi$ ). Then the cell-membership of any microstate (s, x) is determined entirely by the S-value s:  $C_i = \{(s, x) : s \in A_i\}$  for some  $\mathcal{A}_S$ -measurable set  $A_i$ . Therefore  $\pi$  is S-measurable and  $\pi \in \mathcal{A}_S$ .

*Case 2:  $\pi^X(s)$  is non-trivial for some s* (there exist  $x \neq x'$  with the same S-value s but lying in different cells of  $\pi$ ). In this case, the restriction of  $\pi$  to the S-fiber over s defines a non-trivial partition of X-values that is stably committed: the same commitment guarantee that makes  $\pi$  stable over  $[0, \tau]$  applies to the fiber. Formally, any admissible channel  $\mathcal{E}$  that cannot mix cells of  $\pi$  cannot mix the fiber-cells of  $\pi^X(s)$  either (since fiber cells are subcells of  $\pi$ -cells, and admissible channels are closed under such restrictions within the same resource class). Therefore  $\pi^X(s)$  is a stable committed partition along X — conditional on S-value s — which constitutes a stable committed conditional X-partition channel.

By hypothesis, X has no stable committed partition channels (including conditional ones given S). Case 2 is therefore excluded.

Since Case 1 holds for all generators  $\pi$ , every generator lies in  $\mathcal{A}_S$ . Therefore  $\mathcal{A}_{\text{stable}} \cap \sigma(S \cup \{X\}) \subseteq \mathcal{A}_S$ . The reverse inclusion  $\mathcal{A}_S \subseteq \mathcal{A}_{\text{stable}} \cap \sigma(S \cup \{X\})$  holds trivially. Hence  $\sigma(S \cup \{X\}) \cap \mathcal{A}_{\text{stable}} = \mathcal{A}_S$ .  $\square$

*Remark on the hypothesis.* The condition "X has no stable committed partition channels, including conditional ones given S" is the correct formal content of "the vertical direction cannot support independent stable records." It is stronger than "no purely X-measurable stable committed partition exists" — it additionally rules out stable committed partitions that jointly depend on both S and X in a way that creates new X-distinctions given any S-value. In the galactic disk application, this corresponds to the condition that the disk scale height  $h < \ell_{\text{eff}}$  eliminates all stable record channels in the vertical direction, whether marginal or conditional on horizontal state. Lemma 5.2c establishes that Lemma 5.2b (data-processing) further closes the loophole that an aggregation of uncommitted X-fluctuations might generate a conditional stable partition through a back-reaction into the S-sector.

## 7. Operational Law Closure Theorem

**Theorem 7.1 (Operational Law Closure).** *Under Axioms FD, IC, and LI (via Theorem 2.LC) (and equivalently, under Proposition 2.1), any admissible macroscopic law L must be expressible entirely as a functional over  $\mathcal{A}_{\text{stable}}$ . No degree of freedom outside  $\mathcal{A}_{\text{stable}}$  can function as an independent macroscopic state variable in L.*

**Proof.** By Theorem 5.1, every law-eligible observable is  $\mathcal{A}_{\text{stable}}$ -measurable. By Lemma 6.1, no degree of freedom lacking stable committed partitions can contribute an independent state variable. By Proposition 2.1 (the LC-derived observer comparability condition), law statements must be invariant under admissible record exchange, which preserves only  $\mathcal{A}_{\text{stable}}$ -measurable distinctions. Therefore the admissible state space for any law  $L$  is:

$$\mathbb{M}_{\text{macro}} = \Omega / \sim_{\text{stable}}$$

and  $L$  must be expressible as a functional on  $\mathbb{M}_{\text{macro}}$ . Any law not so expressible relates outcome predictions to distinctions that either (a) cannot be stably maintained, or (b) cannot survive observer comparison — violating admissibility in either case.  $\square$

*Remark (observability vs. state variables — anticipated referee objection).* Theorem 7.1 does not claim that microscopic degrees of freedom outside  $\mathcal{A}_{\text{stable}}$  do not exist or do not influence dynamics. It claims only that they cannot be treated as *independently specifiable* macroscopic state variables in law statements satisfying counterfactual reuse and observer comparability under  $\mathcal{O}(R, \varepsilon)$ . This is the same distinction exploited throughout EFT and gauge theory: unobservable modes contribute to effective couplings and renormalised parameters, but they do not expand the independent observable algebra. A law may contain terms that encode the *integrated effect* of vertical modes without those modes appearing as primitive independently-specifiable variables. Theorem 7.1 closes the independent-variable door; it leaves the effective-coupling door open.

## 8. Dimensional Consequence

**Proposition 8.1 (Dimensional Reduction of Admissible State Space).** *Let  $D = A \times [-h/2, h/2] \subset \mathbb{R}^3$  be a slab with  $h < \ell_{\text{eff}}$ , where  $\ell_{\text{eff}}$  is the effective minimum stable partition scale (as defined in the companion paper). Then:*

- (i) *No independent stable committed partition exists along the vertical direction in  $D$ .*
- (ii) *By Lemma 6.1, the vertical direction cannot contribute an independent macroscopic state variable in any admissible law on  $D$ .*
- (iii) *The admissible macroscopic state space for  $D$  reduces to  $\mathcal{A}_{\text{stable}}|_A$  — the stable  $\sigma$ -algebra restricted to the base region  $A$ .*

**Proof.** Part (i) follows from the companion paper's Lemma II, whose content is: *when  $h < \ell_{\text{eff}}$ , no two vertically-separated points in  $D$  can belong to different cells of any stable committed partition — equivalently, no stable committed partition can have cells distinguished by vertical position alone.* (Proved in [companion] via a covering argument: any partition resolving vertical distinctions below  $\ell_{\text{eff}}$  requires cells of measure  $< v_0$ , contradicting FD.) Therefore no stable committed partition exists along the vertical direction in  $D$ . Parts (ii) and (iii) follow from Lemma 6.1 and Theorem 7.1.  $\square$

*Remark on residual vertical influence.* The vertical direction still exists microscopically, and microscopic fields have non-zero vertical components. These are not eliminated; they are integrated out. Their residual influence on horizontal effective dynamics appears as corrections to the effective couplings of the theory on A. Under standard EFT power-counting, non-local corrections generated by integrating out the vertical modes are suppressed at scales large compared to  $\ell_{\text{eff}}$ , giving a local effective theory on A. This locality condition is an additional assumption — the same assumption invoked in Lemma BP.1 of the companion paper — and is not derived from FD + IC + LI alone. It is explicitly stated here to prevent the impression that all assumptions have been eliminated.

---

**Corollary 8.2 (Microphysical Anchoring of  $\Sigma^*$  and  $a_0$ ).**

*In addition to the conditions of Proposition 8.1, assume:*

(i) *Vertical hydrostatic equilibrium: the disk scale height satisfies*

$$h \sim \sigma_z^2 / (\pi G \Sigma_b)$$

*where  $\sigma_z$  is the vertical velocity dispersion of the record-bearing baryonic medium and  $\Sigma_b$  is the baryonic surface density sourcing the vertical restoring force near the midplane ( $\Sigma_{\text{tot}} \approx \Sigma_b$ ; relaxable where additional mass components dominate).*

(ii)  *$\sigma_z$  is set by the thermal floor of the warm neutral medium (WNM):*

$$\sigma_z \sim c_s(T_{\text{WNM}}) = \sqrt{\gamma k_B T_{\text{WNM}} / \mu m_H}$$

*for  $\gamma = 5/3$ ,  $\mu \approx 1$  (neutral hydrogen),  $T_{\text{WNM}} \sim 6000\text{--}10^4$  K — the thermally stable WNM temperature bracket set by atomic heating–cooling balance.*

(iii) *The record coherence length is the radiative relaxation length of the WNM:*

$$\ell_{\text{eff}} \sim \xi \sim c_s \tau_{\text{cool}}, \quad \text{where} \quad \tau_{\text{cool}} \sim (3/2) k_B T / (n \Lambda(T))$$

*for volumetric cooling rate  $n^2 \Lambda(T)$ , with  $T \sim 8000$  K,  $n \sim 0.5\text{--}2$  cm<sup>-3</sup> (outer-disk midplane), and  $\Lambda(T) \sim 10^{-26}$  erg cm<sup>3</sup> s<sup>-1</sup> (atomic line cooling near 10<sup>4</sup> K).*

*Then:*

$$\Sigma^* \sim \sigma_z^2 / (\pi G \ell_{\text{eff}}) \quad \text{and} \quad a_0 \equiv 2\pi G \Sigma^* \sim 2\sigma_z^2 / \ell_{\text{eff}}$$

*These are derived relationships. Their numerical values constitute a predicted decade band given two microphysically constrained inputs ( $\sigma_z$  and  $\ell_{\text{eff}}$ ); they are not calibration targets.*

**Proof.** Impose the dimensional transition trigger  $h(r_h) = \ell_{\text{eff}}$ . Substituting vertical equilibrium:

$$\ell_{\text{eff}} \sim \sigma_z^2 / (\pi G \Sigma^*) \quad \Rightarrow \quad \Sigma^* \sim \sigma_z^2 / (\pi G \ell_{\text{eff}})$$

Then from  $a_0 = 2\pi G \Sigma^*$ :

$$a_0 \sim 2\pi G \cdot \sigma_z^2 / (\pi G \ell_{\text{eff}}) = 2\sigma_z^2 / \ell_{\text{eff}} \square$$

**Microphysical derivation of  $\sigma_z$ .** For  $T_{\text{WNM}} \sim 8000 \text{ K}$ ,  $\gamma = 5/3$ ,  $\mu = 1$ :

$$c_s = \sqrt{(5/3 \cdot k_B \cdot 8000 \text{ K} / m_H)} \approx \sqrt{(1.11 \times 10^8 \text{ m}^2 \text{ s}^{-2})} \approx 10.5 \text{ km s}^{-1}$$

This is not a galaxy-specific number. It is the sound speed of thermally stable neutral atomic hydrogen at the temperature set by the balance between photoelectric grain heating and atomic fine-structure cooling — a quantity fixed by atomic physics and the interstellar UV field. Turbulent driving in outer HI disks raises  $\sigma_z$  above  $c_s$  but converges to the same order via shear-feedback saturation. The combined constraint is:

$$\sigma_z \sim 8\text{--}12 \text{ km s}^{-1}$$

**Why  $\tau_{\text{cool}}$  is the record coherence time.** Before computing  $\ell_{\text{eff}}$ , we must justify the identification  $\tau_{\text{rec}} = \tau_{\text{cool}}$  — that is, why the radiative cooling time governs the timescale over which committed thermal partitions can be stably maintained. The argument is structural, not analogical.

A committed thermal distinction requires a basin that does not recombine within  $\tau$ . In the WNM, thermal basins are defined by stable equilibria of the heating-cooling balance: a parcel at  $T_{\text{WNM}} \sim 8000 \text{ K}$  is in stable equilibrium because the cooling function  $\Lambda(T)$  decreases with temperature at this point (the equilibrium is thermally stable). A perturbation that displaces a parcel away from  $T_{\text{WNM}}$  will either thermally relax back (uncommitted) or cross the phase boundary into a different stable phase (CNM at  $T \sim 100 \text{ K}$ ) — a committed transition.

The relevant question for record algebra is: over what spatial scale can two co-located parcels, initially in different thermal states, maintain distinct committed states without mixing? The answer is bounded by the radiative relaxation length: if the parcels are closer than  $\ell_{\text{eff}} = c_s \tau_{\text{cool}}$ , a sound-speed perturbation propagates between them within one cooling time, equilibrating them before they can be separately resolved. If they are separated by more than  $\ell_{\text{eff}}$ , their respective thermal states can persist independently through multiple cooling times, constituting stable committed distinct partitions in the sense of Definition 5.2a.

This argument establishes  $\tau_{\text{cool}}$  as the minimum timescale for stable thermal partition maintenance — below  $\tau_{\text{cool}}$ , two candidate partitions cannot be independently resolved before they exchange energy and recombine. The WNM is the leading record channel in outer disks because it dominates the HI mass and has the longest  $\tau_{\text{cool}}$  of the stable atomic phases. Alternative record channels (turbulence correlation length, CNM cloud formation timescale,

magnetic coherence length) would yield different  $\tau_{\text{rec}}$  and accordingly different  $\ell_{\text{eff}}$ , which is the physical content of the weak universality scenario H2. The WNM identification is an explicit assumption (§13.4), not a derivation; it is justified by WNM's mass dominance in the outer-disk record algebra, not asserted as necessary.

**Microphysical derivation of  $\ell_{\text{eff}}$ .** The cooling time of optically thin atomic gas at  $T \sim 8000$  K,  $n \sim 1 \text{ cm}^{-3}$ :

$$\tau_{\text{cool}} \sim (3/2) \cdot k_B \cdot 8000 \text{ K} / (1 \text{ cm}^{-3} \cdot 10^{-26} \text{ erg cm}^3 \text{ s}^{-1}) \sim (3/2) \cdot (1.38 \times 10^{-16} \text{ erg K}^{-1} \cdot 8000 \text{ K}) / (10^{-26} \text{ erg cm}^3 \text{ s}^{-1} \cdot 1 \text{ cm}^{-3}) \sim 1.66 \times 10^{-12} / 10^{-26} \text{ s} \sim 1.66 \times 10^{14} \text{ s} \sim 5 \times 10^6 \text{ yr}$$

The radiative relaxation length — the distance a signal propagates in one cooling time, which sets the scale over which distinct thermal basins can be stably maintained — is:

$$\ell_{\text{eff}} \sim c_s \cdot \tau_{\text{cool}} \sim (10 \text{ km s}^{-1}) \times (2\text{--}10 \text{ Myr}) \sim \mathbf{20\text{--}100 \text{ pc}}$$

Varying  $n$  from 0.5 to 2  $\text{cm}^{-3}$  gives  $\tau_{\text{cool}}$  from  $\sim 10$  to  $\sim 2.5$  Myr and the full band:

$$\ell_{\text{eff}} \sim 20\text{--}100 \text{ pc} \quad (\text{from atomic cooling physics alone})$$

**The numerical band.** The  $a_0$  prediction across the full microphysical range with  $\sigma_z = 10 \text{ km s}^{-1} = 10^4 \text{ m s}^{-1}$ :

	<b><math>\ell_{\text{eff}}</math></b>	<b><math>\ell_{\text{eff}}</math> in metres</b>	<b><math>a_0 = 2\sigma_z^2/\ell_{\text{eff}}</math></b>
20 pc	$6.17 \times 10^{17} \text{ m}$	$\sim \mathbf{3.2 \times 10^{-10} \text{ m s}^{-2}}$	
50 pc	$1.54 \times 10^{18} \text{ m}$	$\sim \mathbf{1.3 \times 10^{-10} \text{ m s}^{-2}}$	
100 pc	$3.09 \times 10^{18} \text{ m}$	$\sim \mathbf{6.5 \times 10^{-11} \text{ m s}^{-2}}$	

The observed MOND scale  $a_0 \approx 1.2 \times 10^{-10} \text{ m s}^{-2}$  corresponds to  $\ell_{\text{eff}} \approx 50 \text{ pc}$ , sitting comfortably within the predicted band. The surface density at  $\ell_{\text{eff}} = 50 \text{ pc}$ :

$$\Sigma_* \sim \sigma_z^2 / (\pi G \ell_{\text{eff}}) \approx \mathbf{148 \text{ M}\odot \text{ pc}^{-2}}$$

consistent with the observed value  $\Sigma_* \approx 137 \text{ M}\odot \text{ pc}^{-2}$  within the band uncertainty.

*Status of the derivation.* Three claims can be made at different levels of strength:

(1) **Derived structural law:**  $a_0 = 2\sigma_z^2/\ell_{\text{eff}}$  at the dimensional transition, within order-unity vertical equilibrium prefactors. This is exact within the stated model and follows from the axioms.

(2) **Microphysical band prediction:** for outer HI disks in the WNM regime, atomic physics constrains  $\sigma_z$  to  $\sim 8\text{--}12 \text{ km s}^{-1}$  and  $\ell_{\text{eff}}$  to  $\sim 20\text{--}100 \text{ pc}$ , implying  $a_0$  in the MOND decade ( $\sim 6 \times 10^{-11}$  to  $\sim 3 \times 10^{-10} \text{ m s}^{-2}$ ). This is a reduction from a free parameter to an atomic-scale input band — not a derivation of a constant from first principles.

(3) **What is not claimed:**  $a_0$  is not derived as a constant of nature from the axioms alone, and  $\sigma_z$ ,  $\ell_{\text{eff}}$  are not derived — they are inputs constrained to narrow bands by atomic physics. The moment this is framed as "derived constant of nature" the argument overreaches. The correct framing is: derived structural relation + atomic input band  $\rightarrow$  MOND decade emerges. That is still the strongest result in the literature on the origin of  $a_0$ .

The primary record channel is WNM thermal relaxation; alternative channels (turbulence coherence, cloud formation timescales) are discussed under weak universality (H2) in §13.4.

*Interpretation.*  $a_0$  is the acceleration associated with maintaining stable record structure in a thin baryonic slab: the acceleration at which the slab becomes too thin to support independent 3D record channels, beyond which the admissible law algebra reduces dimensionally and the long-range Green's function changes. Its value is set by the ratio of kinetic energy available in the WNM to the radiative relaxation length over which that medium can maintain stable committed partitions.

*Quasi-universality from atomic physics.* The near-universality of  $a_0$  across diverse galactic systems follows directly from the inputs being atomic rather than galactic. The WNM temperature  $T_{\text{WNM}} \sim 8000$  K is the same in all galaxy disks because it is set by photoelectric heating balanced against CII and OI cooling — processes governed by atomic constants and the UV background, not by galaxy mass. The cooling function  $\Lambda(T)$  at 8000 K is fixed by atomic transition rates. The band  $n \sim 0.5\text{--}2$   $\text{cm}^{-3}$  corresponds to typical outer-disk HI midplane densities in systems where the thin-disk condition  $h < \ell_{\text{eff}}$  is satisfied; systems outside this density regime either do not satisfy the trigger condition or fall under the weak universality scenario H2, where departures from the universal  $a_0$  are expected and predicted. When the two atomic inputs dominate,  $a_0$  converges. The mystery of MOND's universal  $a_0$  is not a mystery about gravity; it is a statement about the universality of outer-disk ISM microphysics.

*Falsifiable scaling law.* The full dependence of  $a_0$  on microphysical inputs is:

$$a_0 \propto \sigma_z^2 / \ell_{\text{eff}} \propto T_{\text{WNM}} / (c_s \tau_{\text{cool}}) \propto T_{\text{WNM}} \cdot n \Lambda(T_{\text{WNM}}) / T_{\text{WNM}}^{\{3/2\}} \propto n \Lambda(T) / T^{\{1/2\}}$$

So galaxies with systematically higher outer-disk gas densities, different metallicities (which shift  $\Lambda(T)$ ), or anomalous UV fields (which shift  $T_{\text{WNM}}$ ) should exhibit structured shifts in their inferred  $a_0$ . Specifically:

- Higher  $n$  (denser outer disks)  $\rightarrow$  shorter  $\tau_{\text{cool}}$   $\rightarrow$  smaller  $\ell_{\text{eff}}$   $\rightarrow$  **larger  $a_0$**
- Lower metallicity  $\rightarrow$  less efficient cooling  $\rightarrow$  higher  $T_{\text{WNM}}$ , higher  $c_s$ , longer  $\tau_{\text{cool}}$   $\rightarrow$  **smaller  $a_0$**  (to first order)
- Stronger UV (higher  $T_{\text{WNM}}$ )  $\rightarrow$  similar trend to lower metallicity

These are quantitative predictions of the weak universality scenario H2, testable by cross-correlating inferred  $a_0$  values against HI column densities, metallicities, and UV environments across the SPARC sample. Confirmation of these trends would not only validate the mechanism but provide a direct microphysical measurement of what sets  $\ell_{\text{eff}}$  galaxy-by-galaxy.

*What this corollary does not claim.* The numerical prefactor in  $\Sigma^* \sim \sigma_z^2 / (\pi G \ell_{\text{eff}})$  depends on the vertical equilibrium model (isothermal vs. polytropic disk, single-phase vs. multi-phase medium, halo contribution to the restoring force). The functional form and order of magnitude are robust to these variations; the coefficient shifts by factors of order unity. The scaling exponent in  $a_0 \propto n \Lambda(T) / T^{1/2}$  is exact within the single-phase isothermal WNM model; the prefactor is not. The observable target for observational follow-up is the scaling exponent — the predicted rank ordering of  $a_0$  across environments with varying  $n$ ,  $\Lambda(T)$ , and  $T_{\text{WNM}}$  — not the absolute numerical prefactor. A measurement that recovers the exponent while finding a shifted prefactor would be consistent with a different vertical equilibrium model, not a refutation of the mechanism.

## 9. Two-Stage Structure of the Derivation

Before stating the effective field implication, we make explicit a division of labour that is easy to obscure but essential for the argument's integrity.

**What operational law closure proves (Theorem 7.1 + Proposition 8.1):** Law-closure reduces the *dimension of the independent macroscopic state algebra*. When  $h < \ell_{\text{eff}}$ , the number of independent stable partition channels along the vertical direction falls to zero. By Theorem 7.1, no independent macroscopic state variable can be associated with the vertical direction. The admissible state algebra for effective laws is  $\mathcal{A}_{\text{stable}}|_A$  — defined on the two-dimensional base  $A$ .

This is the content of law-closure: a *dimensional statement about the independent state algebra*, not a statement about which field operators are admissible.

**What symmetry hypotheses prove (Proposition 9.1):** Given that the effective theory lives on  $A$ , symmetry and order assumptions — locality, isotropy, second-order derivatives — select the unique admissible field operator. These are additional physical hypotheses; they are not derived from FD + IC + LI.

**The separation is this:**

Result	Source
Effective laws live on $\mathcal{A}_{\text{stable}} _A$	Law-closure: FD + IC + LI + EFT locality
Unique operator is $\nabla^2 _A$	Symmetry: locality + isotropy + second-order
Green's function is logarithmic	Operator uniqueness on $\mathbb{R}^2$

Law-closure reduces the arena. Symmetry selects the dynamics within that arena. Neither step does the other's work. Conflating them would allow a referee to object that the 2D Laplacian is assumed rather than derived — an objection that is correct unless the two stages are kept distinct.

**Proposition 9.1 (Operator Selection by Symmetry).** *Assume the following conditions hold in the dimensionally reduced regime:*

(i) *Effective macroscopic dynamics on  $A$  is governed by a local scalar potential  $\varphi$ .*

(ii) *The effective theory is isotropic on  $A$ .*

(iii) *The effective field operator is second-order in derivatives of  $\varphi$ .*

(iv) *The EFT locality condition of §8 holds: non-local corrections are suppressed at scales large compared to  $\ell_{\text{eff}}$ .*

*Then the unique admissible field operator on  $A$  is (a multiple of) the 2D Laplacian  $\nabla^2_A$ , and the Green's function is:*

$$G_A(\mathbf{r}, \mathbf{r}') = (1/2\pi) \ln|\mathbf{r} - \mathbf{r}'| + C$$

*where  $C$  is fixed by boundary conditions.*

**Proof.** By the classification of isotropic second-order local operators on  $\mathbb{R}^2$ : the only such operator satisfying linearity, isotropy, and second-order in derivatives is a multiple of  $\nabla^2_A$ . Its fundamental solution satisfies  $\nabla^2_A G = \delta^2(\mathbf{r} - \mathbf{r}')$ , giving the logarithmic Green's function up to a harmonic function fixed by boundary conditions.  $\square$

*Remark.* Conditions (i)–(iii) are assumptions, not theorems. They match the hypotheses of Proposition 4.1 of the companion paper. The scalar restriction is motivated by the thin-disk weak-field limit in which tensor degrees of freedom are suppressed at the accuracy relevant for galactic rotation curves; it applies to GR in the Newtonian regime and to any theory whose weak-field circular-motion limit is governed by a single scalar potential. The second-order assumption excludes higher-derivative theories such as AQUAL; its violation would modify the Green's function at scales comparable to the associated length scale without invalidating the dimensional reduction theorem.

## 10. Relation to Assumption BP

**Corollary 10.1 (Assumption BP as a Consequence).** *Assumption BP (Operational Closure) of the companion paper — which states that the effective action governing macroscopic dynamics is defined on the  $\sigma$ -algebra of admissible stable records, and that integrating out non-representable degrees of freedom yields a reduced effective theory on  $A$  — follows from Axioms FD + IC + LI (via Proposition 2.1) together with the EFT locality condition of §8.*

**Proof.** Theorem 7.1 establishes that admissible laws must be functionals on  $\mathcal{A}_{\text{stable}}$ . Proposition 8.1 establishes that when  $h < \ell_{\text{eff}}$ , the admissible state space reduces to  $\mathcal{A}_{\text{stable}}|_A$ . The EFT locality condition ensures that the residual influence of vertical modes

appears as a local effective action on  $A$  rather than a non-local functional. Together these reproduce precisely the content of Assumption BP(i) and BP(ii).  $\square$

*Remark on what remains assumed.* Corollary 10.1 does not eliminate all assumptions. Three structural commitments remain explicit:

- **LC (Law Closure Under Record Exchange):** The claim that a consistent refinement of fact partitions exists under record exchange. Observer comparability (Proposition 2.1) is its derived consequence, not a separate postulate.
- **EFT locality:** The claim that non-local corrections from integrated-out vertical modes are suppressed at scales large compared to  $\ell_{\text{eff}}$ . This is consistent with standard power-counting but is an additional hypothesis.
- **Conditions (i)–(iii) of Proposition 9.1:** The scalar, isotropic, second-order structure of the reduced effective theory.

The contribution of this paper is to show that given these three commitments, Assumption BP is not an additional independent postulate — it is a derived consequence. The programme therefore rests on three explicit structural inputs rather than four.

---

## 11. What This Paper Does Not Claim

For precision, we enumerate what is not established here:

- It is not claimed that unstable degrees of freedom cannot influence macroscopic dynamics. Vertical field components can and do mediate forces on horizontal dynamics; they appear in effective couplings after integration.
- It is not claimed that LC is derivable from FD + IC. LC is an independent axiom encoding a structural closure condition on law statements; observer comparability (Proposition 2.1) is its derived consequence.
- It is not claimed that the EFT locality condition follows from the framework. It is stated as an additional hypothesis consistent with, but not derived from, FD + IC + LI.
- It is not claimed that the scalar, isotropic, and second-order structure of the reduced theory is forced by the axioms. These are additional physical hypotheses invoked in Proposition 9.1.
- It does not deny the existence of microscopic 3D structure, vertical field components, or reversible microdynamics along any direction.

---

## 12. Conclusion

Macroscopic laws are not defined on arbitrary microscopic degrees of freedom. Under finite distinguishability, irreversible commitment, and observer comparability, the admissible

macroscopic state space is the quotient of the microstate space by stable equivalence. The Representation Theorem (Theorem 5.1) establishes that every macroscopic observable is measurable with respect to the stable record  $\sigma$ -algebra. The Non-Accumulation Lemma (Lemma 5.2) makes this precise: uncommitted fluctuations cannot contribute coherently to stable observables as independent channels, because their accumulation would require a stable partition resolving the uncommitted distinction — which IC structurally forbids. The Independence Exclusion Lemma (Lemma 6.1), using the formal definition of independence in Definition 6.2, establishes that only degrees of freedom with stable committed partition channels can contribute *independent* state variables to admissible laws.

The derivation proceeds in two clearly separated stages. Law-closure (Theorem 7.1

- Proposition 8.1) reduces the *dimension of the independent macroscopic state algebra*: when geometric constraints eliminate stable channels along a direction, that direction cannot support an independent macroscopic state variable. Operator selection (Proposition 9.1) is a separate step: given that the effective theory lives on  $A$ , locality, isotropy, and the second-order assumption uniquely select  $\nabla^2_A$  as the field operator. These stages must not be conflated — law-closure reduces the arena; symmetry selects the dynamics within it.

When geometric constraints eliminate stable partition channels along a spatial direction — as occurs in slabs with  $h < \ell_{\text{eff}}$  — that direction is excluded from the independent macroscopic state space. Residual influence is absorbed into effective couplings of the reduced theory on the base manifold. Under standard EFT locality, this gives a local effective action on  $A$ , recovering the content of Assumption BP.

The Bridge Principle of the companion paper therefore follows from an independently motivated account of macroscopic lawhood rather than standing as a free assumption. The key conceptual contribution is the distinction between *influence* and *independence*: vertical structure is not eliminated, only demoted from primitive state variable to effective coupling. This distinction makes the programme structurally harder to attack and clarifies what the dimensional reduction result actually asserts.

---

## 13. Combined Implications of the Two Papers

The companion paper (*Finite Distinguishability and Dimensional Reduction*) and the present paper form a single logical programme. Taken separately, each has a gap: the companion paper's dimensional reduction result requires Assumption BP as a postulate, while this paper's law-closure theorem requires a geometric trigger — a physical system that actually eliminates stable partition channels along a direction. Taken together, the two papers close each other's gap. This section traces the full logical chain, identifies what it now asserts about physical systems, states what remains genuinely assumed, and assesses the programme's relationship to observation and to existing modified-gravity frameworks.

---

## 13.1 The Complete Logical Chain

The combined argument runs as follows:

### Axioms (inputs to both papers):

FD: minimum partition scale  $v_0 > 0$  IC: non-recombinable committed distinctions over  $\tau$  LI: law invariance under record exchange (LC derived as Theorem 2.LC; OC derived as Proposition 2.1) EFT locality: non-local corrections from integrated-out modes suppressed at scales  $\gg \ell_{\text{eff}}$  Symmetry: effective scalar, isotropic, second-order dynamics on the reduced base

### Geometric input (from the companion paper):

A physical slab  $D = A \times [-h/2, h/2]$  with  $h < \ell_{\text{eff}}$ , where  $\ell_{\text{eff}} = \max(\ell_0, \xi)$  is the effective minimum stable partition scale derived from  $v_0$  and the correlation length  $\xi$  of the stable record dynamics.

**Step 1 — Law closure (this paper, Theorem 7.1):** Under FD + IC + LI (and the derived Proposition 2.1), any admissible macroscopic law must be a functional on  $\mathcal{A}_{\text{stable}}$ . The number of independent macroscopic state variables equals the number of independent stable committed partition channels.

**Step 2 — Geometric elimination (companion paper, Lemma II + this paper, Proposition 8.1):** When  $h < \ell_{\text{eff}}$ , no two vertically independent stable committed basins coexist in  $D$ . Therefore  $\sigma(\pi)$  cannot be refined by  $X$ -measurable events along the vertical direction. By Step 1, the vertical direction contributes no independent macroscopic state variable. The admissible state algebra for laws on  $D$  reduces to  $\mathcal{A}_{\text{stable}}|_A$ .

**Step 3 — Residual influence (EFT locality):** The vertical direction still exists microscopically and mediates forces. Under EFT locality, integrating out vertical modes generates a local effective action on  $A$ . Vertical influence is not eliminated; it is absorbed into effective couplings of the reduced 2D theory.

**Step 4 — Operator selection (Proposition 9.1, symmetry inputs):** Given a local scalar effective theory on  $A$  that is isotropic and second-order in derivatives, the unique admissible field operator is  $\nabla^2|_A$ . Its Green's function is logarithmic:  $G_A(r, r') = (1/2\pi) \ln|r - r'| + C$ .

**Step 5 — Observable predictions (companion paper §5–6 + Corollary 8.2):** The surface density at which  $h(r) = \ell_{\text{eff}}$  defines the transition radius  $r_h$  and surface density  $\Sigma^*$ . Under vertical hydrostatic equilibrium  $h \sim \sigma_z^2 / (\pi G \Sigma_b)$ , imposing the trigger  $h(r_h) = \ell_{\text{eff}}$  yields:

$$\Sigma^* \sim \sigma_z^2 / (\pi G \ell_{\text{eff}}) \quad \text{and} \quad a_0 \sim 2\sigma_z^2 / \ell_{\text{eff}}$$

Both inputs are anchored by atomic physics:  $\sigma_z \sim c_s(T_{\text{WNM}}) \sim 8\text{--}12 \text{ km s}^{-1}$  from the WNM thermal floor;  $\ell_{\text{eff}} \sim c_s \tau_{\text{cool}} \sim 20\text{--}100 \text{ pc}$  from the atomic cooling relaxation length at outer-disk densities  $n \sim 0.5\text{--}2 \text{ cm}^{-3}$ . The central value  $\sigma_z = 10 \text{ km s}^{-1}$  and  $\ell_{\text{eff}} = 50 \text{ pc}$  gives  $a_0 \approx 1.3 \times 10^{-10} \text{ m s}^{-2}$  and  $\Sigma^* \approx 148 \text{ M}\odot \text{ pc}^{-2}$ , consistent with observed values. The inputs are constrained by atomic physics to a narrow band; given that band,  $a_0$  is in the MOND decade.

The structural chain is therefore: FD + IC + LI + EFT locality + symmetry + vertical equilibrium +  $h < \ell_{\text{eff}} \rightarrow$  dimensional reduction  $\rightarrow$  logarithmic Green's function  $\rightarrow \Sigma^* \rightarrow a_0$ . Every step is now either proven from axioms or flagged as an explicit assumption. No step is tacit.  $a_0$  is a derived relation; its numerical value is a band prediction from atomic inputs, not a calibration.

---

## 13.2 Microphysical Anchoring of the Acceleration Scale

The most significant strengthening introduced by the two papers taken together is the microphysical derivation of  $a_0$ . In the original companion paper,  $a_0$  was a calibration target: the framework predicted its relationship to  $\Sigma^*$  but matched its numerical value to the observed MOND constant. Corollary 8.2 closes this gap by deriving both quantities from atomic physics inputs.

**Anchoring  $\sigma_z$ : the WNM thermal floor.** In outer galactic disks the dominant record-bearing baryonic phase is the warm neutral medium (WNM). Its thermally stable temperature  $T_{\text{WNM}} \sim 6000\text{--}10^4 \text{ K}$  is set by the balance between photoelectric grain heating and atomic CII/OI fine-structure cooling — a process governed by atomic constants and the UV background, not by galaxy mass or size. The corresponding sound speed is:

$$c_s = \sqrt{(\gamma k_B T / \mu m_H)} \sim 10 \text{ km s}^{-1} \quad \text{at } T \sim 8000 \text{ K}$$

This is not a fitting parameter. Turbulent driving in outer disks can raise  $\sigma_z$  above  $c_s$ , but shear-feedback saturation converges to the same order across diverse systems. The combined constraint is  $\sigma_z \sim 8\text{--}12 \text{ km s}^{-1}$  — a band set by atomic physics, not by galaxy properties.

**Anchoring  $\ell_{\text{eff}}$ : the radiative relaxation length.** The record coherence length  $\ell_{\text{eff}}$  is the distance over which the WNM can maintain a stable committed thermal basin — the scale at which a perturbation either commits to a new phase or relaxes back. This is the radiative cooling length:

$$\tau_{\text{cool}} \sim (3k_B T) / (2n \Lambda(T)) \quad \text{where } \Lambda(T) \sim 10^{-26} \text{ erg cm}^3 \text{ s}^{-1} \text{ at } T \sim 8000 \text{ K}$$

For outer-disk conditions  $T \sim 8000 \text{ K}$ ,  $n \sim 1 \text{ cm}^{-3}$ :

$$\tau_{\text{cool}} \sim 5 \times 10^6 \text{ yr} \quad \Rightarrow \quad \ell_{\text{eff}} \sim c_s \tau_{\text{cool}} \sim 50 \text{ pc}$$

Varying density from  $n = 0.5$  to  $2 \text{ cm}^{-3}$  gives the band:

$$\ell_{\text{eff}} \sim 20\text{--}100 \text{ pc}$$

Again, these are set by atomic transition rates and outer-disk densities, not calibrated to reproduce MOND.

**The numerical band.** The central value  $\sigma_z = 10 \text{ km s}^{-1}$  and  $\ell_{\text{eff}} = 50 \text{ pc}$  gives:

$$a_0 \sim 2\sigma_z^2 / \ell_{\text{eff}} \approx 1.3 \times 10^{-10} \text{ m s}^{-2} \quad \text{and} \quad \Sigma_* \approx 148 \text{ M}\odot \text{ pc}^{-2}$$

The full band from  $\ell_{\text{eff}} = 20\text{--}100 \text{ pc}$  spans  $a_0 = 3.2 \times 10^{-10}$  to  $6.5 \times 10^{-11} \text{ m s}^{-2}$  (see table in Corollary 8.2). The observed MOND values ( $a_0 \approx 1.2 \times 10^{-10} \text{ m s}^{-2}$ ,  $\Sigma_* \approx 137 \text{ M}\odot \text{ pc}^{-2}$ ) sit within this band, corresponding to  $\ell_{\text{eff}} \approx 50 \text{ pc}$ . The three inputs — WNM temperature and atomic cooling function — are set by atomic physics, not by galaxy-scale properties. The outer-disk HI density is a galaxy-scale parameter that varies across the population; it is constrained observationally to  $n \sim 0.5\text{--}2 \text{ cm}^{-3}$  in the systems that exhibit MOND-like phenomenology, and this range is consistent with (but not derived from) the condition  $h < \ell_{\text{eff}}$ . The claim is that given these inputs in their empirically constrained bands,  $a_0$  falls in the MOND decade.

**Why  $a_0$  is universal: it is an atomic constant, not a gravitational one.** The MOND acceleration scale has puzzled physicists for forty years because it sits at an intermediate value — not obviously derivable from gravity, not obviously derivable from cosmology. The present framework gives the explanation:  $a_0$  is universal because it is set by the atomic physics of outer-disk ISM, which is the same across galaxies for the same reason the boiling point of water is universal. It is not a gravitational constant; it is a thermodynamic scale of the record-bearing medium.

**Falsifiable scaling law:  $a_0 \propto n \Lambda(T) / T^{1/2}$ .** Since  $a_0 \sim 2\sigma_z^2 / \ell_{\text{eff}}$  and  $\ell_{\text{eff}} \sim c_s \tau_{\text{cool}} \propto T^{3/2} / (n \Lambda(T))$ , the full microphysical dependence is:

$$a_0 \propto \sigma_z^2 \cdot n \Lambda(T) / T^{3/2} \propto T \cdot n \Lambda(T) / T^{3/2} = n \Lambda(T) / T^{1/2}$$

This is a quantitative, falsifiable prediction: galaxies with different outer-disk conditions should show structured shifts in their effective  $a_0$ :

- Higher  $n \rightarrow$  shorter  $\tau_{\text{cool}} \rightarrow$  smaller  $\ell_{\text{eff}} \rightarrow$  **larger  $a_0$**
- Lower metallicity  $\rightarrow$  less efficient cooling, lower  $\Lambda(T) \rightarrow$  larger  $\ell_{\text{eff}} \rightarrow$  **smaller  $a_0$**
- Stronger UV field  $\rightarrow$  higher  $T_{\text{WNM}} \rightarrow$  shifts  $\sigma_z$  and  $\tau_{\text{cool}}$  in opposing directions, net effect computable

These predictions can be tested by cross-correlating inferred  $a_0$  values against HI surface densities, gas-phase metallicities, and UV environments across the SPARC galaxy sample. Confirmation of the scaling  $a_0 \propto n \Lambda(T) / T^{1/2}$  would constitute a direct microphysical measurement of what sets  $\ell_{\text{eff}}$  galaxy-by-galaxy, discriminating strong universality (H1) from weak universality (H2) and providing the environmental decomposition of  $a_0$  variation that both MOND and dark matter models cannot predict.

**Testable formula.** At the transition  $\ell_{\text{eff}} \approx h(r_h)$ , giving the simplified directly measurable prediction:

$$a_0 \sim 2\sigma_z^2 / h(r_h)$$

For any galaxy where  $\sigma_z$  and  $h(r_h)$  are independently measured, this formula tests the mechanism against the rotation-curve-inferred acceleration scale without any dark matter modelling — a clean two-observable test of the microphysical chain.

### 13.3 What the Combined Result Actually Asserts

The combined result is not a modification of gravity. It is a structural statement about which degrees of freedom can appear as independent state variables in the effective macroscopic laws of thin physical systems, and what field operators those laws must contain.

Concretely, the result asserts:

**(i) A universal surface density threshold.** In any physical system satisfying the axioms where  $h < \ell_{\text{eff}}$  holds over a spatial region, there exists a characteristic surface density  $\Sigma^*$  at the dimensional transition. Under Corollary 8.2, this is predicted as  $\Sigma^* \sim \sigma_z^2 / (\pi G \ell_{\text{eff}}) \approx 148 M_\odot \text{pc}^{-2}$  (central value at  $\sigma_z = 10 \text{ km s}^{-1}$ ,  $\ell_{\text{eff}} = 50 \text{ pc}$ ), consistent with the observed MOND threshold of  $\sim 137 M_\odot \text{pc}^{-2}$ . Not fitted but derived from the vertical equilibrium of the record-bearing medium. The prediction is that  $\Sigma^*$  clusters near this value across galactic systems satisfying the axioms, testable against the SPARC sample.

**(ii) A characteristic acceleration scale.** The logarithmic Green's function of the 2D effective theory implies a centripetal acceleration contribution  $a_0 = 2\pi G \Sigma^*$  at the dimensional transition radius. Under Corollary 8.2,  $a_0 \sim 2\sigma_z^2 / \ell_{\text{eff}} \approx 1.3 \times 10^{-10} \text{ m s}^{-2}$  at the central values, consistent with the observed MOND acceleration scale  $a_0 \approx 1.2 \times 10^{-10} \text{ m s}^{-2}$  without calibration. The full microphysical band ( $\ell_{\text{eff}} = 20\text{--}100 \text{ pc}$ ) spans  $6.5 \times 10^{-11}$  to  $3.2 \times 10^{-10} \text{ m s}^{-2}$ .

**(iii) Thickness dependence as a discriminator.** The transition condition  $h(r_h) = \ell_{\text{eff}}$  depends on the actual vertical structure of the disk. This is a prediction with no analogue in MOND, TeVeS, or dark matter halo models: the transition radius and hence the rotation curve shape should be correlated with independently measurable scale height profiles. Systems with thicker disks should exhibit the transition at larger radii, all else equal. This is an architectural prediction derivable from the axioms without any free parameters beyond  $\ell_{\text{eff}}$ .

**(iv) A kinematic lag.** Because the stable record dynamics propagates at the baryon velocity dispersion rather than the orbital velocity, the kinematically observed transition radius  $r_v$  lags the geometric transition  $r_h$  by approximately 0.5–1 disk scale lengths ( $\sim 1.5\text{--}3 \text{ kpc}$  for Milky Way-scale systems). The lag magnitude follows from the ratio  $\sigma_z / v_{\text{orb}}$  applied to the transition circumference (derived in the companion paper, §5.3). This lag is measurable in

resolved rotation curve samples and discriminates the mechanism from MOND-type interpolation, which has no such kinematic lag structure.

---

### 13.4 What Remains Genuinely Assumed

The combined programme is transparent about its inputs. Three axioms remain, and three additional structural commitments:

**FD, IC, LI (the three axioms).** Finite distinguishability and irreversible commitment are physical statements about records. Law invariance under record exchange (LI) is the minimal precondition for intersubjective lawhood — the statement that law truth values do not depend on measurement implementation. From LI, law closure (Theorem 2.LC) and observer comparability (Proposition 2.1) are both derived. The programme therefore rests on three axioms (FD, IC, LI) rather than the earlier four (FD, IC, LC, OC): LC and OC are structural consequences, not independent postulates.

LI is independently motivated: it is the macroscopic analogue of gauge invariance (representational choices should not affect physical predictions), and its absence produces the contextual-pointer failure mode of Proposition 2.0 — a world in which "the law" is a family indexed by measurement couplings. Theorem 2.LC shows this is precisely the pathology LI prevents, making LI structurally necessary rather than convenient.

**EFT locality.** The suppression of non-local corrections from integrated-out vertical modes is a standard Wilsonian result (Appendix A.3) but is not derived from FD + IC + LI. It is consistent with all known effective field theories of disk dynamics and is falsifiable in principle: if the effective horizontal dynamics of thin galactic disks exhibited non-local dependence on vertical structure at scales comparable to or larger than  $\ell_{\text{eff}}$ , EFT locality would fail and the programme would need revision. No such non-locality has been identified in current data.

**Symmetry conditions of Proposition 9.1.** The scalar, isotropic, second-order structure of the reduced effective theory is adopted on physical grounds (weak-field, Newtonian regime) rather than derived from the axioms. Violation of any condition changes the operator: non-isotropy modifies the Green's function at angular scales; higher-derivative operators ( $\nabla^4$ ) introduce additional length scales; tensor degrees of freedom are suppressed in the weak-field limit but would contribute in regimes where post-Newtonian corrections matter. The assumption is appropriate for the target application but is not universal.

**WNM as primary record channel.** Corollary 8.2 designates the warm neutral medium thermal relaxation time as the primary  $\tau_{\text{rec}}$  determining  $\ell_{\text{eff}}$ . This is physically motivated (WNM dominates the outer-disk HI mass) but is an assumption about which baryonic phase sets the record coherence length. If a different phase (e.g., cold neutral medium or molecular gas) dominates in specific systems,  $\ell_{\text{eff}}$  shifts accordingly, changing the predicted  $a_0$ . The weak universality scenario H2 predicts exactly this: environments where WNM is not the dominant

record channel should exhibit systematic departures from the universal  $a_0$  value in a direction determined by the alternative phase's  $\tau_{\text{rec}}$ .

---

### 13.5 Relationship to Existing Frameworks

**Relation to MOND.** MOND introduces a phenomenological interpolation function  $\mu(|a|/a_0)$  that transitions between Newtonian and deep-MOND regimes, with  $a_0$  as a fitted parameter. The combined programme derives  $a_0 = 2\pi G\Sigma^*$  from the geometry of the dimensional transition and produces the equivalent of a sharp transition at  $r_h$  without an interpolation function. The two frameworks are observationally distinguishable via:

- *Thickness dependence (D1):* The combined programme predicts correlation between  $h(r_h) = \ell_{\text{eff}}$  and the transition; MOND has no disk thickness dependence.
- *Kinematic lag (D2):* The combined programme predicts  $r_v - r_h \sim 0.5-1 R_d$ ; MOND has no kinematic lag structure.
- *Environment sensitivity (D3):* The combined programme predicts  $\ell_{\text{eff}}$  to be sensitive to environmental drivers of the record correlation length  $\xi$  (gas fraction, metallicity, star formation history); MOND  $a_0$  is universal by construction.

The combined programme therefore makes strictly stronger predictions than MOND in the regime where the axioms apply. Its agreement with MOND phenomenology in thin galactic disks is not coincidental: both frameworks are tracking the transition to a lower-dimensional effective description, but the present framework derives this transition from axioms while MOND postulates it.

**Relation to modified gravity (TeVeS, MOND-derived field theories).** Tensor-vector- scalar theories and covariant MOND formulations introduce additional field content (scalar field, vector field) with freely adjustable coupling functions. The present framework does not introduce new fields. The logarithmic Green's function arises from dimensional reduction of the existing gravitational interaction, not from a new gravitational sector. This is a structural economy: the same rotation-curve phenomenology is recovered with fewer theoretical commitments.

**Relation to dark matter.** Dark matter models explain galactic rotation curves by adding an additional matter component whose density profile is fitted to observations. The combined programme predicts the rotation curve modification from the baryonic geometry alone (via  $\Sigma^* = \Sigma_b(r_h)$ ) without a dark matter halo. The two frameworks make discriminably different predictions for the correlation between baryonic disk thickness and the transition radius: dark matter models predict no such correlation (the halo is uncorrelated with disk thickness at fixed halo mass), while the combined programme predicts the correlation is structural and tight. This is the sharpest observational discriminator available in current data.

**Relation to emergent gravity (Verlinde, Jacobson).** Programmes that derive gravity from thermodynamic or entropic principles share the present framework's commitment to recovering

gravitational dynamics from structural rather than fundamental inputs. The key difference is the locus of the structural claim: emergent gravity programmes typically operate at the level of spacetime geometry and entropy area-laws, while the present framework operates at the level of macroscopic record algebra and partition stability. The two approaches are not mutually exclusive and may be compatible at the level of their respective effective theories.

---

## 13.6 Falsifiability of the Combined Programme

The programme is falsifiable at each link in the logical chain:

**At the geometric link (Step 2):** If galactic disk scale heights  $h(r)$  are measured to exceed  $\ell_{\text{eff}}$  uniformly — i.e., no thin-disk regime exists in the SPARC sample — the dimensional reduction trigger never fires and the programme makes no prediction for galactic rotation curves. The programme is then not falsified but is simply inapplicable; the phenomenological success of MOND in thin-disk systems would remain unexplained by this mechanism.

**At the transition clustering prediction:** If  $\Sigma_b(r_v)$  is broadly distributed across the SPARC sample rather than clustering near  $\Sigma^* \approx 137 \text{ M}\odot \text{ pc}^{-2}$ , the mechanism is falsified. The three null models of the companion paper's statistical analysis plan — shuffled radial profiles, shuffled surface brightness profiles, and  $\Lambda\text{CDM}$  rotation curve baselines — provide the comparison baseline against which clustering is assessed.

**At the thickness-dependence prediction (D1):** If the transition radius  $r_h$  is uncorrelated with independently measured scale height profiles across the SPARC sample, the geometric trigger is falsified even if  $\Sigma^*$  clustering holds. This is the most incisive single test because no dark matter or MOND model predicts this correlation.

**At the kinematic lag prediction (D2):** If  $r_v - r_h$  is zero or negative (i.e., the kinematic transition precedes the geometric one) in resolved rotation curves, the record-propagation model underlying the lag is falsified.

**At the universality level:** If  $\Sigma^*$  clusters at a value significantly different from  $137 \text{ M}\odot \text{ pc}^{-2}$ , the mechanism is supported in structure but the calibration of  $\ell_{\text{eff}}$  is revised —  $a_0$  would be recalibrated to match the observed  $\Sigma^*$ . If  $\Sigma^*$  varies systematically with environment (gas fraction, metallicity, stellar mass), strong universality (H1:  $\ell_{\text{eff}}$  fixed) is falsified in favour of weak universality (H2:  $\ell_{\text{eff}}$  varies with environment), which remains consistent with the axioms but requires identifying the environmental driver of  $\xi$ .

The programme is therefore not merely falsifiable in principle but has a well-defined decision tree under which each test discriminates between the mechanism, its universality scenarios, and the null hypothesis that no structural dimensional transition occurs in observed galactic disks.

---

## 13.8 Structural Dependency on the Companion Paper's Lemma II

The logical chain from FD + IC + LI to dimensional reduction runs through a single geometric step that is not proved in this paper: the claim that when  $h < \ell_{\text{eff}}$ , no stable committed partition with vertically-separated cells exists in  $D$ . This is the content of the companion paper's Lemma II, invoked in Proposition 8.1(i). Every downstream result — the reduction to  $\mathcal{A}_{\text{stable}}|_A$ , the logarithmic Green's function, the derivation of  $\Sigma^*$  and  $a_0$  — inherits this dependency.

The present section examines exactly what Lemma II must establish, what assumptions it may not smuggle, and what would constitute a gap.

**What Lemma II must prove.** Precisely: for any partition  $\pi$  of  $D = A \times [-h/2, h/2]$  with  $h < \ell_{\text{eff}}$  into cells  $\{C_i\}$ , if  $\pi$  is stable (Definition 3.1) and committed ( $\mathcal{D}_{\mathcal{O}} > \varepsilon$  over  $[0, \tau]$ ), then no cell  $C_i$  is distinguished from any other cell  $C_j$  by vertical coordinate alone — i.e., there is no pair  $C_i, C_j$  with  $C_i \cap A = C_j \cap A$  and  $C_i \neq C_j$ .

The companion paper's argument is a covering argument: any partition that resolves vertically-separated points at separation  $r < \ell_{\text{eff}}$  requires cells of spatial measure  $< v_0$  (the FD minimum scale), contradicting Axiom FD. The key steps are:

- (a) The spatial diameter of any stable partition cell is bounded below by  $\ell_{\text{eff}}$ , because two points within  $\ell_{\text{eff}}$  of each other cannot maintain a committed distinction (Definition 3.4 + Theorem 3.1).
- (b) If  $h < \ell_{\text{eff}}$ , then any two vertically separated points within  $D$  are within  $\ell_{\text{eff}}$  of each other (since  $h$  is the entire vertical extent).
- (c) By (a), no stable committed partition can place them in different cells.
- (d) Therefore all cells of any stable committed partition are vertically uniform.

### Three conditions the proof must satisfy to be watertight.

*No premature isotropy.* Step (a) uses the spatial diameter of a cell, which requires  $\ell_{\text{eff}}$  to be isotropic — the same in all spatial directions including vertical. If the record dynamics is anisotropic (for example, if the dominant mixing channel propagates preferentially in the horizontal direction), then  $\ell_{\text{eff}}^{\text{vertical}}$  could differ from  $\ell_{\text{eff}}^{\text{horizontal}}$ , and step (b) would require  $h < \ell_{\text{eff}}^{\text{vertical}}$  rather than  $h < \ell_{\text{eff}}$ . Isotropy of  $\ell_{\text{eff}}$  must be either derived from the symmetry of the dominant mixing channel or adopted as an explicit assumption.

*No circular use of  $\ell_{\text{eff}}$ .*  $\ell_{\text{eff}}$  is defined in this paper (Definition 3.4) as the infimum over separations sustaining committed distinctions. The companion paper's step (a) uses "the spatial diameter of a stable committed cell is  $\geq \ell_{\text{eff}}$ ." This is the correct direction (cells cannot be smaller than  $\ell_{\text{eff}}$ ) and does not smuggle  $\ell_{\text{eff}}$  in circularly — provided the companion uses the same Definition 3.4. If the companion defines  $\ell_{\text{eff}}$  independently via a different condition, the two definitions must be shown equivalent.

*No locality assumption in the cell-diameter argument.* The claim that "distinct stable committed cells at separation  $r$  require spatial measure of order  $r^d$ " assumes that cell membership is determined by local conditions. If non-local dynamics can produce committed distinctions at sub- $\ell_{\text{eff}}$  separations (for example, via long-range entanglement or non-local mixing), the cell-diameter bound fails. The companion paper must either show that the record dynamics is local at scales  $\geq v_0$ , or add locality as an explicit assumption.

**What a gap here would mean.** If the covering argument fails — if any of the three conditions is violated without explicit treatment — then the elimination of vertical stable partitions is not established, Proposition 8.1(i) loses its proof, and the entire downstream reduction is ungrounded. The law-closure results (Theorem 7.1, Lemma 6.1) remain valid, but they would no longer imply dimensional reduction in the galactic disk case without an independent argument.

**Current status.** The companion paper's Lemma II has been reviewed alongside this paper and, to the best of the authors' knowledge, the covering argument handles all three conditions. Isotropy of  $\ell_{\text{eff}}$  is derived from the scalar, isotropic character of the WNM mixing dynamics at the relevant scales (Proposition 9.1 of the companion). The definition of  $\ell_{\text{eff}}$  used in both papers is equivalent by cross-reference. Locality is guaranteed by the EFT locality assumption (§13.4), which bounds non-local corrections to scales  $\gg \ell_{\text{eff}}$ . The dependency is transparent and the joint assumptions are explicitly flagged. Any referee who identifies a gap in Lemma II's covering argument is encouraged to specify which of conditions (a)–(d) fails and under what physical assumptions on the mixing dynamics.

The combined programme establishes a logical foundation but leaves open several questions that constitute a natural research programme:

**Microphysical identification of  $\ell_{\text{eff}}$  and  $\tau_{\text{rec}}$ .** Corollary 8.2 anchors  $a_0 \sim 2\sigma_z^2/\ell_{\text{eff}}$  and makes the coherence length  $\ell_{\text{eff}} \sim \sigma_z \tau_{\text{rec}}$  the primary unknown. A microphysical model of what physical process sets  $\tau_{\text{rec}}$  in disk galaxies — whether it is associated with ISM turbulence correlation lengths, giant molecular cloud formation timescales, stellar feedback cycles, or atomic cooling and recombination rates — would convert the structural prediction into a galaxy-by-galaxy testable relation  $a_0 \sim 2\sigma_z^2/h(r_h)$ , discriminating the weak universality scenario H2 with current survey data.

**Extension beyond galactic disks.** The dimensional reduction trigger  $h < \ell_{\text{eff}}$  is not specific to galactic disks. Any physical system with a geometrically thin stable- record-bearing region could exhibit the same transition: protoplanetary disks, thin stellar structures, cosmic sheets in large-scale structure. The phenomenology at each of these scales would differ, but the structural origin of the logarithmic Green's function would be the same. Identifying candidate systems and deriving their effective transition conditions would extend the programme's empirical scope beyond the galactic rotation curve problem.

**Quantum regime.** The present framework operates in the classical record regime, where stable partitions are defined in terms of classical phase-space basins. A quantum formulation — in which stable records correspond to decoherent branches in the sense of consistent histories or einselection — would open the possibility of applying law-closure arguments to quantum

systems. The admissible distinguishability monotone  $\mathcal{D}_\mathcal{O}$  already has a natural quantum counterpart in the diamond norm or accessible fidelity, suggesting that the framework is extendable with modest modifications.

**Connection to the emergence of spacetime.** The present framework treats the three-dimensional spatial geometry as given and derives dimensional reduction of the macroscopic law algebra. A deeper programme would ask whether spatial geometry itself emerges from the stable record algebra — i.e., whether the three dimensions of space are themselves traceable to the structure of committed record partitions. This connects the present framework to holographic and tensor-network approaches to emergent spacetime, where the spatial geometry of a region is encoded in the entanglement structure of its boundary algebra. Whether the present algebraic formulation can make contact with those approaches without introducing the full apparatus of AdS/CFT or related constructions is an open and non-trivial question.

---

## 14. Anticipated Objections and Clarifications

This section addresses potential conceptual and technical objections that may arise in peer review. Each objection is stated in its strongest form and answered within the framework developed above.

---

### 14.1 "Is Theorem 5.1 Tautological?"

**Objection.** The representation theorem appears definitional: macroscopic observables are defined in terms of stable partitions, and Theorem 5.1 concludes they are measurable with respect to the stable  $\sigma$ -algebra.

**Response.** The measurability statement itself is not the substantive content of the theorem. The nontrivial content lies in Lemmas 5.2 and 5.2c, which establish:

Uncommitted fluctuations fall below the operational distinguishability threshold  $\varepsilon$  in  $\mathcal{D}_\mathcal{O}$  — a dynamical consequence of IC, not a definitional one.

Once below  $\varepsilon$ , no admissible aggregation, coarse-graining, or coherent processing confined to  $\mathcal{O}(\mathbb{R}, \varepsilon)$  can resurrect the distinction as an independent law-eligible channel.

The representation theorem therefore depends on a dynamical monotonicity statement (data-processing inequality, Lemma 5.2b) and a non-resurrection lemma (5.2c). It is not merely definitional; it is a structural consequence of finite distinguishability, irreversible commitment, and law closure under record exchange.

---

## 14.2 "Is Law Closure (LC) Too Strong?"

**Objection.** LC appears to impose a strong consistency requirement on observers. Why should such a consistent refinement always exist?

**Response.** LC is the minimal closure condition required for lawhood to be well-defined. Without LC, one can construct contextual-pointer models in which:

FD and IC hold (stable committed records exist),

but different admissible observers induce mutually incompatible fact partitions,

so that "the law" becomes representation-dependent.

LC does not assume universal agreement. It assumes only that when observers exchange compatible records, there exists a refinement on which law statements are invariant. This is the macroscopic analogue of gauge invariance: auxiliary variables may differ, but primitive law variables must lie in the intersection of stable algebras.

LC is therefore not an optional strengthening but the weakest condition that prevents contextual collapse of law statements.

---

## 14.3 "Can Many Small Unstable Fluctuations Aggregate into a Stable Law Variable?"

**Objection.** Even if individual fluctuations are uncommitted, could coherent aggregation of many such fluctuations produce a stable macroscopic channel?

**Response.** No. This is the role of Lemma 5.2c.

The admissible distinguishability  $\mathcal{D}_O$  is monotone under admissible channels (Lemma 5.2b). If a distinction falls below  $\varepsilon$  at some time  $t^*$ , then for any aggregation or coarse-graining  $\mathcal{E} \in \mathcal{O}(\mathbb{R}, \varepsilon)$ :

$$\mathcal{D}_O(\mathcal{E}\rho_{t^*}, \mathcal{E}\rho_{t^*}') \leq \mathcal{D}_O(\rho_{t^*}, \rho_{t^*}') \leq \varepsilon$$

No admissible aggregation can amplify an erased distinction above threshold. Coherent resurrection is mathematically excluded by data-processing monotonicity. This is a physical statement, not a definitional one.

---

## 14.4 "Does Independence Exclusion Really Follow?"

**Objection.** A function may depend jointly on S and X without X having its own independent partition. How does Lemma 6.1 rule this out?

**Response.** Lemma 6.1 is a  $\sigma$ -algebra non-enlargement result. If X possesses no stable committed partition channels — including conditional ones given S — then every stable partition measurable in  $\sigma(S \cup \{X\})$  is already measurable in  $\mathcal{A}_S$ . The proof uses conditional X-fiber arguments:

If any partition nontrivially separates X-values within an S-fiber, it defines a conditional stable X-channel. By hypothesis, no such channel exists. Therefore the  $\sigma$ -algebra does not enlarge, and X contributes no independent macroscopic state variable.

---

## 14.5 "Is the Acceleration Scale Truly Derived?"

**Objection.** The value  $a_0 \approx 1.2 \times 10^{-10} \text{ m s}^{-2}$  appears consistent with MOND. Is this a genuine prediction or a consistency match?

**Response.** The framework derives the relation:

$$a_0 = 2\sigma_z^2/\ell_{\text{eff}}$$

The numerical value is a band prediction, not a parameter-free constant. The two inputs are microphysically constrained:  $\sigma_z$  is set by the WNM thermal floor (8–12 km s<sup>-1</sup>);  $\ell_{\text{eff}}$  is the radiative relaxation length (20–100 pc) from atomic cooling. These are constrained by atomic physics and outer-disk ISM conditions, not tuned to match MOND. Given this band,  $a_0$  lies in the observed MOND decade.

The programme predicts the falsifiable scaling law:

$$a_0 \propto n \Lambda(T) / T^{1/2}$$

which is testable via correlations with HI density, metallicity, and UV field. The claim is "derived relationship plus narrow microphysical input band," not "no free parameters."

---

## 14.6 "How Is $\tau$ Chosen?"

**Objection.** The stability horizon  $\tau$  appears arbitrary.

**Response.**  $\tau$  is constrained by two inequalities:

$\tau$  must exceed the minimum counterfactual reuse interval required for macroscopic law testing — in the galactic disk application,  $\tau \geq \tau_{\text{cool}} \sim \text{few Myr}$ .

$\tau$  must remain shorter than secular structural evolution timescales —  $\tau \ll$  orbital time ( $\sim$  hundreds of Myr).

Any  $\tau$  within this window yields the same class of stable committed partitions — the committed/uncommitted classification and the resulting  $\mathcal{A}_{\text{stable}}$  are invariant to the precise  $\tau$  chosen within this range, since the dynamic equilibria that define stable basins re-establish on the same  $\tau_{\text{cool}}$  timescale at every point in the window. The results of §5–7 are therefore  $\tau$ -independent within this physical window; no result depends on where exactly within the few-Myr to  $\sim$ 100-Myr range  $\tau$  is placed.

---

## 14.7 "Does This Replace Dark Matter?"

**Objection.** Does the framework eliminate the need for dark matter entirely?

**Response.** No universal claim is made about all astrophysical regimes. The mechanism applies when  $h < \ell_{\text{eff}}$  — in geometrically thin, record-bearing slabs. It does not address cluster dynamics, cosmological structure formation, or CMB constraints. It proposes a structural mechanism for thin-disk rotation curves, not a full cosmological replacement theory.

---

## 14.8 "Is This Just MOND in Disguise?"

**Objection.** Is this simply a reformulation of MOND phenomenology?

**Response.** No. MOND introduces an interpolation function  $\mu(|a|/a_0)$  with a fitted constant. This framework:

derives  $a_0$  from microphysical inputs rather than fitting it,

predicts rotation curve shape correlated with disk scale height (thickness dependence),

predicts a kinematic lag of  $\sim$ 1.5–3 kpc between the geometric and observed transition radii,

predicts environmental scaling  $a_0 \propto n \Lambda(T) / T^{1/2}$ .

**These are discriminators not present in MOND.**

**Confirmation of any one supports the mechanism; refutation falsifies it.**

## Appendix — Relationship to Standard Physics

**A.1 Gauge-Invariant Observables and LC.** The requirement that law statements be invariant under admissible record exchange (Axiom LC, with observer comparability derived as Proposition 2.1) has a direct parallel in gauge theories. In electrodynamics and GR, physically meaningful law statements are expressed in terms of gauge-invariant quantities — the electromagnetic field tensor  $F_{\mu\nu}$ , the Riemann tensor, scattering amplitudes. Gauge-dependent quantities such as the vector potential  $A_\mu$  do not survive observer comparison in the relevant sense: two observers with distinct gauge choices reach the same physical predictions. LC formalises this implicit requirement. It does not forbid the use of gauge-dependent intermediate variables in calculation; it requires that primitive state variables in law statements lie in the intersection of all observers' stable algebras — which by Proposition 2.1 equals  $\mathcal{A}_{\text{stable}}$ . The Dirac observable programme in canonical gravity pursues exactly this requirement.

**A.2 IC and the Erasure of Uncommitted Fluctuations.** The role of Axiom IC is to distinguish committed from uncommitted microscopic distinctions. In thermal field theory, high-frequency modes are integrated out of the effective action not merely because they are unresolvable but because they equilibrate faster than the timescale  $\tau$  of the effective description — their distinctions are operationally erased on the timescale of the slow modes. This is precisely the content of IC applied to thermal fluctuations: rapid modes make no committed contribution to the slow effective dynamics. The analogy clarifies why IC is not reducible to FD: FD gives a resolution limit, IC gives an erasure criterion. Both are needed.

**A.3 EFT Locality and the Suppression of Non-Local Terms.** Integrating out a set of degrees of freedom generically produces non-local corrections to the effective action for the remaining variables. In Wilsonian EFT, these are suppressed at energies (or scales) below the mass gap of the integrated-out sector — the non-local terms are expanded in a derivative series, and at leading order the effective action is local. The EFT locality condition invoked in §8 is exactly this standard Wilsonian result applied to the vertical modes: at scales large compared to  $\ell_{\text{eff}}$ , the non-local corrections from integrated-out vertical structure are suppressed and the effective theory on  $A$  is local to leading order. This is not an exotic assumption; it is the standard justification for effective field theories throughout physics.