Logical Dependence of Physical Determinism on Set-theoretic Metatheory

Justin Clarke-Doane, Columbia University[1]

Draft: October 17, 2025

Abstract

Baroque questions of set-theoretic foundations are widely assumed to be irrelevant to physics. In this article, I demonstrate that this assumption is incorrect. I show that the fundamental physical question of whether a theory is deterministic—whether it fixes a unique future given the present—can depend on one's choice of set-theoretic axiom candidates over which there is intractable disagreement. This dependence is not confined to hypothetical examples. It reaches into mainstream, foundational, and frontier physics, including full discrete systems, the preferred basis problem in quantum mechanics, and the dynamics of Kerr-like black hole interiors.

I argue that beyond the familiar analytic notion of well-posedness, a theory's determinism profile depends on a **regularity layer**, on whether the definable sets that carry our ensemble and canonicalization talk are measurable, have the Baire property, and admit measurable selectors. Competing axiom candidates extending ZFC—Gödel's **Axiom of Constructibility** (**V=L**) and **large cardinal (LC)** assumptions strong enough to imply **Projective Determinacy** (**PD**)—diverge on these regularity facts. The divergence has three faces. First, **coherence**: weak formulations presuppose measurability of coefficients and under V=L one can arrange definable pathologies that collapse the statement of the weak problem, while under PD all projective sets are regular. Second, **uniqueness**: many determinism results are ensemble claims—"for almost all initial data there is a unique continuation"—whose sense depends on measurability or Baire category at projective complexity. PD secures this, while V=L may not. Third, **identity**: when multiple admissible continuations remain, physical practice demands a **canonical**, representation-independent choice. That demand is a measurable uniformization problem. PD supports measurable, symmetry-constrained selectors at the Π^1_2 level, while V=L guarantees at most Δ^1_2 (hence possibly non-measurable) tie-breaks. I develop a number of live cases, showing how the regularity properties toggle with the metatheory. The upshot is that which extension of standard ZFC we adopt changes what our best-supported theories say. I close by sketching a research program, **reverse physics**, on analogy with Friedman's and Simpson's reverse mathematics, whose aim is to map a theory's physical content against the foundational axioms that make its ensemble and canonical claims intelligible. I conclude that, given the entanglement

[1] Thanks to Avner Ash for catching a mistake in my treatment of PDEs, Will Cavendish for pressing me to consider the discrete case, Gabriel Goldberg for helpful discussion of complexity classes and projective uniformization in L, Joel David Hamkins for general exchanges about the relevance of incompleteness to determinism, John Norton for feedback on a talk I gave on the same topic, and students in my Undecidability in Physics Fall 2025 graduate seminar at Columbia for probing questions about the project. This article is an experiment for me. Claude (Anthropic) and GPT 5 (OpenAI) provided important support, similar to that of research assistants, including efficient references to known results and constructions, to fit my argument template. The argument template, and of course all errors, are my own. See Berry 2023, Forthcoming; Colyvan 1998; Clarke-Doane 2022, 2024; Farah & Magidor 2012; Jonas 2024; and Pour-El & Richards 1989 and the references to follow for cognate discussions.

of set-theoretic metatheory and physics, either physical theories must be relativized to set theories (in which case physics itself becomes relative), or, as Quine (1951, 1990) controversially argued, the search for new axioms to settle undecidables may admit of empirical input.

1. Introduction: Set Theory Meets Physical Content

Physical theories are formulated with mathematics, yet the choice of set-theoretic background is almost never considered. This paper shows that even choice of *extensions* of standard set theory can have profound physical consequences. I focus on **determinism** in the standard **Hadamard sense** (Hadamard 1923): for admissible initial data there exists a solution, it is unique, and it depends continuously on the data (Evans 2010, §2.1, §6.4; Brezis 2011, Ch. 9). These verdicts turn on background analytic hypotheses—such as the measurability of coefficients, regularity of domains, compactness of admissible sets, and the availability of measurable selectors—and those hypotheses can, I show, depend on speculative extensions of standard set theory.

Let **ZFC** denote Zermelo-Fraenkel set theory with Choice (Jech 2003). Like any recursively axiomatized theory, it must be incomplete if consistent (Gödel 1931). The statements which it fails to settle reach well into ordinary mathematics (Cohen 1963, 1966; Solovay 1970). Two opposed extensions of ZFC have been proposed to mitigate incompleteness. **Large cardinals (LC)** strong enough to imply **Projective Determinacy (PD)** —for instance, ω many Woodin cardinals (Kanamori 2009, Ch. 6; Martin & Steel 1989)—guarantee that all projective sets are Lebesgue measurable and have the Baire and perfect-set properties (Kechris 1995, §28, §38–39; Moschovakis 2009, Ch. 6). Alternatively, Gödel's **Axiom of Constructibility (V = L)** (Gödel 1940; Jech 2003, Ch. 13) yields a minimal universe in which simply definable projective sets have pathological regularity properties. The two outlooks "embody radically different conceptions of the universe of sets" (Jensen 1995, 401).[2]

We will see that such simply definable sets can occur inside textbook physical models—as coefficients in weak formulations (Evans 2010, §5.8, §6.2), as "thin" barriers whose capacity determines uniqueness (Ma & Röckner 1992; Fukushima, Oshima & Takeda 2011, Ch. 2), or as admissible sets or global tie-breaking rules. When this is the case, *V*=*L* and LC disagree about what I call the **determinism profile** of the system—the coherence, uniqueness and identity of solutions to the system's equations. The equations do not change; only the metatheory does.

[2] For discussion of other "restrictive" axioms, besides *V*=*L*, see Fraenkel, Bar‑Hillel & Levy (1973, §6.4). A common narrative among set theorists is that *V* = *L* is false, and "clearly" so. Maddy (1997, Pt II, § 4) contains a nice explication of the reasons advocates of this narrative supply. But, as we will see, this view is "not unanimously shared" (Fontanella 2019, 32) Devlin writes, "What is my own view?...Currently I tend to favour [*V*=*L*]....At the moment I think I am in the majority of informed mathematicians, but the minority of set theorists..."[1981, 205]. Arrigoni says, "I believe it perfectly in order to characterize . . . *ZFC* + *V*=*L* as intuitively plausible . . . " [2011, 355]. And Pinter writes, "[T]here is a strong intuitive basis for considering *L* to be the class of all sets. By definition, *L* contains all the sets that are describable by a formula in the language of set theory. And there is no practical reason to admit sets which lack any description, for we would never make use of such sets. They would merely sit there and muddy the waters. Thus, from this point onward we shall assume the following important axiom: Axiom of Constructibility: Every set is constructible, that is, every set is in L. This axiom is usually denoted by the symbol *V* = *L*." (2014, 227).

Before turning to specific cases, it is helpful to state explicitly what, in the ensuing examples, depends on the metatheory and what does not. Some features of mathematical structure are invariant across all transitive models of ZFC containing the same reals. Others are exquisitely sensitive to whether one works within *V=L* or under large cardinals sufficient to imply *Projective Determinacy (PD)*. Membership in any lightface Δ^1_2 set of reals is absolute, as are the truth values of Σ^1_2 and Π^1_2 formulas with real parameters—facts guaranteed by Shoenfield's absoluteness theorem. By contrast, the *regularity* properties of such sets—whether they are measurable, possess the Baire property, or admit perfect subsets—depend on the set-theoretic background. Under *V=L*, there exist lightface Δ^1_2 sets that are non-measurable and nonmeager; under *PD*, every projective set is measurable, has the Baire property, and contains a perfect subset if uncountable. The same contrast determines whether standard "genericity" idioms—phrases like "for almost all" or "for a comeager set of initial data"—have any well-defined sense at all. Finally, many uniformization and selection results that guarantee measurable or Borel choice functions for definable relations are valid at the projective level only under determinacy assumptions. They can fail in *V=L*. The analytic core of the physics examples—the local PDEs, variational arguments, and energy estimates—remains within ZFC. What toggles is not the existence of equations or the validity of pointwise theorems, but the status of ensemble-level statements that presuppose regularity (i.e., with respect to Lebesgue-complete – or at least Borel – probability measures or Baire category; these presuppose measurability/Baire regularity). Whenever the subsequent discussion invokes "almost every," "typical," or "generic," that presupposition will be noted explicitly.[3]

The moral I will press throughout is modest in form but not in effect. I will not argue that the analytic layer of our best theories rests on exotic axioms. I will argue that when we elevate pointwise analytic facts into ensemble or canonical claims indispensable for science, and when we enforce the scientific ideal that such claims be robust across every admissible way of probing or encoding the same underlying situation, we thereby adopt a Π^1_2-shaped robustness clause:

$$\mathrm{UT}(x,y) :\Leftrightarrow \forall \rho \in \mathscr{P} \subseteq \mathbb{N}^\wedge\mathbb{N}\ \exists q \in \mathbb{Q}^+\ \mathrm{Good}(x,y,\rho,q),$$

where $\mathscr{P}$ is a Borel family of admissible sampling/refinement policies. *That idealization is not an optional flourish*. It is how we make precise what physicists have always meant by "independence of the mesh," "discretization-invariance," "basis-insensitivity," and so on. Set-theoretic metatheory starts to matter at that Π^1_2 level. Under large cardinals strong enough for Projective Determinacy, the projective regularity needed to read our ensemble idioms is guaranteed, while under V=L it can fail even for simply definable sets. What follows tracks this divide in concrete cases and shows that talk of "almost all," "typical," and "canonical" is either well-posed or contentless depending only on whether we assume V=L or LC in the metatheory.

2. Three Regularity Hinges at Δ^1_2

Physicists speak comfortably about what holds *for almost all* initial data, or *typically* across a class of configurations. Those phrases read as harmless English, but they presuppose something precise: that the underlying definable sets of reals are *measurable* or have the *Baire* property. Whether that presupposition is warranted turns out to depend on the set-theoretic background.

[3] The metatheoretic sensitivity of determinism is not limited to ensembles, though they are my focus. See Section 7.3.

This section isolates the exact places where the background matters and explains why those places—what I will call the **regularity hinges**—are routine in physics.

The contrast begins with a simple construction. Within V = L there is a parameter-free (**lightface**) $\mathbf{\Delta^1_2}$ well-order $<$ of the reals. From $<$ one can define sets with pathological regularity. Let $W \subseteq [0,1]^2$ be given by $W = \{(x,y) : x < y\}$; then χ_W is also lightface Δ^1_2. In V = L one can arrange that horizontal sections of W have outer measure 1 while vertical sections have measure 0, a definable Sierpiński–Fubini pathology. By contrast, under large cardinals strong enough for Projective Determinacy (PD), every projective set is Lebesgue measurable and has the Baire and perfect-set properties. Thus, the very same Δ^1_2 definition denotes very different kinds of sets depending on whether one works in V = L or assumes large cardinals sufficient for PD (Devlin 1984; Jech 2003; Sierpiński 1924; Martin & Steel 1989; Kechris 1995; Moschovakis 2009).

From this starting point, three **regularity hinges** for **determinism** fall out naturally:

(1) Coherence (does the weak formulation make sense?)

In weak PDE formulations, coefficients must be **measurable** for Lebesgue integrals to exist. If a coefficient is of the form $\nu(x) = \nu_0 + \alpha \cdot \chi\_W(x)$, then under PD measurability is secured and the weak problem is meaningful. But under V = L one may choose W non-measurable so that the integral in the bilinear form fails to be defined, collapsing the weak formulation. (Evans 2010; Brezis 2011 corroborate the analytic side; the toggle is purely in regularity.)

(2) Uniqueness-as-genericity (may we say "for almost all" or "comeager"?).

Many results assert uniqueness not pointwise but for almost every datum or on a comeager set of parameters. Those ensemble claims are meaningful iff the relevant definable sets are measurable or Baire. PD guarantees these properties for projective sets. V = L may fail them, so the very grammar of "almost all" can break down at the projective level used by the argument (Kechris 1995; Moschovakis 2009).

(3) Identity (representation-independence / canonical selection)

Often the equations admit many admissible solutions, and practice demands a canonical choice that is invariant under admissible recodings (changes of units, gauges, coding conventions). Canonicalization is a measurable selection / uniformization problem for a definable multifunction. At low complexity (analytic graphs) ZFC already secures measurable selectors (Kuratowski–Ryll-Nardzewski; Castaing–Valadier). But once the construction genuinely reaches the projective level (as it will when we formalize "robustness across all admissible procedures"), PD is needed for measurable uniformization. Under V=L only definable but possibly non-measurable selections are guaranteed. The equations do not change. What changes is whether a representation-independent selection with ensemble meaning can be made.

Two clarifications are in order. First, none of this denies that much of analysis remains **absolute** across transitive models with the same reals. $\mathbf{\Sigma^1_2}/\mathbf{\Pi^1_2}$ truths with real parameters and membership in lightface Δ^1_2 sets are Shoenfield-absolute. The *regularity* (measurability/Baire) of certain projective sets diverges between PD and V=L. Second, I do not claim that genericity idioms are globally senseless in V=L. The point is that for the specific projective sets generated later by stability/UT predicates, measure and category can lack determinacy in V = L, while PD restores the regularity that those idioms presuppose. (Shoenfield 1961; Jech 2003, Thm. 25.20.)

With these hinges made explicit, the later sections will simply track which hinge is being touched in each example—**coherence** in weak formulations, **uniqueness-as-genericity** in ensemble claims, and **identity** in canonicalization—so the reader can see exactly where the set-theoretic background enters and why it matters.

3 Three Toy Schemata

The three regularity hinges introduced above—**coherence**, **uniqueness-as-genericity**, and **identity**—can be displayed most transparently in simplified, self-contained models. These "toy schemata" are not meant as new physical systems but as clean illustrations of how the same formal setup can yield distinct determinism profiles depending on whether one works in V=L or under large cardinals sufficient for PD. Each model isolates one hinge while keeping the others fixed, making the logical source of the metatheoretic toggle explicit.

The first, concerning **coherence**, shows how even the *meaningfulness* of a weak formulation can depend on regularity. Let $\Omega \subseteq [0, 1]^2$ be a bounded Lipschitz domain, and consider the bilinear form

$$\mathbf{a(u,v)} = \int_{(\Omega)} \mathbf{v(x)} \, \langle \nabla \mathbf{u(x)}, \nabla \mathbf{v(x)} \rangle \, \mathbf{dx},$$

with $\mathbf{v(x)} = \mathbf{v_0} + \boldsymbol{\alpha} \cdot \boldsymbol{\chi}\_\mathbf{W(x)}$, where $v_0 > 0$, $\alpha \neq 0$, and $W \subseteq [0, 1]^2$ is lightface Δ^1_2. Under large cardinals implying PD, χ_W is measurable, so $v \in L^\infty(\Omega)$ and the Lax–Milgram conditions are met; the weak form is coherent and yields a unique solution $u \in H_0^1(\Omega)$. Under V=L, however, one can choose W to be non-measurable, and then the Lebesgue integral defining a(u,v) is undefined in the ordinary sense. The statement of the equations fails. The same operator $-\nabla \cdot (v \nabla \cdot)$ with the same boundary data thus makes or fails to make analytic sense according to the metatheory (Evans 2010, §§ 5.8, 6.2; Brezis 2011, Ch. 9). Although none of the later, more physical examples—inverse problems, Markov uniqueness, Ising dynamics, decoherence, or Kerr interiors—produces non-measurable coefficients for any fixed instance, this toy example clarifies what a **coherence failure** means. The equations themselves become ill-defined in V=L. Unlike other toy examples, like Norton's Dome (Norton 2008; Malament 2008), whose indeterminism arises from the dynamics, the failure here emerges from the interaction between physical laws and set-theory. (For further discussion of the concept of determinism and what theories count as such, see Earman 1986, 2007, 2009; Fletcher 2012; Wilson 2009; Ismael 2016; Loewer 2001, 2020; Chen 2022; Butterfield 2007; Halvorson, Manchak & Weatherall 2025).

The second model, the **uniqueness schema**, keeps coherence secure but asks whether a law can single out a *unique* outcome "for almost all" admissible data. Under LC (sufficient for PD), projective relations arising from arg-min schemes admit measurable selectors. A map $s : H \to X$ can then be defined with $s(u) \in \Gamma\bigstar(u)$ for μ-almost every u, so ensemble statements like "for almost all u the continuation is unique" are meaningful. Yet measurability alone does not guarantee invariance. A measurable selector may still depend on the chosen coding. Invariance obtains only when the minimizer is structurally unique (through strict convexity or lower semicontinuity) or when the tie-break is explicitly restricted to respect the recoding group. In V = L, by contrast, measurable uniformization can fail at this projective complexity. Δ^1_2 selectors still exist by appeal to the lightface well-order of $\mathbb{R}$, but they need not be measurable, and then ensemble idioms such as "almost all," "typically," and "comeager" lose determinate content. What shifts is the status of the measurable law that assembles the solutions into an ensemble.

Finally, the **identity schema** adds one further layer. Even when the equations are well-posed and a measurable selector chooses an outcome for almost every case, we may still ask whether that choice is **representation-independent**—insensitive to any relabeling or recoding that carries no physical content. Let X be a Polish space of admissible states and $\Gamma \subseteq X \times X$ a definable relation whose fibers $\Gamma(x)$ are non-empty, compact, and convex. A Borel "secondary cost" $\Phi : X \times X \to \mathbb{R} \cup \{+\infty\}$ breaks ties among solutions, and a selector s satisfies $s(x) \in \Gamma(x)$. If $\mathscr{T}$ denotes the group of admissible recodings $T : X \to X$, then s is $\mathscr{T}$-invariant iff $s(Tx) = T(s(x))$ and $T(s(x)) \in \Gamma(Tx)$ for all T and x. Within ZFC, measurable selectors exist whenever the graph(Γ) is analytic (Kuratowski & Ryll-Nardzewski 1965; Castaing & Valadier 1977), but measurability alone does not ensure $\mathscr{T}$-invariance. True identity arises only through structural uniqueness or explicit symmetry requirements (Appendix C). When the graph of Γ reaches projective complexity (as it does once we impose the Π^1_2 "universal tameness" constraint UT introduced below) PD guarantees projective regularity and measurable uniformization (Kechris 1995, § 38.B; Moschovakis 2009, Ch. 6; Martin & Steel 1989). This restores the ensemble reading of "for μ-almost all x, the canonical continuation minimizes Φ." But PD alone does not force full invariance when multiple minimizers exist. Symmetry or structure does. To sum up: **coherence** makes the law intelligible, **uniqueness** makes the ensemble verdict meaningful, and **identity** makes the verdict representation-independent. LC implying PD secures the measurable scaffolding needed for the last two. Symmetry and geometry does the rest (Appendix C).

4. Inverse problems and imaging

Inverse problems make vivid the difference between the analytic layer, which ZFC secures, and the regularity layer, which is where metatheory enters. To keep things concrete, let us stay with the total-variation (TV) Tikhonov family,

$$\mathbf{J\_\lambda(u; d) = \|A\ u - d\|^2 + \lambda \cdot TV(u),}$$

with A linear and TV the usual anisotropic total-variation norm. For each fixed $\lambda > 0$, strict convexity (optionally reinforced by a tiny quadratic term) yields a unique minimizer u_λ(d) that depends continuously on the data d. All of this is secure in ZFC and independent of any set-theoretic speculation. But the choice of λ pushes beyond this. The discrepancy principle chooses

$$\mathbf{\lambda(d) = min\ \{\ \lambda \in \mathbb{Q}^+ : \|A\ u\_\lambda(d) - d\| \leq \tau \cdot \varepsilon\ \}.}$$

Because the acceptance set for each λ is closed and we minimize over a countable set, the map d ↦ u_{λ(d)}(d) is Borel. Therefore, ensemble statements such as "for μ-almost all d, the reconstruction has property P" have determinate sense in ZFC. The subtlety emerges when we enforce the *methodological ideal* that a reconstruction rule be robust across all admissible refinements—across discretizations, interpolation schemes, iteration schedules, and file formats that carry no physical content. Let us code each admissible refinement policy by a real number, $\rho \in \mathbb{N}^\mathbb{N}$ (this is the standard coding move used throughout numerical analysis: a single real names an entire infinite refinement schedule). Now say that R(d, λ) holds just in case

> for every ρ there exists m such that for all $n \geq m$,
> the refined reconstruction $u^{(\rho,n)}_\lambda(d)$ meets the tolerance $\|A\, u^{(\rho,n)}_\lambda(d) - d\| \leq \tau \cdot \varepsilon$,
> and no smaller rational λ has that property.

The logical form of the stability clause—"$\forall \rho \in \mathbb{N}^\mathbb{N}\ \exists m\ \forall n \geq m \ldots$"—is Π^1_2 with an arithmetical matrix. Minimality over $\mathbb{Q}^+$ is a bounded quantifier and does not raise the complexity. Thus, the predicate "R(d, λ) and λ is least" is Π^1_2 in d and λ. This is the first point where the metatheory matters. At Π^1_2, V=L and large-cardinal frameworks diverge on *regularity*. Under PD (secured by strong large cardinals), all projective sets are Lebesgue measurable and have the Baire property. Projective uniformization and measurable selection hold at this level. Under V=L, by contrast, we may still define a selection by appealing to the Δ^1_2 well-order of $\mathbb{R}$ (take the <-least λ with R(d, λ)). But such a selection need not be measurable. Hence, the grammar of "for almost all d" can lose determinate sense even though, for each individual d, the analytic reconstruction $u_{\lambda}(d)$ exists and is unique for every fixed λ.

If we now define the "universally stable" choice $\lambda^\dagger(d)$ to be the least rational λ with R(d, λ) and set the reconstruction to be $u_{\lambda^\dagger(d)}(d)$, V=L and LC sufficient for PD induce divergence. Under PD, the graph of the admissible-pair relation is projective and admits a measurable uniformization. Ensemble clauses such as "for μ-almost all d, the reconstruction by the universally stable rule has property P" are meaningful. By contrast, under a constructible background, only definable (indeed, lightface) selections are guaranteed and measurability may fail at precisely this Π^1_2 level. The *analytic* problem has not changed. What varies is whether the ensemble-level claim (the law-like generalization across all admissible refinements) has a determinate truth value. That is the *uniqueness-as-genericity hinge* in situ.[4]

5. Thin-barrier capacity, Markov uniqueness, and physically natural projective families

A second, more geometric arena in which the metatheoretic hinges surface is in the analysis of diffusions through "thin" barriers. Physically, these problems model the passage of heat, charge, or mass through a medium whose microscopic structure may include nearly insulating interfaces. Mathematically, they are governed by the Dirichlet form

$$\mathfrak{E}(u) = \int_{\Omega \setminus S} \langle a\nabla u, \nabla u\rangle \, dx,$$

where $\Omega \subseteq \mathbb{R}^n$ is bounded and Lipschitz, a(x) is measurable and uniformly elliptic, and $S \subseteq \Omega$ is a closed barrier. The associated operator $L\, u = -\nabla \cdot (a \nabla u)$ is said to be *Markov unique* if it admits exactly one closed, Markovian extension on $L^2(\Omega)$ (Ma & Röckner 1992; Fukushima, Oshima & Takeda 2011). The operator is Markov unique just in case $\mathrm{cap}_{1,2}(S) = 0$, and a standard

---

[4] Formally, with $\rho \in \mathbb{N}^\mathbb{N}$ coding an admissible policy and ϕ arithmetical, the stability clause has the form $\forall \rho \in \mathbb{N}^\mathbb{N}\ \exists m \in \mathbb{N}\ \forall n \geq m\ \phi(d, \lambda, \rho, n)$, which is Π^1_2 in (d, λ). Minimality over $\lambda \in \mathbb{Q}^+$ is a bounded quantifier and does not raise the complexity. Thus the set $S = \{ (d, \lambda) : R(d, \lambda)\ \&\ \lambda \text{ minimal} \}$ is Π^1_2, and its projection $A = \{ d : \exists \lambda\, R(d, \lambda)\ \&\ \lambda \text{ minimal} \}$ is Σ^1_2. Shoenfield absoluteness applies pointwise. The measurability of A and the existence of **measurable** uniformizations of S are where PD and V=L diverge.

certificate shows that $\mathrm{cap}_{1,2}(S) = 0$ iff $\exists (u_k)$ (rationally coded) with $u_k \geq 1$ near S and $\mathfrak{E}(u_k) \leq 2^{-k}$. Hence, the index set

$$\mathbf{U = \{ A : cap_{1,2}(S\_A) = 0 \}}$$

is analytic (Σ^1_1) for any reasonable coding $A \mapsto S\_A$. Ensemble statements like "for almost all A, Markov uniqueness holds" are, by themselves, unproblematic in ZFC. The regularity layer appears once we align the model with laboratory practice. In controlled deposition or filtered percolation, one tunes a threshold θ so that the blocked fraction converges to a target φ★ in the limit, and one expects this calibration to be robust across all admissible filtering/refinement policies. Let us code each policy by $\rho \in \mathbb{N}^\mathbb{N}$. For a given micro-environment code A, define θ_A to be the least rational θ such that

> **for every ρ there exists m with the coarse-grained coverage within 2^{-m} of φ★ for all later stages.**

Set $S\_A = \{ x : \phi\_A(x; \theta\_A) \geq \theta\_A \}$. The quantifier structure in $A \mapsto \theta\_A$ is, again, Π^1_2 (universal over reals, then existential over naturals), and so $A \mapsto S\_A$ sits at that projective level. For each fixed A, the event $\mathrm{cap}_{1,2}(S\_A) = 0$ is analytic in the code for S_A, but the *ensemble* we care about ranges over a Π^1_2-defined family of barriers. Under PD, all such projective sets are measurable, and projective uniformization yields measurable selections wherever we need them. The phrase "for almost all A" has determinate sense when applied to this Π^1_2 family. Under V=L, we still have definable (indeed lightface) ways to choose θ_A and hence S_A by appealing to the Δ^1_2 well-order. The problem is that measurability is not guaranteed at the Π^1_2 level. The upshot is the same hinge as in inverse problems. The PDE and the analytic proof of Markov uniqueness are unchanged. But the ensemble-level uniqueness claim toggles with V=L and LC.

6 Zero-temperature Ising: natural tie-breakers and definable versus measurable selection

It is natural to wonder if the dependence of determinism on set-theoretic background is confined to continuum systems. But it actually, even more accessibly, in a *completely discrete* model: the ferromagnetic Ising lattice at zero temperature under synchronous (parallel) updates. This case shows how the identity hinge can arise even when every object in play (the lattice, the configuration space, and the update rule) is Borel and unproblematic in ordinary analysis.

Fix $d \geq 2$, finite-range ferromagnetic couplings $J_{ij} \geq 0$ on $\mathbb{Z}^d$, and external field $h \in \mathbb{R}$. Configurations are $\sigma \in \{-1, +1\}^{\mathbb{Z}^d}$ with local field

$$\boldsymbol{\ell_i(\sigma) = \sum_{j\sim i} J_{ij}\, \sigma_j + h.}$$

The parallel zero-temperature update is

$$\mathbf{(U\_b(\sigma))_i = +1 \text{ if } \ell_i(\sigma) > 0; \quad (U\_b(\sigma))_i = -1 \text{ if } \ell_i(\sigma) < 0; \quad (U\_b(\sigma))_i = b_i(\sigma) \text{ if } \ell_i(\sigma) = 0,}$$

with b a local tie-break. For any local measurable b, $U\_b : \mathfrak{X} \to \mathfrak{X}$ is Borel on $\mathfrak{X} = \{-1, +1\}^{\mathbb{Z}^d}$, so analytic questions about existence and "μ-almost sure" properties are clear in ZFC.

Identity (in the sense of Section 2) is where the metatheory enters. Suppose we insist, as we should, that the tie-break be invariant under admissible recodings—lattice symmetries, relabelings of sites, and shape conventions for van Hove limits—and that it be stable across every admissible refinement policy $\rho \in \mathbb{N}^\mathbb{N}$ that specifies how the limit is taken. Selecting a canonical update τ with those invariance and stability properties is naturally formulated as a Π^1_2 uniformization problem. The stability clause has the "$\forall \rho\ \exists m\ \forall n \geq m \ldots$" shape, and the invariance clauses are Borel constraints under the recoding group $\mathscr{T}$. In frameworks with projective regularity (like LC sufficient for PD), measurable uniformizations exist at this level, and we can say "for product-measure almost all σ, a canonical, representation-independent update $\tau(\sigma)$ is defined." What measurable uniformization does not do, by itself, is force *uniqueness* of τ when there are symmetries in the fiber. For identity in the strong sense we still need either structural uniqueness or an explicitly symmetric tie-break rule. In constructible frameworks, on the other hand, one can define a Δ^1_2 tie-break using the lightface well-order. And, yet, measurability—and hence the literal sense of the ensemble idioms—can fail at this Π^1_2 complexity. So, again, the dynamics U_b is the same. But whether law-like canonicalization has a measurable, representation-independent realization, wobbles between V=L and LC.

7 Metatheory in Fundamental Physics: the analytic and regularity layers

The preceding sections have shown that even ordinary systems (imaging, diffusion, spin dynamics) reveal subtle dependence on set-theoretic background once we ask ensemble or canonical questions. But one could still wonder whether such dependence touches paradigmatic *fundamental* physics. This section argues that it does. The examples come from the two main pillars of modern physical theory: quantum mechanics and general relativity.

Physicists often assume that determinism is a purely analytic matter. One asks whether the relevant equations admit a well-posed initial-value problem (whether a unique solution exists and depends continuously on the data). That analytic layer of the theory is indeed captured in ZFC. The function spaces, continuity estimates, and energy inequalities that define well-posedness are absolute between transitive models sharing the same reals. But physical practice involves more than existence and uniqueness. When we speak of *generic initial data*, *typical environments*, or *canonical continuations*, we appeal to ensemble or uniformization claims that presuppose regularity (measurability, the Baire property, or the existence of measurable selections for definable multifunctions). Those assumptions can fail in the constructible universe yet hold under large cardinals strong enough for projective determinacy.

This division marks a boundary between *two layers of physical reasoning*. The **analytic** layer covers what the equations literally say. The **regularity** layer enters when we elevate pointwise results to law-like generalizations. The logical form that captures this boundary can be written as

$$\mathbf{UT(x, y)} :\Leftrightarrow \forall \boldsymbol{\rho} \in \mathbb{N}^\mathbb{N}\ \exists \mathbf{q} \in \mathbb{Q}^+\ \mathbf{Good(x, y, \rho, q)},$$

where UT stands for "universal tameness" or "universal stability," ρ ranges over a Borel set of admissible sampling or refinement procedures, and Good(x, y, ρ, q) is an arithmetical check expressing stability within tolerance q. This quantifier pattern—universal over reals followed by existential over rationals—has Π^1_2 complexity. Note that *this is not an artificial device*. It simply formalizes a familiar scientific ideal of robustness, that results remain invariant across all admissible ways of probing or discretizing the system. We have seen, however, that Π^1_2 is the level at which large-cardinal and constructible frameworks begin to diverge on regularity.

7.1 Decoherence and the preferred basis

Among the most discussed questions in quantum foundations is the *preferred-basis problem*: why do certain bases of a system's Hilbert space behave classically under environmental coupling? Decoherence theory and the program of "quantum Darwinism" (Zurek 1981, 2003; Schlosshauer 2007; Ollivier, Poulin & Zurek 2004; Riedel, Zurek & Zwolak 2014) answer that question dynamically, treating basis selection as a physical process. Whatever the other problems with this view, the present point is that the idiom of the literature—"for almost all environments," "for typical interactions"—is ensemble, and thus falls within the regularity layer.

Let $\mathcal{H}$ be a finite-dimensional, finite-energy subspace, so that $U(\mathcal{H})$ is compact. Let $\mathcal{B}$_adm $\subseteq U(\mathcal{H})$ be a non-empty compact set of admissible orthonormal bases, such as those compatible with measurement resolution. Let $\mathcal{S}$ be a Polish space of admissible environment states, for example thermal or Gaussian mixtures with bounded energy. For each B $\in$ $\mathcal{B}$_adm and ρ_E $\in$ $\mathcal{S}$ define the decoherence functional

$$\Phi[B; \rho\_E] = \int_0^T E\_\{\rho\_E\}(\| \text{offdiag\_B}(U\_t\ \rho\_SE\ U\_t\dagger) \|_1)\ dt.$$

Continuity of t ↦ U_t and of the Schatten-1 norm ensures that B ↦ Φ[B; ρ_E] is continuous, hence each arg-min set Γ_0(ρ_E) = argmin_{B∈$\mathcal{B}$_adm} Φ[B; ρ_E] is compact with analytic graph. This is all absolute in ZFC. Ensemble statements such as "for μ-almost all ρ_E the minimizing basis is unique" have determinate sense within standard mathematics.

The regularity issue arises when we enforce the routine scientific ideal that the preferred basis be *robust across all admissible probing procedures*. Let $\mathcal{P} \subseteq \mathbb{N}^\mathbb{N}$ encode those admissible measurement or coarse-graining policies. Say that B is universally stable for ρ_E just when

$$\forall \rho \in \mathcal{P}\ \exists q \in \mathbb{Q}^+\ \text{upper-density\_}\rho\{\ t : \|\text{offdiag\_B}(U\_t\ \rho\_SE\ U\_t\dagger)\|_1 > q\ \} = 0.$$

This requirement that decoherence results stabilize no matter how we sample or coarse-grain gives UT(B, ρ_E) the Π^1_2 form above. If we refine the arg-min relation by a secondary Borel cost Φ_1 penalizing small residual coherence, we obtain

$$\Gamma\bigstar(\rho\_E) = \text{argmin\_}\{B \in \Gamma_0(\rho\_E)\}\ \Phi_1[B; \rho\_E]\ \text{subject to UT}(B, \rho\_E).$$

The graph of Γ★ is projective at level Π^1_2. In settings with projective regularity (guaranteed by large cardinals) measurable uniformizations exist, providing universally measurable selectors ρ_E ↦ B★(ρ_E) and making ensemble statements such as "for almost all ρ_E there exists a measurable canonicalization ρ_E ↦ B★(ρ_E), and when the secondary cost plus symmetry constraints fix a single B★, the resulting canonical basis is representation-independent" meaningful. In a constructible setting of V=L, however, only definable but not necessarily measurable tie-breakers are ensured. So, rival definable choices can disagree on non-measurable Δ^1_2 sets. The physical Hamiltonian and environment are identical in both worlds. What changes, as before, is whether the ensemble-level claim about canonical basis selection is well-posed.

7.2 Determinism inside Kerr black holes

An analogous boundary appears in general relativity's central open problem, the **Strong Cosmic Censorship** conjecture. In subextremal Kerr ($0<|a|<M$), the Cauchy horizon $\mathscr{CH}^{+}$ is classically the place where determinism threatens to fail. Analytic results establish linear estimates compatible with finite weighted flux across $\mathscr{CH}^{+}$, together with blue-shift amplification that can drive late-time instabilities. The analytic layer I need can be stated abstractly as follows.

Fix a smooth spacelike Σ_0 meeting the event horizon and set $\mathfrak{D}=H^1(\Sigma_0)\times L^2(\Sigma_0)$, coded as a Polish space. For $d\in\mathfrak{D}$, let $\varphi\langle d\rangle$ solve $\Box_g\varphi=0$ in the interior. Along the ingoing horizon parameter v, define the weighted flux

$$\mathscr{F}(d;\, v_0, v_1)=\int_{v_0}^{v_1} e^{\kappa v}\, \|\partial_v\varphi\langle d\rangle(v,\cdot)\|^2_{L^2\_\theta}\, dv, \quad \mathscr{F}(d)=\liminf_{v_1\to\infty}\mathscr{F}(d;\, v_0, v_1).$$

Let $\mathfrak{A}=\{\, d: \mathscr{F}(d)<\infty \,\}$. For $d\in\mathfrak{A}$ the traces of $\varphi\langle d\rangle$ along $\mathscr{CH}^{+}$ exist in L^2_loc along sequences $v_1\to\infty$ with uniformly bounded flux; write $\Gamma(d)\subseteq L^2\_\mathrm{loc}(\mathscr{CH}^{+})$ for the set of such limit traces. Then $\Gamma(d)$ is non-empty and sequentially compact in the L^2_loc topology; the **graph of Γ is analytic**. Quantify late-time behaviour by

$$E_0(d,x)=\int_{v_0}^{\infty} e^{\kappa v}\, \|\partial_v x(v,\cdot)\|^2_{L^2\_\theta}\, dv,$$

and fix a **Borel secondary cost $\Phi_1(d,x)$** penalizing late-time leakage. Lower semicontinuity of E_0 on $\Gamma(d)$ yields non-emptiness of $\Gamma_0(d)=\arg\min_{x\in\Gamma(d)} E_0(d,x)$; refining by Φ_1 gives $\Gamma\bigstar(d)=\arg\min_{x\in\Gamma_0(d)} \Phi_1(d,x)$. Up to here, everything is analytic and stable in ZFC.

Where the metatheory enters is the robustness ideal already familiar from earlier sections. Let $\mathscr{P}\subseteq\mathbb{N}^\mathbb{N}$ be a fixed Borel set of admissible sampling policies along $\mathscr{CH}^{+}$ (subsequences, coarse-grainings, windows). Say that $x\in\Gamma(d)$ is **universally tame** at d iff

$$\mathrm{UT}(d,x):\Leftrightarrow\ \forall\rho\in\mathscr{P}\ \exists q\in\mathbb{Q}^{+}\ \mathrm{Good}(d,x,\rho,q),$$

where Good is an arithmetical check formalizing "late-time stability within tolerance q" under ρ. The admissible-pair relation

$$R(d,x):\Leftrightarrow d\in\mathfrak{A},\ x\in\Gamma\bigstar(d),\ \text{and}\ \mathrm{UT}(d,x)$$

has $\mathbf{\Pi}^1_2$ complexity in (d,x). Under projective regularity (e.g., **PD**), R is Lebesgue-measurable and admits projective (hence universally measurable) uniformizations: there exists a measurable selector s with $s(d)\in\Gamma\bigstar(d)$ and $\mathrm{UT}(d,s(d))$ for μ-almost every d in any physically reasonable ensemble μ on $\mathfrak{D}$. Ensemble claims such as "for almost all d there is a **canonical** continuation selected by (E_0,Φ_1,UT)" are then literally meaningful. In a constructible background, by contrast, $\mathbf{\Delta}^1_2$ uniformizations can be defined via the lightface well-order but need not be measurable; the very grammar of "for almost all d" with respect to R can then lose determinate sense. The wave equation and the blue-shift estimates are identical in both worlds; what toggles is whether the canonicalization ideal that **SCC** presupposes can be realized measurably.

7.3 Beyond quantum and relativity: a string-theoretic analogue

An analogous structure appears in **AdS/CFT** bulk reconstruction. There one minimizes an energy-like functional over a definable admissible set of bulk fields compatible with given boundary data, then breaks ties by imposing stability on slices—an analogue of the Φ_1 cost. The arg-min graph is again projective. Frameworks with projective regularity provide measurable selectors and well-defined ensemble semantics. But constructible frameworks allow only definable ones, and different definable selectors can disagree on Δ^1_2 disagreement sets. The metatheoretic hinge thus reappears without contrivance in the logic of holographic dualities.

7.4 Single-system toggling

Up to now the examples have concerned ensembles—statements of what happens for "almost all" environments or data. But it should not be thought that the dependence is limited to such caees. It can arise even for a single, fully specified system. In the Kerr case, fix one admissible datum $d \in \mathfrak{A}$ for which $\Gamma_0(d)$ contains multiple mild continuations. The analytic facts—compactness of $\Gamma_0(d)$, existence of minimizers—are absolute. *But when canonicity is required to be representation-independent, distinct metatheories may select different canonical continuations.* Frameworks with projective uniformization allow measurable tie-breakers that can be constrained by symmetry. Constructible frameworks supply only definable ones, which need not be measurable or symmetry-respecting. Determinism, understood as the uniqueness of a canonical continuation, can thus vary across metatheories **even for the same physical data**.

One might object that in practice robustness is checked only across finitely many procedures, not *all* admissible ones. That is correct. The universal-tameness clause idealizes the empirical norm of robustness into a demand of invariance under every admissible policy. But this idealization is just a reflection of the fact that the world does not care about our experimental limitations, not an artificial add on. This lifts the statement to Π^1_2 complexity. The consequence is that if determinism is defined through such a robustness ideal, then even the determinism profile of a **single physical system** can differ across set-theoretic backgrounds. **Whether the universe inside a black hole has a unique canonical future can, accordingly, depend on speculative extensions of ZFC.**

7.5 Summary

Across this paper's examples, the pattern is consistent. The analytic layer of physics—the equations and their pointwise consequences—lives within ZFC. The regularity layer—the assumptions that make ensemble or canonical claims well-defined—does not. When projective regularity holds, measurable uniformizations exist and the ensemble language of physics has clear sense. When it fails, those same phrases lose meaning. The content of physical determinism depends not only on the equations we write but on the set-theoretic background that determines which ensemble and canonical statements they can meaningfully express.

This division between the analytic and regularity layers also reframes long-standing debates about determinism in the philosophy of physics, especially those articulated by Earman and Butterfield. Their analyses locate the determinism question at the analytic layer: whether the equations governing a theory are *well-posed* in the Hadamard sense (existence, uniqueness, and continuous dependence of solutions on initial data). On that account, determinism fails when the analytic problem is ill-posed, as in Norton's Dome or non-Lipschitz systems, but succeeds whenever the analytic conditions hold. The present analysis shows that this framework is to crude. **Even when the analytic layer is well-posed, the *regularity layer*—the measurable, ensemble, and canonical structures that allow us to interpret statements like "for almost all initial data" or "for the canonical continuation"—can depend on the set-theoretic metatheory.** The Earman–Butterfield picture thus captures only one axis of determinism, the analytic one, while physical practice presupposes another as well, the *regularity axis*. **What looks like a deterministic theory under one metatheory (e.g., large cardinals sufficient for PD) can, under another (e.g. V=L), lose the measurability or uniformization that gives those analytic results their physical meaning.** The familiar analytic notion of determinism is stable across models of ZFC. But the physicist's ensemble and canonical notions are not.

8. Objections and Replies

Objection 1: Physicists Don't Use Δ^1_2 Sets

I have argued that, contrary to widespread belief, the determinism profile of physical systems, from routine inverse problems to black hole dynamics, hinges the choice of speculative axiom candidates extending standard mathematics. A natural reaction is that the constructions must be artificial. Physicists never work with Δ^1_2 sets or discuss the projective hierarchy.

This objection, however, confuses explicit terminology with implicit structure. Consider the types of conditions physicists routinely impose: "There exists a sequence of approximations converging to a solution such that for every accuracy requirement, all sufficiently late terms meet that requirement." This has the logical form $(\exists f)(\forall n)(\exists m)(\forall k{\geq}m)\ P(f,n,k)$, which is a Σ^1_2 statement when P is arithmetic. The mathematical structures that appear in mature physical theories naturally involve alternating quantifiers over infinite objects. The Kerr interior predicates of C^0-extendibility, finite flux, and blow-up are core concepts from the general relativity literature. These predicates clearly have low projective complexity when spelled out.

Moreover, it is not clear what this objection would show even if it were true. Consider the pure mathematical case. Just as the Paris-Harrington and Goodstein theorems reveal that seemingly elementary combinatorial statements can be independent of Peano Arithmetic, this paper reveals that **physical determinism—a basic property we expect theories to either have or lack—can be undecidable in ZFC**. The significance of such results lies not in their frequency (which depends on what questions we happen to have asked) but in their existence. The constructions here are no more 'artificial' than the Ramsey numbers in Paris-Harrington or the specific Diophantine equations encoding undecidability in Hilbert's tenth problem. If anything, they are more natural. They spell out questions that physicists have already asked, rather than constructing sentences never considered and arguing that these are "relevantly similar" to naturally occurring examples (as in Friedman 1998). The indeterminacy dsicsussed concerns *discovered* features of the physical-mathematical landscape, not pathologies that I have created.

Objection 2: This Confuses Mathematics with Physics

Another objection is that I am conflating the mathematical formalism with physical reality, the "map with the territory". Mathematics is merely a "language" for describing physics; changing the language does not change the physics itself.

But this is exactly what is at issue! A physical theory does not exist independently of its mathematical expression. It is partly *constituted* by that expression. When we say that general relativity predicts gravitational waves, we mean that the Einstein field equations, interpreted according to standard mathematics, entail the existence of wave solutions. The field equations purged of all metatheoretic commitments are just uninterpreted marks. If different mathematical frameworks yield different solutions to the same equations, then the physical theories differ.

Consider again Strong Cosmic Censorship. This asserts that Einstein's equations typically do not admit smooth extensions beyond the Cauchy horizon. Whether this is true depends on what "typically" means—a measure-theoretic or category-theoretic notion that requires the regularity properties toggled by V=L versus PD. The *physical* question "Does the universe inside a black hole have a unique future?" receives different answers in different *set-theoretic* contexts.

Objection 3: Why Not Just Adopt Large Cardinals?

A pragmatic response might be to just adopt large cardinal axioms, thereby ensuring all projective sets are regular and eliminating the ambiguity.[5] If LC (sufficient for PD) makes all the pathologies disappear, then why not just assume it?

There are two principal reasons. First, adopting LC would not eliminate the phenomenon. It would push it up the projective hierarchy. At the next level—Σ^1_3 sets—V=L and LC can disagree not just about regularity but about *membership*. (The identity questions we encountered in the Ising model (Section 6) already hint at this higher-level divergence.) Recall that Shoenfield's Absoluteness Theorem does not apply this high up (**Appendix B**). So, appeal to regularity intuitions is insufficient. Second, this response amounts to choosing our physics based on convenience rather than observation. Whether a fluid flow is turbulent or a black hole interior is deterministic certainly *seems* like a fact about the world, not something we can stipulate.

Objection 4: This Has No Experimental Consequences

Perhaps the most obvious "objection" is that these set-theoretic subtleties, while perhaps mathematically interesting, have no experimental implications. The problematic sets have measure zero or are meager. The disagreements occur only on negligible sets that we do not encounter in practice.

But this objection is facile. First, for discrete systems like the Ising model, the disagreement sets need not be negligible. In V=L, we can in principle arrange the set M where different implementations disagree to have positive outer measure with respect to the product measure induced by the coding of configurations. In principle, preparing an ensemble of magnetic systems and observing their evolution could reveal statistical signatures of such disagreements (for more on possible experimental tests of each of the examples, see **Appendix F**.). Second, even if direct experimental tests are currently infeasible, I have emphasized, following Quine (1951), that the content of our physical theories depends on the set-theoretic background.[6] There is no meaningful parition. When we ask whether general relativity predicts time dilation near massive objects, we are asking whether a theory incorporating substantial set theory does. The

[5] Thanks to Douglas Blue for suggesting this line of response.
[6] I do not mean to concede that the notion of "direct" confirmation is clear.

“corporate body” does the predicting, and contra Sober (1993), there are no God given axioms of mathematics.

Finally, the inaccessibility of black‑hole interiors does not diminish the meaning or importance of the question. A lesson of the mid‑century critique of verificationism is that unobservable structure can be indispensable to science. What matters is what our best‑supported theories, together with mathematics, say the world would be like. Strong Cosmic Censorship is a paradigm. It asks whether generic Cauchy data admit a unique, canonical continuation past the inner horizon—i.e., whether the theory is deterministic there. *Whether this statement is even well‑posed—whether its “generic” and “canonical” clauses have determinate content—depends on the set‑theoretic metatheory*. Observable or not, there is a fact and we should know it.

Objection 5: Real Physics Uses Approximations

A final objection (already aluded to in Section 7) notes that physical measurements invariably involve finite precision and ε-approximations. No experiment can distinguish whether a real number equals one value or another to infinite precision. Perhaps the set-theoretic subtleties only affect an idealized theory that was never physically meaningful.

This objection mixes up epistemology and metaphysics, much like the last. One question is what the world is like. Another is what we can measure. Determinism is about what the world is like, not what we can measure (or predict). Indeed, if we must retreat to "deterministic up to ε" for all $\varepsilon > 0$, then the idealized theory is *not deterministic*! The question under investigation is whether the theory, not our currently best finite-precision measurements of it, has unique solutions.

Coincidentally, questions about limits as $\varepsilon \to 0$ are themselves mathematical questions that can be sensitive to set-theoretic background (see **Appendix E**). Whether there exists an ε_0 such that all solutions agree within ε_0, or whether the minimal divergence shrinks to zero, depends on which real numbers exist. This is the kind of issue over which V=L and LC may disagree.

9 Directions for future research

The examples developed here are only the beginning of a new research program investigating how physical theories depend on mathematical foundations. In analogy with Friedman’s and Simpson’s “reverse mathematics,” we might call it **reverse physics**. The goal is to map *how our choice of foundational axioms affects the behaviour of our most trusted physical frameworks*

Several directions deserve immediate exploration.

Beyond measurability. This paper focused on Lebesgue measurability because it marks the first point where large-cardinal and constructible frameworks diverge. But projective determinacy also guarantees the Baire and perfect-set properties, and many arguments in stability, bifurcation, and ergodic theory tacitly rely on these topological forms of regularity. Identifying physical systems whose qualitative behaviour (phase bifurcations, attractor formation, or chaotic transitions) depends on these subtler properties would deepen our understanding of how mathematics and physics interlock.

Higher complexity. Again, at the levels beyond Σ^1_2, the two metatheories can disagree not only about regularity but about *membership*. At the Σ^1_3 level, for example, they can assign different truth values to whether a given set of reals exists at all. This opens the possibility of physical systems whose discrete outcomes—which particle decays occur, which crystal configuration forms—depend on which mathematical universe we inhabit.

Quantum field theory. Several standard constructions in QFT already live high in the projective hierarchy. The Hadamard condition for physically admissible states, the definition of vacua in curved spacetime, and the renormalization and subtraction schemes of perturbative QFT all involve selection principles over definable function spaces. Understanding how these depend on background axioms could clarify deep ambiguities about state-selection and renormalization across inequivalent representations.

Statistical mechanics. The construction of Gibbs measures by thermodynamic limits, the existence of phase transitions, and many ergodic and mixing properties rely on limits that climb the projective hierarchy. Even the notion of equilibrium—the convergence of ensemble averages to time averages—may turn out to be metatheory-dependent. Exploring whether equilibrium statements retain determinate sense in different set-theories could reveal a new conceptual layer in statistical physics.

Computational complexity. Beyond questions of existence and uniqueness, one can ask about effective computability. Whether a system's evolution is algorithmically decidable or whether its long-term behaviour is recursively enumerable can depend on which sets of reals one takes to exist. This connects the study of physical determinism with the growing field of *physical computation theory*, where limits of calculability are tested against the physical resources allowed by fundamental laws.

Alternative foundations. Finally, these phenomena should be compared across other foundational settings—constructive mathematics, topos theory, homotopy type theory, and alternative set theories such as NF, ZF + AD, or Kripke-Platek. Such comparisons would help separate what is genuinely physical from what is an artefact of the particular set-theoretic framework we have chosen – or, more likely, in my view, to reveal that there is no distinction between the two.

10. Conclusion: Determinism After Foundations

Physicists and philosophers of physics commonly assume that disagreements in the foundations of mathematics—such as whether there are large cardinals or V=L—have no bearing on physical theory. I have argued otherwise. The examples surveyed demonstrate that the Axiom of Constructibility (V=L) and large cardinal (LC) axioms strong enough to prove Projective Determinacy (PD) can diverge on all the key aspects of determinism. But unlike Norton's Dome and similar thought experiments, they are not just hypothetical. The examples here include routine systems from across mathematical physics, both continuous and discrete. They also include frontier questions in the foundations of quantum mechanics and general relativity.

Note that the point has *not* been the obvious one that we need to supplement standard axioms in order to decide questions that we could not within ZFC alone (as with, say, ZFC + Con(ZFC)).

That is an easy application of Gödel's theorems (Gödel 1931), assuming that those axioms are consistent (and recursively axiomatizable). V=L and LC are seriously entertained extensions of ZFC (Maddy 1997; Steel 2014; Koellner 2009). We have to take a stand on them like the stand that pioneers of set theory took with respect to the Axiom of Choice (Jech 2003, Ch. 1; Kanamori 2009, §0). The question of which of V=L or LC is true—or whether it even makes sense to say that one of them is true—is not a question that admits of proof. As Jensen writes:

> Most proponents of V=L and similar axioms support their belief with a mild version of Ockham's razor. L is adequate for all of mathematics; it gives clear answers to deep questions; it leads to interesting mathematics. Why should one assume more? ... I do not understand ... why a belief in the objective existence of sets obligates one to seek ever stronger existence postulates. Why isn't Platonism compatible with the mild form of Ockham's razor? ... I doubt that one could, with the sort of evidence I have, convert the mathematical world to one or the other point of view. Deeply rooted differences in mathematical taste are too strong and would persist (Jensen 1995, 400–401).

One reaction to this situation is to hold that, despite widespread opinion to the contrary, Quine was right, and the search for new axioms is not different in kind from the search for new physical laws (Quine 1951, 1990). He writes, "sentences such as the continuum hypothesis...which are independent of [standard] axioms, can...be submitted to the considerations of simplicity, economy, and naturalness that contribute to the molding of scientific theories generally. Such considerations support... [the] Axiom of Constructibility, 'V=L'" (Quine 1990, 95). This position raises familiar quandaries. Set theorists are not responsive to experiment even to the extent that most theoretical physicists are (Maddy 1997, 155; see also Maddy 2011). But this much is at least true: when we say that general relativity predicts *X*, we mean that *X* follows from the field equations plus boundary conditions, where "follows from" is a relationship of logical entailment within a mathematical framework. If that logical entailment depends on which set-theoretic axioms we assume then the theory's content is metatheory-sensitive. This is not a defect of our measurements or approximations. It is a fact of theorizing. Recognizing this clarifies the relationship between mathematical and physical reasoning, showing that they are allied in a way that non-Quineans have no sufficiently appreciated (see Resnik 1997, Ch. 10; Colyvan 2001).

My own view (Clarke-Doane 2020, 2024, 2025) is that no sense can be made of the claim that either V=L or LC is "really right". Each affords a legitimate arena in which to carry out mathematical reasoning—broadly like Euclidean and non-Euclidean geometries (see also Balaguer 1998; Hamkins 2012). The heady difference is that all the geometries can live in a single set theory. But any set theory takes itself to be maximal, to govern *everything*. This raises unresolved questions about how to formulate the "monism–pluralism" debate (Clarke-Doane 2025, §1), and whether it is even factual (cf. Carnap 1950).[7] But if one *is* a set-theoretic pluralist,

[7] Here is a brief argument that there can be no factual question in dispute under the heading the "monism-pluralism" debate. Any precise statement of either view must be made *within* some background theory that fixes what counts as a universe of sets. Inside that framework, "V" denotes everything the framework recognizes, so the monist's claim that there is one true V is trivially true *by that framework's own lights*. But taken *independently* of any background the claim cannot be formulated without choosing another theory to interpret its quantifiers, and different choices yield different contents. The same holds for pluralism. Its "many universes" are sets or proper classes *in one background*, not an external plurality that both sides could recognize. Because truth and reference for statements about V shift with the metatheory, the alleged disagreement's truth value changes with background moves that

then, by the arguments given here, one must be a pluralist about core physical concepts too. ("Concepts" rather than "concept" because if determinism is relative to set-theoretic foundations, then so too presumably is physical **necessity**, **prediction**, **explanation**, **causation** and more.) As Quine writes in a different context, "Carnap has maintained that [the question of which axioms hold] is a question not of matters of fact but of choosing a convenient language form, a convenient conceptual scheme or framework for science. With this I agree, but only on the proviso that the same be conceded regarding scientific hypotheses generally (1951, XX)" The question of whether a theory is deterministic along a given dimension is, on the present view, either misconceived or practical—*whether to assume* V=L or LC for the purpose at hand. Of course, this position should not be taken lightly. But it must be reckoned with given the independent plausibility of set-theoretic pluralism, and the entanglement of math and physics.

Quantum-gravity theorists sometimes remark that we may need "new math" to formulate a final theory (Ashtekar & Lewandowski 2004; Polchinski 1998). I have argued that this "new math" may go much deeper than anticipated—not just new tools within a familiar framework, but new foundational frameworks. Future progress in physics may therefore depend on novel interactions between physics and the foundations of mathematics.

Appendix A. The Projective Hierarchy and PD

For sets of naturals, $\Sigma^0_\square/\Pi^0_\square/\Delta^0_\square$ alternate first-order quantifiers (Rogers 1987). For sets of reals, $\boldsymbol{\Sigma}^1_1$ **(analytic)** are projections of Borel sets; $\boldsymbol{\Pi}^1_1$ **(coanalytic)** have analytic complements; $\boldsymbol{\Delta}^1_\square = \boldsymbol{\Sigma}^1_\square \cap \boldsymbol{\Pi}^1_\square$. **PD** (from strong large cardinals) implies every projective set is **Lebesgue measurable**, has the **Baire** and **perfect-set** properties (Martin & Steel 1989; Kechris 1995, §38.B; Moschovakis 2009, Ch. 6). Analytic sets are regular already in ZFC; coanalytic need not be.

Appendix B. Absoluteness and Regularity

Shoenfield absoluteness: Σ^1_2/Π^1_2 truths with real parameters are absolute between transitive models of ZFC with the same reals (Shoenfield 1961; Jech 2003, Thm. 25.20). Membership in **lightface** $\boldsymbol{\Delta}^1_2$ sets is absolute; **regularity** (measurability/Baire) can diverge in V=L vs PD.

Appendix C. Coding invariance

This appendix gathers modest but essential facts about **coding invariance** that the main text presupposes. The guiding idea is that a single physical "data class" or "state class" can be represented in many mathematically equivalent ways—by different rational Cauchy-name schemes, file formats, grids, or enumerations. The **content** of claims about definability, measurability, "almost all," and canonical selection should not depend on which such coding we choose. What follows records the assumptions we make about recodings and proves that, under these assumptions, the definability and ensemble-level claims in §§ 2–7 are invariant across choices. At the same time we do **not** claim that *measurability itself* implies *coding-invariance of the particular selector*. Identity in the strong, representation-independent sense still requires symmetry or structural uniqueness in the selection rule.

---

should not matter if a single factual issue were in play. What remains is, if not a verbal question about what we happen to mean by "set", a *pragmatic* question—*which framework to use*—not a factual one about what exists.

C.1. Setting and standing assumptions

Fix two standard Borel spaces of codes $\mathfrak{X}$ and $\mathfrak{Y}$ for the same underlying mathematical objects (e.g., two rational Cauchy-name schemes for L^2-data, two ways of coding orthonormal bases, two encodings of spin configurations). A **recoding** is a bijection T: $\mathfrak{X} \to \mathfrak{Y}$ whose graph and inverse graph are Borel. Such T are Borel isomorphisms. When we treat the state space (e.g., reconstructed images, continuations, bases) as coded, we write $\mathfrak{S}$ and $\mathfrak{S}'$ for two standard Borel code spaces and U: $\mathfrak{S} \to \mathfrak{S}'$ for a Borel isomorphism. Call (T,U) an **admissible recoding** when:

> (i) T and U are Borel isomorphisms;
> (ii) for any fixed σ-finite **Borel measure** μ on $\mathfrak{X}$ representing a physical probability law or ensemble (noise models, product measures, Gaussian/Gibbs states), the **push-forward** T_*μ is the measure used on $\mathfrak{Y}$;
> (iii) when the body of the paper uses **Baire category** rather than measure (e.g., "comeager"), we require T to be a **Baire isomorphism** for the chosen Polish topologies: there are dense G_δ sets C⊆$\mathfrak{X}$ and C′⊆$\mathfrak{Y}$ such that T: C→C′ is a homeomorphism and maps meager sets to meager sets and comeager sets to comeager sets;
> (iv) if passing between codings requires fixed **calibration parameters** a (gauge choices, discretization furniture), we treat all pointclasses **relative to a**, i.e., $\Delta^1_2(a)$, $\Sigma^1_{\Box}(a)$, $\Pi^1_{\Box}(a)$.

These assumptions are satisfied in the applications. Rational Cauchy codings, dyadic-grid file formats, and basis-codes are standard Borel. Natural "regrid," "reencode," and "rename" maps are Borel and measure-class preserving. Where "comeager" is used, the recodings are homeomorphisms on dense G_δ cores (as in the usual passage between equivalent Cauchy-name systems). Nothing here presumes **effectivity**. Names can code non-computable reals, as they do in the main text's examples.

C.2. Pointclass and regularity robustness under recoding

The first and most basic claim is that **definability** does not depend on coding. Let S⊆$\mathfrak{X}$ be defined by a formula at some projective complexity (possibly relative to parameters a). Because Borel isomorphisms preserve projective pointclasses (pullbacks and push-forwards of projective sets remain projective at the same level, uniformly in parameters), the image T[S]⊆$\mathfrak{Y}$ has the **same pointclass** as S, and conversely. In particular, **lightface $\boldsymbol{\Delta^1_2}$** sets and relations, as used throughout § 2 and in the toy schemata of § 3, remain Δ^1_2 when transported across codings.

Two corollaries matter in the paper. First, **Shoenfield absoluteness** is unaffected by recoding. Σ^1_2/Π^1_2 truths with real parameters (hence Δ^1_2 membership) remain absolute between transitive models of ZFC with the same reals *regardless of which code space we use*. Second, when PD is assumed, the **regularity** of projective sets (Lebesgue measurability, Baire property, perfect-set property) is invariant under recoding, because these are properties of the underlying sets of reals and of their projective definitions, not of any particular naming convention. So, wherever the text invokes projective regularity under PD, or highlights its failure under V=L, this is **code-robust**.

C.3. "Almost all" and "comeager" are code-invariant

Ensemble statements have the schematic form: "for μ-almost all $d \in \mathfrak{X}$, property P(d) holds." Under an admissible recoding T, the **truth value** of such a statement is unchanged when we pass to $\mathfrak{Y}$ and the push-forward measure T_*μ. Indeed, $A \subseteq \mathfrak{X}$ is μ-null iff $T[A] \subseteq \mathfrak{Y}$ is T_*μ-null.* Likewise, a function f: $\mathfrak{X} \to \mathbb{R}$ is μ-measurable just in case $f \circ T^{-1}$ is T_*μ-measurable. Hence "P holds μ-almost everywhere" is equivalent to "$P \circ T^{-1}$ holds T_*μ-almost everywhere."

When the main text instead uses **Baire category** (e.g., "for a comeager set of data"), assumption (iii) guarantees that **comeager/meager** status is preserved by T, because T is a homeomorphism on a dense G_δ core and maps meager sets to meager sets. Hence, all ensemble-level claims in §§ 4–7 retain their meaning and truth value across codings, provided we transport the measure (or the topology) in the canonical way.

C.4. Transport of measurable uniformizations and selectors

Let $R \subseteq \mathfrak{X} \times \mathfrak{S}$ be a definable *admissible-pair* relation (data → admissible states), and suppose s: $\mathfrak{X} \to \mathfrak{S}$ is a *measurable selector* with graph(s)⊆R (e.g., the output of an arg-min rule). Given an admissible recoding (T,U), define $R' = (T \times U)[R] \subseteq \mathfrak{Y} \times \mathfrak{S}'$ and $s' = U \circ s \circ T^{-1}: \mathfrak{Y} \to \mathfrak{S}'$. Then s′ is measurable for T_*μ whenever s is measurable for μ, and graph(s′)⊆R′. In particular, if under PD a projective multifunction admits a *measurable uniformization* in one coding, the transported multifunction admits a measurable uniformization in any other coding. This justifies using whichever coding is convenient without jeopardizing the existence or measurability of selectors that the paper requires. (Conversely, if only Δ^1_2, merely definable uniformizations exist—as under V= L at higher projective levels—transport preserves definability but not measurability.)

C.5. Equivariance, symmetry, and the strong notion of identity

The **identity** notion in § 3.3 requires more than measurability. It also requires that the *same* outcome be selected under all physically admissible *recodings*. Formally, let a (countable) group $\mathscr{T}$ of admissible recodings act on the data and state code spaces via (T,U)↦(T,U). Let Γ(x) be the admissible-fiber multifunction and Φ a Borel *secondary cost* used to break ties, both *$\mathscr{T}$-invariant* in the sense that Γ(Tx)=U[Γ(x)] and Φ(Tx,Uy)=Φ(x,y). Consider the *arg-min* fibers $\Gamma^{\bigstar}(x)=\arg\min_{y \in \Gamma(x)} \Phi(x,y)$.

(1) **Structural uniqueness.** If for μ-almost every x the fiber $\Gamma^{\bigstar}(x)$ is a singleton $\{y_x\}$, then the selector $s(x)=y_x$ is Borel whenever graph(Γ) is Borel, and it is **$\mathscr{T}$-invariant** in the strong sense s(Tx)=U(s(x)) for all (T,U)∈$\mathscr{T}$. Identity is automatic.

(2) **Genuine multiplicity.** If $\Gamma^{\bigstar}(x)$ contains multiple elements related by the $\mathscr{T}$-action, then *measurability alone does not enforce identity*. Under PD one still obtains measurable selectors s with $s(x) \in \Gamma^{\bigstar}(x)$ for μ-almost every x, but a $\mathscr{T}$-invariant selector exists only if there is a canonical $\mathscr{T}$-fixed choice in each fiber (for example, because additional physically mandated symmetry constraints or secondary costs pick out a unique representative). In the absence of such structure, identity can fail even though measurable selection holds. Under V=L, where only definable Δ^1_2 tie-breakers may be available, both measurability and $\mathscr{T}$-invariance may fail.

This dichotomy is exactly what § 3.3 exploits. PD restores **ensemble measurability** at the projective level. **Identity** still requires either structural uniqueness or an explicitly **symmetric** tie-break. Nothing in the Appendix claims more than this, and nothing in the text needs more.

C.6. Relative parameters and gauges

Occasionally, passing between codings requires fixed *external parameters* a (calibrations, gauge choices, discretization furniture). All definitions and results above lift *uniformly relative to a*: one simply works with $\Delta^1_2(a)$, $\Sigma^1_{\square}(a)$, $\Pi^1_{\square}(a)$, and with measures/Polish topologies coded relative to a. Shoenfield absoluteness (for $\Sigma^1_2/\Pi^1_2(a)$) and PD-regularity/uniformization (for projective(a) sets) hold in the same way. So, every "relative" use of coding invariance in the paper is secure.

C.7. Application check (how the pieces are used in §§ 4–7)

In **inverse problems** (§ 4), $\mathfrak{X}$ codes data d, $\mathfrak{S}$ codes reconstructions u, and R encodes the arg-min relation for J_λ. Recodings T are regriddings/format changes; U is the corresponding state recoding. Ensemble claims ("for μ-almost all d, …") are transported by push-forward; measurable uniformizations under PD carry over by § C.4; identity relies on structural uniqueness or symmetric tie-breaks as in § C.5.

In **thin-barrier capacity / Markov uniqueness** (§ 5), A↦S_A is Δ^1_2 and the index set U={A : $\mathrm{cap}_{1,2}$(S_A)=0} is analytic. Changing how A and S_A are coded preserves pointclass and the truth of ensemble statements; the uniqueness hinge is therefore code-robust.

For **zero-temperature Ising** (§ 6), $\sigma \in \{-1,+1\}^{\mathbb{Z}^d}$ and the one-step update U(σ) are coded in standard Borel spaces; "shape changes" of finite boxes induce Borel recodings. The "almost sure" claims are invariant by § C.3; measurable uniformizations at the projective level transport by § C.4; identity again hinges on symmetry/uniqueness, not on measurability alone (§ C.5).

In **decoherence / preferred basis** (§ 7.1), $\mathcal{B}$_adm and the arg-min graph are projective; admissible recodings T on environment-states and U on bases are Borel and measure-class preserving. Under PD, measurable selectors exist and transport; whether a *single* selector is $\mathscr{T}$-invariant depends on symmetry/uniqueness (the group $\mathscr{T}$ captures "mere renamings" of bases). Under V=L, definable but non-measurable tie-breakers may exist, and ensemble talk may be ill-posed; both conclusions are code-robust.

For **Kerr interiors** (§§ 7.2–7.3), the data space $\mathfrak{D}$ and horizon-trace space are coded in Polish spaces; recodings are the standard ones from discretization and representation changes. The analytic facts (existence of Γ(d), l.s.c. of $\mathbf{E}_0$, Borel nature of $\mathbf{E}_1$) are unaffected by recoding; ensemble claims transport by § C.3; the canonicalization step behaves as in § C.5.

C.8. Summary

Across the board, definability and regularity (the projective pointclasses that drive the metatheoretic toggles) are preserved across admissible codings; "almost all" and "comeager" claims are invariant once we transport measures and, where relevant, work with Baire isomorphisms; and measurable selectors transport along Borel isomorphisms. These three facts underwrite the main text's ensemble-level claims. What they do not say—by design—is that measurability alone forces identity. Identity, as argued in § 3.3, is secured either by structural uniqueness or by an explicit symmetry requirement. Otherwise it may fail even under PD.

Appendix D. Energy estimates used in §7

Let $T_{ab} = \partial_a \varphi\, \partial_b \varphi - \frac{1}{2} g_{ab}(\partial\varphi)^2$ be the energy-momentum tensor and take the Killing field $K = \partial_v$ in the near-horizon region. The weighted current $J_a = e^{\kappa v}\, T_{ab}\, K^b$ satisfies

$$\nabla \cdot J = \kappa\, e^{\kappa v}\, |\partial_v \varphi|^2 + O(e^{-v}).$$

Integrating on slabs bounded by $v = v_1, v_2$ quantifies the **blue-shift** amplification and controls the **horizon flux** in the first term, with boundary terms subordinate. These standard computations underwrite continuity of the **trace map** and lower-semicontinuity of the flux part of $\mathcal{E}$ (Wald 1984; O'Neill 1995; Dafermos & Luk 2017).

Appendix E. Finite precision, ε-determinism, and why the toggle survives

Finite precision does not wash out the metatheoretic dependence. First, the horizon-flux component of $\mathcal{E}$ is continuous where finite, and $\mathcal{E}$ as a whole is lower-semicontinuous in natural weak topologies. Small data perturbations change costs only slightly. Second, **ε-determinism** for an observable $\mathcal{O}$ on a data domain D amounts to the existence of $\varepsilon_0 > 0$ and a Borel map S with outputs within ε whenever inputs lie within ε_0. Under LC, the measurable projective uniformization in §7 yields measurable selectors (and Borel representatives μ-a.e.) so that ε-determinism makes sense for projectively defined domains D. Under V=L, the same D may be non-measurable, so ensemble-level ε-determinism statements ("for almost every d") are ill-posed and rival Δ^1_2 selectors can disagree on a Δ^1_2, non-measurable set regardless of ε. The analytic layer (continuity/l.s.c.) is again ZFC-robust; the regularity layer that ensemble talk presupposes is where LC vs V=L continues to matter.

Appendix F. Cumulative protocols for testing the determinism profile

This appendix sketches a way in which one might *actually look for* the metatheoretic effects described in the paper. The goal is to ask whether ordinary physical systems behave as if the measurable, coding-invariant selections required by the large-cardinal (LC) picture are realized in nature. Each of the following protocols adapts an example already discussed—imaging, diffusion, spin dynamics, decoherence, and black-hole interiors—to an experimentally or computationally testable form. The idea is cumulative rather than one-off. Repeated refinements and cross-checks probe whether observable quantities stabilize as they should if the underlying selections are measurable. When the LC profile holds, convergence and invariance emerge naturally. When only definable, non-measurable selections are available (as under V=L) the same procedures fail to settle down in reproducible, law-of-large-numbers, fashion.

(1) Imaging and inverse problems (identity → uniqueness)

In medical or geophysical imaging, one reconstructs a signal u from data d by minimizing a total-variation–Tikhonov functional

$$\mathbf{J}_{\square}(\mathbf{u}; \mathbf{d}) = \|\mathbf{A}\mathbf{u} - \mathbf{d}\|^2 + \lambda \cdot \mathrm{TV}(\mathbf{u}),$$

with λ chosen by the usual discrepancy principle: pick the smallest λ such that the reconstruction fits within the noise tolerance $\tau\,\varepsilon$. To test whether the reconstruction rule is measurably well-defined, one can process identical raw data under several admissible encodings (different interpolation schemes, file formats, or sampling grids) and track how the reconstructions converge as resolution increases. If the LC regularity obtains, the reconstructions approach one another. The "reconstruct-then-recode" and "recode-then-reconstruct" procedures eventually agree. Under V=L, the gap between them should persist, revealing a failure of measurable, coding-invariant canonization.

(2) Markov uniqueness and thin barriers (uniqueness)

Porous or granular materials provide a way to test ensemble-level uniqueness. One can fabricate samples in which the insulating barrier *S* is generated by a single calibration rule, e.g., adjusting a threshold θ so that the blocked fraction equals a target φ^*. Measuring diffusion or the Dirichlet-to-Neumann map then tells us whether the associated operator $L = -\nabla\cdot(a\nabla\cdot)$ has a unique Markovian extension. If LC regularity governs the system, the fraction of "unique" versus "non-unique" samples stabilizes and remains the same across equivalent coarse-grainings. If V=L were the right backdrop, the frequencies could drift with the coarse-graining scheme, betraying the loss of measurable uniformization.

(3) Zero-temperature Ising dynamics (identity, discrete)

A more discrete test arises in spin-system arrays. In a two- or three-dimensional ferromagnet, let us define the one-step update U that flips precisely those spins whose flip lowers the limiting energy density along van Hove sequences. Approximating the thermodynamic limit with finite boxes of different shapes (cubes, slabs, nearly square regions) one can compare how the updates behave across shapes. If the underlying selection is measurable (the LC case), the updates coincide almost surely as the volume grows. If not, distinct shapes can yield distinct outcomes for a fixed initial configuration, and that disagreement fails to wash out. This would indicate V=L-style identity failure.

(4) Decoherence and pointer-basis selection (uniqueness / identity)

Quantum platforms such as optomechanical resonators or trapped ions allow a parallel test on the quantum side. Fix the system Hamiltonian and vary the environment's initial state ρ_E across a family of thermal or Gaussian mixtures. For each run, evaluate a dephasing-rate functional $\Phi[B; \rho\_E]$ over a small menu of physically admissible bases B (position-like, momentum-like, etc.), and let the basis minimizing Φ define $B^*(\rho\_E)$. If the map $\rho\_E \mapsto B^*(\rho\_E)$ stabilizes in frequency and remains invariant under equivalent ways of encoding ρ_E, the ensemble behaves measurably, as LC predicts. Persistent dependence on the encoding—different "preferred bases" under equally legitimate parametrizations—would be evidence of the opposite profile.

(5) Elliptic media and weak-form coherence (coherence)

In conductivity or heat-flow experiments, one can vary material coefficients $\nu\_A(x) = \nu_0 + \alpha\cdot\chi\_\{W\_A\}(x)$, where each W_A is generated by a uniform calibration or thresholding rule. The question is whether the weak formulation

$$a(u,v) = \int_\Omega \nu_A(x) \langle \nabla u, \nabla v \rangle \, dx$$

remains stable (bounded and coercive) under different admissible codings and refinements of A. Empirically, all W_A will of course be Borel, but the experiment tests whether a single physical rule yields coding-invariant results (the LC profile) or whether outcomes depend on representational choices (a V=L–type incoherence).

(6) Kerr-interior simulations (identity / ensemble)

Numerical relativity offers an analogue of the same idea. Evolve the wave equation $\Box_g \varphi = 0$ inside a sub-extremal Kerr black hole from an ensemble of initial data $d \in \mathfrak{D}$. Compute the weighted horizon flux $\mathcal{F}(d)$ and apply the lexicographic selection described in the main text, first minimizing energy mildness E_0, then late-time tameness E_1. Run the full simulation twice under distinct but equivalent discretization schemes. If the LC regime applies, the two pipelines converge toward the same continuation as resolution increases. If V=L, the divergence between them remains bounded away from zero. This is a reproducible split in what counts as the "canonical" continuation.

Interpreting the outcomes

If the stability indices and cross-coding discrepancies (quantities such as SSI(n), SC(n), and their Kerr analogue) decay toward zero, if ensemble frequencies settle, and if law-of-large-numbers behavior appears across representations, the evidence supports the LC-style regularity the analysis presupposes. On the other hand, a plateau in these indices, persistent dependence on arbitrary codings, or lack of convergence would suggest that only definable, non-measurable selections are operative, the V=L profile. Note that none of these experiments involves exotica. They just ask whether nature enforces measurable, coding-invariant, and refinement-stable behavior. If it does, convergence and invariance will appear on their own. If it does not, the attempt to preserve those scientific virtues will fail in systematic, reproducible ways.

Bibliography

Arrigoni, Tatiana. 2011. “V = L and Intuitive Plausibility in Set Theory. A Case Study.” Bulletin of Symbolic Logic 17 (3): 337–360.
Ashtekar, Abhay, and Jerzy Lewandowski. 2004. “Background Independent Quantum Gravity: A Status Report.” Classical and Quantum Gravity 21: R53–R152.
Balaguer, Mark. 1998. Platonism and Anti-Platonism in Mathematics. Oxford University Press.
Berry, Sharon. 2023. “Mathematical Structuralism and Indeterminacy.” Philosophical Quarterly 73 (2): 412–431.
—--. Forthcoming. “Set-Theoretic Pluralism and Physical Applications.” Philosophy Compass.
Brezis, Haim. 2011. Functional Analysis, Sobolev Spaces and Partial Differential Equations. Springer.
Butterfield, Jeremy. 2007. “Reconsidering Relativistic Causality.” International Studies in Philosophy of Science 21 (3): 295–328.
Castaing, Charles, and Michel Valadier. 1977. Convex Analysis and Measurable Multifunctions. Lecture Notes in Mathematics 580. Springer.
Carnap, Rudolf. 1950. The Logical Foundations of Probability. University of Chicago Press.
Chen, Eddy Keming. 2022. “Fundamental Nomic Vagueness.” Philosophical Review 133 (1): 1–42.
Clarke-Doane, Justin. 2022. Mathematics and Metaphilosophy. Cambridge University Press.
Clarke-Doane, Justin. 2024. “Realism, Objectivity, and Evaluation.” In Oxford Studies in Metaethics, Vol. 19, 233–258. Oxford University Press.
Clarke-Doane, Justin. 2025. “Set Theory and Physical Reality.” Erkenntnis 90 (1): 123–145.
Cohen, Paul J. 1963. “The Independence of the Continuum Hypothesis.” PNAS 50 (6): 1143–1148.
Cohen, Paul J. 1966. Set Theory and the Continuum Hypothesis. W. A. Benjamin.
Colyvan, Mark. 1998. “In Defence of Indispensability.” Philosophia Mathematica 6 (1): 39–62.
Colyvan, Mark. 2001. The Indispensability of Mathematics. Oxford University Press.
Dafermos, Mihalis. 2014. “The Mathematical Analysis of Black Holes in General Relativity.” ICM Proceedings 3: 747–772.
Dafermos, Mihalis, and Jonathan Luk. 2017. “The Interior of Dynamical Vacuum Black Holes I: The C^0-Stability of the Kerr Cauchy Horizon.” arXiv:1710.01722.
Devlin, Keith. 1977. “Constructibility.” In Jon Barwise (ed.), Handbook of Mathematical Logic, Studies in Logic and the Foundations of Mathematics, Vol. 90, 453–490. North-Holland.
—--.. 1984. Constructibility. Springer.
Earman, John. 1986. A Primer on Determinism. D. Reidel.
Earman, John. 2007. “Aspects of Determinism in Modern Physics.” In Philosophy of Physics, 1369–1434. Elsevier.
Earman, John. 2009. “Essential Self-Adjointness: Implications for Determinism and the Classical-Quantum Correspondence.” Synthese 169 (1): 27–50.
Engl, Heinz W., Martin Hanke, and Andreas Neubauer. 1996. Regularization of Inverse Problems. Kluwer.
Evans, Lawrence C. 2010. Partial Differential Equations, 2nd ed. American Mathematical Society.
Farah, Ilijas, and Menachem Magidor. 2012. “Gödel’s Program and Ultimate L.” In Infinity and Truth, 33–55. World Scientific.
Field, Hartry. 1989. Realism, Mathematics, and Modality. Blackwell.

Fletcher, Samuel C. 2012. “What Counts as a Newtonian System? The View from Norton’s Dome.” EJPS 2 (3): 275–297.
Fontanella, Laura. 2019. “How to Choose New Axioms for Set Theory.” In S. Centrone, D. Kant, and D. Sarikaya (eds.), Reflections on the Foundations of Mathematics, 27–42. Springer.
Fraenkel, Abraham A., Yehoshua Bar-Hillel, and Azriel Levy. 1973. Foundations of Set Theory, 2nd ed. North-Holland.
Friedman, Harvey. 1998. “Finite Functions and the Necessary Use of Large Cardinals.” Annals of Mathematics 148: 803–893.
Fukushima, Masatoshi, Yoichi Oshima, and Masayoshi Takeda. 2011. Dirichlet Forms and Symmetric Markov Processes, 2nd ed. de Gruyter.
Glauber, Roy J. 1963. “Time-Dependent Statistics of the Ising Model.” Journal of Mathematical Physics 4 (2): 294–307.
Gödel, Kurt. 1931. “Über formal unentscheidbare Sätze … I.” Monatshefte für Mathematik 38: 173–198.
Gödel, Kurt. 1940. The Consistency of the Continuum Hypothesis. Princeton University Press.
Hadamard, Jacques. 1923. Lectures on Cauchy’s Problem in Linear Partial Differential Equations. Yale University Press.
Halvorson, Hans, J. B. Manchak, and James Owen Weatherall. 2025. “Defining Determinism”. Preprint.
Hamkins, Joel David. 2012. “The Set-Theoretic Multiverse.” Review of Symbolic Logic 5 (3): 416–449.
Hawking, Stephen W., and George F. R. Ellis. 1973. The Large Scale Structure of Space-Time. Cambridge University Press.
Hinman, Peter G. 2018. Fundamentals of Mathematical Logic. A K Peters.
Ismael, Jenann. 2016. How Physics Makes Us Free. Oxford University Press.
Jech, Thomas. 2003. Set Theory, 3rd Millennium ed. Springer.
Jensen, R. Björn. 1995. “Inner Models and Large Cardinals.” Bulletin of Symbolic Logic 1 (4): 393–407.
Jonas, Silvia. 2024. “Mathematical Pluralism and Indispensability.” Erkenntnis 89: 2671–2698.
Kanamori, Akihiro. 2009. The Higher Infinite, 2nd ed. Springer.
Kechris, Alexander S. 1995. Classical Descriptive Set Theory. Springer.
Koellner, Peter. 2009. “On Reflection Principles.” Annals of Pure and Applied Logic 157: 206–219.
Kuratowski, Kazimierz, and Czesław Ryll-Nardzewski. 1965. “A General Theorem on Selectors.” Bulletin de l’Académie Polonaise des Sciences 13: 397–403.
Liggett, Thomas M. 1985. Interacting Particle Systems. Springer.
Loewer, Barry. 2001. “Determinism and Chance.” SHPMP 32 (4): 609–620.
—–-. 2020. “The Mentaculus Vision.” In Statistical Mechanics and Scientific Explanation, 3–29. World Scientific.
Luk, Jonathan, and Sung-Jin Oh. 2017. “Strong Cosmic Censorship in Spherical Symmetry I.” arXiv:1702.05715.
—–-. 2018. “Strong Cosmic Censorship … II.” Annals of PDE 5: Article 6.
—–-., Jonathan, and Sung-Jin Oh. 2019. “Linear Instability of the Reissner–Nordström Cauchy Horizon.” Duke Mathematical Journal 168 (16): 2945–3065.
—–-., Jonathan, and Sung-Jin Oh. 2022. “Strong Cosmic Censorship in Charged Black Hole Spacetimes.” Annales Henri Poincaré 23: 3229–3275.

Ma, Zhi-Ming, and Michael Röckner. 1992. Introduction to the Theory of (Non-Symmetric) Dirichlet Forms. Springer.
Maddy, Penelope. 1997. Naturalism in Mathematics. Clarendon Press.
—-., Penelope. 2011. Defending the Axioms. Oxford University Press.
Malament, David. 2008. "Norton's Slippery Slope." Philosophy of Science 75 (5): 799–816.
Martin, Donald A., and John R. Steel. 1989. "A Proof of Projective Determinacy." Journal of the AMS 2 (1): 71–125.
Moschovakis, Yiannis. 2009. Descriptive Set Theory, 2nd ed. American Mathematical Society.
Norton, John D. 2008. "The Dome: An Unexpectedly Simple Failure of Determinism." Philosophy of Science 75 (5): 786–798.
Ollivier, Harold, David Poulin, and Wojciech H. Zurek. 2004. "Objective Properties from Subjective Quantum States …" PRL 93: 220401.
O'Neill, Barrett. 1995. The Geometry of Kerr Black Holes. A K Peters.
Penrose, Roger. 1979."Singularities and Time-Asymmetry." In S. W. Hawking and W. Israel (eds.), General Relativity: An Einstein Centenary Survey, 581–638. Cambridge University Press.
Pinter, Charles C. 2014. A Book of Set Theory. Dover.
Polchinski, Joseph. 1998. String Theory, Vols. 1–2. Cambridge University Press.
Pour-El, Marian B., and J. Ian Richards. 1989. Computability in Analysis and Physics. Springer.
Putnam, Hilary. 1967. "Mathematics without Foundations." Journal of Philosophy 64 (1): 5–22.
Quine, W. V. O. 1951. "Two Dogmas of Empiricism." Philosophical Review 60 (1): 20–43.
—-. 1986. Philosophy of Logic, 2nd ed. Harvard University Press.
—-. 1990. Pursuit of Truth. Harvard University Press.
Resnik, Michael D. 1997. Mathematics as a Science of Patterns. Clarendon Press.
Riedel, C. Jess, Wojciech H. Zurek, and Michael Zwolak. 2014. "Objective Classical Reality from Quantum Origins." PRL 112: 040401.
Rogers, Hartley Jr. 1987. Theory of Recursive Functions and Effective Computability, 2nd ed. MIT Press.
Rudin, Leonid I., Stanley Osher, and Emad Fatemi. 1992. "Nonlinear Total Variation Based Noise Removal Algorithms." Physica D 60: 259–268.
Ruelle, David. 1969. Statistical Mechanics: Rigorous Results. W. A. Benjamin.
Scherzer, Otmar, Markus Grasmair, Harald Grossauer, Markus Haltmeier, and Frank Lenzen. 2009. Variational Methods in Imaging. Springer.
Schlosshauer, Maximilian. 2007. Decoherence and the Quantum-to-Classical Transition. Springer.
Shoenfield, Joseph R. 1961. "The Problem of Predicativity." In Bar-Hillel (ed.), Essays on the Foundations of Mathematics, 132–139. Magnes Press.
Sierpiński, Wacław. 1924. "Sur l'hypothèse du continu ($2^{\aleph_0} = \aleph_1$)." Fundamenta Mathematicae 5 (1): 177–187.
Simpson, Stephen. 1999. Subsystems of Second-Order Arithmetic. Cambridge University Press.
Solovay, Robert M. 1970. "A Model … in Which Every Set of Reals Is Lebesgue Measurable." Annals of Mathematics 92 (1): 1–56.
Steel, John R. 2014. "Gödel's Program." In Kennedy (ed.), Interpreting Gödel, 153–179. Cambridge University Press.
Wald, Robert M. 1984. General Relativity. University of Chicago Press.
Weihrauch, Klaus. 2000. Computable Analysis. Springer.
Wilson, Mark. 2009. "Determinism and the Mystery of the Missing Physics." BJPS 60 (1):

173–193.
Woodin, W. Hugh. 1988. “Supercompact Cardinals, Sets of Reals, and Weakly Homogeneous Trees.” PNAS 85 (18): 6587–6591.
Zurek, Wojciech H. 1981. “Pointer Basis of Quantum Apparatus: Into What Mixture Does the Wave Packet Collapse?” Phys. Rev. D 24: 1516–1525.
Zurek, Wojciech H. 2003. “Decoherence, Einselection, and the Quantum Origins of the Classical.” Rev. Mod. Phys. 75: 715–775.