

# Monotone FIM Collapse Along the Tao Cascade: Closing the Equidistribution Gap Without Cellular Automaton Mixing

D. Christian<sup>1,\*</sup>

<sup>1</sup>*The Axioms of Pattern Ontology (APO) Initiative*

(Dated: March 12, 2026)

The independence results for averaged Navier–Stokes regularity in [1] identify an “equidistribution gap”: the information-geometric (FIM) route to undecidability requires that a specific deterministic orbit equidistributes, but only measure-generic equidistribution is established. We close this gap for the averaged system by analyzing the FIM directly along Tao’s frequency cascade, bypassing cellular automaton mixing theory entirely.

Tao’s blowup mechanism is a self-replicating machine that transfers energy from frequency scale  $n$  to scale  $n + 1$ , erasing the previous scale. We prove that this erasure forces monotone FIM contraction:  $\lambda_1(t_{n+1}) \leq (1 - \eta)\lambda_1(t_n) + \mathcal{O}(\epsilon^n)$  for an explicit  $\eta > 0$ , giving  $\lambda_1(t_n) \rightarrow 0$  geometrically. The key mechanism is the strict convexity of Fisher information under mixture: during each cascade step, energy passes through a two-scale mixture, which necessarily decreases the spectral gap.

This completes the FIM route for averaged NS, giving a second proof of undecidability that is fully information-geometric. For exact (physical) NS, the situation is fundamentally different: Tao’s engineered nonlinearity is unavailable, and the cascade analysis reveals precisely what the computational encoding requires—logic gates built from a chosen nonlinearity, velocities driven to infinity before dissipation can act. The exact nonlinearity  $(u \cdot \nabla)u$  does not have this structure.

The C2 equivalence and Shoenfield absoluteness of [1] constrain the exact question to a four-way dichotomy (R/F/I/I’), with the truth value forcing-invariant. We examine each scenario against the Church–Turing theorem and the C2 equivalence, and find that the answer to the Clay problem determines whether a fundamental PDE is computationally universal—a structural characterization that explains why the problem has resisted resolution and what each possible outcome would mean for the relationship between physics, computation, and formal proof.

## I. INTRODUCTION

### A. The Problem

The companion paper [1] establishes the undecidability of averaged Navier–Stokes regularity via two independent routes. The *direct route* uses Tao’s universality theorem [3]—averaged NS regularity from datum  $u_0^M$  is equivalent to halting of machine  $M$ —combined with the Church–Turing theorem. This requires no additional machinery.

The *geometric (FIM) route* introduces the Fisher Information Matrix spectral gap  $\lambda_1(t)$  and proves:

(B) **Backward:**  $\lambda_1(t) \rightarrow 0$  implies blow-up (Theorem 3.2 of [1]).

(F) **Forward:** Blow-up should imply  $\lambda_1(t_n) \rightarrow 0$ , providing the geometric mechanism for FIM collapse.

The backward direction (B) is proven unconditionally for all strong NS solutions. The forward direction (F) for averaged NS was established in [1] only via a chain: non-halting machine  $\rightarrow$  CA runs forever  $\rightarrow$  CA equidistributes  $\rightarrow P_n \rightarrow P_{\text{unif}} \rightarrow \lambda_1 \rightarrow 0$ . This chain encounters the *equidistribution gap*: equidistribution of a specific deterministic CA orbit is not guaranteed by the standard measure-theoretic results.

The weaker conclusion  $\inf_n \lambda_1(t_n) = 0$  follows from the backward direction by contrapositive (Lemma 4.3(NH) of [1]), but the full monotone convergence  $\lambda_1(t_n) \rightarrow 0$  remained open.

### B. Why the CA Framing Is Wrong

The equidistribution gap persists because the CA framing misidentifies the structure of Tao’s construction. As analyzed in [2], the conditions required for CA-based equidistribution—spatial  $\alpha$ -mixing and  $\beta$ -normality—apply to

---

\* daniel@freereason.org

classical shift-space cellular automata on  $\{0, 1\}^{\mathbb{Z}}$ . But Tao’s construction is not a spatially extended CA in this sense. It is a *frequency cascade*: a system of ODEs coupling adjacent frequency shells, where the “computation” happens in the interaction between scales, not in the spatial arrangement of bits.

Specifically, Tao constructs [3]:

- Four scalar fields  $(X_{1,n}, X_{2,n}, X_{3,n}, X_{4,n})$  at each frequency scale  $n$ , acting as logic gates (pump, amplifier, rotor, timing).
- Quadratic interactions coupling scale  $n$  to scale  $n + 1$ , energy-conserving by design.
- A self-replication mechanism: at time  $t_n$ , the system at scale  $n$  “programs” a copy of itself at scale  $n + 1$  (which is  $(1 + \epsilon_0)$  times finer), then erases the original.

The initial datum  $u_0$  is constructed via the ansatz

$$u(t, x) \approx \sum_n X_n(t) \lambda^{3n/5} \psi(\lambda^{2n/5} x) \quad (1)$$

for a fixed Schwartz function  $\psi$  with Fourier transform supported away from the origin, and  $\lambda = (1 + \epsilon_0)^{5/2}$ . The “tape” encoding the Turing machine lives in the ODE coefficients  $X_{j,n}(0)$ , not in a spatial binary sequence.

There is no natural shift structure, no Bernoulli measure, and no spatial notion of word frequency. The classical CA equidistribution theory (Hedlund [4], Peyrière [5]) does not apply.

### C. The New Approach

We bypass the CA framing entirely and analyze  $\lambda_1(t)$  directly along the cascade. The key observation is that the replication-and-erasure mechanism of the Tao cascade has a direct consequence for the FIM: each replication step transfers the velocity field’s energy to a finer scale while destroying the spatial structure at the previous scale. This destruction of spatial structure is precisely the loss of parametric distinguishability that drives  $\lambda_1$  to zero.

The argument proceeds in three steps:

1. **Spectral localization** (Section III): At time  $t_n$ , the velocity field  $u(t_n, \cdot)$  is concentrated in a frequency shell of radius  $\sim \lambda^{2n/5}$ , with energy fraction  $\geq 1 - C\epsilon^n$  at scale  $n$ .
2. **FIM and spectral support** (Section IV): The FIM spectral gap  $\lambda_1$  is controlled by the spatial gradients of the velocity field relative to its amplitude. When the velocity field is localized at frequency scale  $k$ , the FIM is determined by the structure within that shell.
3. **Erasure implies FIM contraction** (Section V): The replication step transfers energy from shell  $n$  to shell  $n + 1$  while erasing shell  $n$ . The score function—the logarithmic derivative of the velocity distribution—undergoes a rescaling that contracts  $\lambda_1$  by a factor  $(1 - \eta)$  per step.

Together, these give the main result:

**Theorem I.1** (Monotone FIM collapse along the cascade). *Under the Tao encoding of a non-halting Turing machine  $M$ , the FIM spectral gap of the averaged NS flow satisfies*

$$\lambda_1(t_{n+1}) \leq (1 - \eta) \lambda_1(t_n) + R_n, \quad (2)$$

where  $\eta = \eta(\epsilon_0) > 0$  depends only on the cascade parameter  $\epsilon_0$ , and  $R_n = \mathcal{O}(\epsilon^n)$  is an error from the residual energy in off-scale modes. In particular,  $\lambda_1(t_n) \rightarrow 0$  geometrically.

**Corollary I.2** (FIM route succeeds). *The forward spectral equivalence holds for averaged NS from Tao-encoded data: if  $M$  does not halt, then  $\lambda_1(t_n) \rightarrow 0$ . Combined with the backward spectral equivalence (Theorem 3.2 of [1]), this gives a complete information-geometric proof of the halting–regularity equivalence, and hence of undecidability, without invoking cellular automaton equidistribution.*

## II. SETUP AND NOTATION

### A. What Tao Built

Tao’s construction [3] is often summarized as “an averaged Navier–Stokes equation that blows up.” This summary obscures the mechanism. We now describe the construction at the level of detail needed to understand both why the FIM cascade works and why it cannot work for exact NS.

### 1. The averaged nonlinearity

The exact Navier–Stokes equation has the form  $\partial_t u = \Delta u + B(u, u)$ , where  $B$  is the Euler bilinear operator: the Leray projection of the advection term  $(u \cdot \nabla)u$ . In Fourier space, the trilinear form  $\langle B(u, v), w \rangle$  is determined by a fixed tensor  $\Lambda_{\xi_1, \xi_2, \xi_3}$  evaluated on the Fourier modes at frequencies  $\xi_1 + \xi_2 + \xi_3 = 0$ .

Tao replaces  $B$  with an *averaged* operator  $\tilde{B}$ , defined by

$$\langle \tilde{B}(u, v), w \rangle = \mathbb{E}[\langle B(m_1 R_1 D_1 u, m_2 R_2 D_2 v), m_3 R_3 D_3 w \rangle], \quad (3)$$

where the expectation is over random rotations  $R_i \in \text{SO}(3)$ , random dilations  $D_i$  (bounded away from 0 and  $\infty$ ), and random real Fourier multipliers  $m_i$  of order zero. The cancellation  $\langle \tilde{B}(u, u), u \rangle = 0$  is preserved by construction, so the averaged equation obeys the same energy identity as the true NS equation.

The averaging is the critical engineering step. Every Sobolev-space estimate that holds for  $B$  also holds for  $\tilde{B}$  (by Minkowski’s inequality and the Hörmander–Mikhlin theorem), so  $\tilde{B}$  is “no stronger” than  $B$  from a harmonic analysis perspective. But the averaging introduces *freedom*: by choosing the probability distribution on  $(R_i, D_i, m_i)$  carefully, Tao selects which frequency interactions are active and which are suppressed. This freedom is what allows the construction of logic gates.

### 2. The dyadic model and the four-component gates

The blow-up solution is constructed via a dyadic model—an ODE system for mode amplitudes  $(X_{j,n})_{j=1}^4$  at each frequency scale  $n$ , with  $\lambda = (1 + \epsilon_0)^{5/2}$ . The four components at each scale serve as logic gates:

- A *pump* that stores energy;
- An *amplifier* that transfers energy to the next scale;
- A *rotor* that controls the timing of the transfer;
- A *timing gate* that prevents premature energy leakage.

The interactions between these components are quadratic and energy-conserving. The gate design ensures that at each scale  $n$ , the system runs a “quadratic computer” that, after time  $t_{n+1} - t_n$ , abruptly replicates itself at scale  $n + 1$  (which is  $(1 + \epsilon_0)$  times finer and  $(1 + \epsilon_0)^{5/2}$  times faster) while erasing the copy at scale  $n$ .

This is the von Neumann machine: a self-replicating construct built within the laws of the averaged equation.

### 3. Why it outruns dissipation

The viscous dissipation at scale  $n$  acts at rate  $\lambda^{4n/5} \sim (1 + \epsilon_0)^{2n}$ —the natural decay rate for frequency  $\sim \lambda^{2n/5}$ . The cascade time between replication steps is  $t_{n+1} - t_n \sim \lambda^{-(5/2 - \mathcal{O}(\epsilon_0))n}$ . Since  $5/2 > 4/5$  (supercriticality), the cascade gets faster *relative to dissipation* at each step. The replication outruns the viscous decay, and the energy concentrates at ever-finer scales, producing blow-up at the accumulation time  $T^* = \lim t_n < \infty$ .

This is the supercriticality barrier: in three dimensions, the energy identity controls a supercritical norm, so the dissipative term cannot prevent the energy cascade from running to infinity. Tao’s theorem proves that any argument relying only on the energy identity and Sobolev estimates for the nonlinearity is defeated by this barrier, because  $\tilde{B}$  satisfies all the same estimates as  $B$ .

### 4. Why it cannot work for exact NS

The exact nonlinearity  $B$  is fixed by the physics: it is the Leray projection of  $(u \cdot \nabla)u$ , with a specific tensor  $\Lambda_{\xi_1, \xi_2, \xi_3}$  determined by the geometry of  $\mathbb{R}^3$ . There is no freedom to choose which frequency interactions are active. All trilinear interactions permitted by the divergence-free condition are present, with amplitudes determined by the fixed tensor.

This means:

- No gate structure can be imposed. The four-component design requires specific interactions to be present and others to be absent. The exact  $B$  has all interactions on simultaneously.

- No timing control is possible. The rotor and timing gate rely on suppressing unwanted energy transfers during the replication step. With all interactions active, energy leaks to neighboring scales before the replication completes.
- The energy barrier is not overcome. Without the gate structure, the energy identity  $\frac{d}{dt} \|u\|_{L^2}^2 = -2\nu \|\nabla u\|_{L^2}^2 \leq 0$  caps the total energy, and any finite-energy datum exhausts its computational capacity in finitely many steps (Proposition 6.5 of [1]).

Tao himself identifies this as the key open problem: constructing logic gates from ideal (inviscid, non-averaged) fluid. No one has done this for the Euler or Navier–Stokes equations on  $\mathbb{R}^3$ . The averaged nonlinearity  $\tilde{B}$  is a mathematical object engineered to permit gate construction. The exact nonlinearity  $B$  is a physical object that does not.

## B. The Cascade in Summary

The averaged Navier–Stokes equation (??) with the engineered  $\tilde{B}$  admits solutions of the form (1). The ODE system for the mode amplitudes has the schematic form

$$\dot{X}_{j,n} = -\lambda^{4n/5} X_{j,n} + Q_{j,n}(X_{\cdot,n-1}, X_{\cdot,n}), \quad (4)$$

where the first term is viscous dissipation and  $Q_{j,n}$  implements the gate logic.

**Definition II.1** (Cascade times and energy concentration). *The cascade times  $t_0 < t_1 < t_2 < \dots$  are the sequence converging to  $T^* < \infty$  such that:*

1. At time  $t_n$ , the energy is concentrated at scale  $n$ :

$$E_n(t_n) := \sum_j X_{j,n}(t_n)^2 \geq (1 - C_0 \epsilon_0^n) \sum_{k,j} X_{j,k}(t_n)^2 =: (1 - C_0 \epsilon_0^n) E(t_n). \quad (5)$$

2. The time intervals satisfy  $t_{n+1} - t_n \sim \lambda^{-(5/2 - \mathcal{O}(\epsilon_0))n}$ , which is shorter than the dissipation time  $\lambda^{-4n/5}$  at scale  $n$  for large  $n$ .
3. Between  $t_n$  and  $t_{n+1}$ , the replication step transfers energy from scale  $n$  to scale  $n+1$ :

$$E_{n+1}(t_{n+1}) \geq (1 - C_1 \epsilon_0) E_n(t_n), \quad E_n(t_{n+1}) \leq C_1 \epsilon_0 E_n(t_n). \quad (6)$$

Property (6) is the *erasure property*: after replication, most of the energy at scale  $n$  has been transferred to scale  $n+1$ , and the mode amplitudes at scale  $n$  are small. This is the core mechanism of Tao’s construction—the von Neumann machine creates a copy of itself at a finer scale and erases itself.

## C. The Velocity Distribution and FIM

We use the velocity distribution  $P_t(x) = |u(x,t)|^2 / \|u(\cdot,t)\|_{L^2}^2$  and the FIM as defined in [1]. The parameterization  $\theta \mapsto u_0^\theta$  is a smooth family of initial data near  $u_0^M$ , and

$$g_{ij}(\theta, t) = \int_{\mathbb{R}^3} (\partial_i \ln P_t)(\partial_j \ln P_t) P_t dx, \quad \lambda_1(t) = \min_{|v|=1} g_{ij} v^i v^j. \quad (7)$$

The score function is  $s_i = \partial_{\theta_i} \ln P_t = 2(\partial_{\theta_i} u \cdot u) / |u|^2 - c_i$ , where  $c_i = \mathbb{E}_{P_t}[2(\partial_{\theta_i} u \cdot u) / |u|^2]$ .

## D. Frequency-Shell Decomposition

For each scale  $n$ , define the projection onto the  $n$ th frequency shell:

$$u_n(x, t) = \sum_j X_{j,n}(t) \lambda^{3n/5} \psi_j(\lambda^{2n/5} x), \quad (8)$$

where the  $\psi_j$  are the basis functions at scale  $n$ . The total velocity field decomposes as  $u = \sum_n u_n$  with the shells approximately orthogonal in  $L^2$  (the frequency supports are disjoint up to small overlaps controlled by the Schwartz tails of  $\psi$ ).

The velocity distribution then decomposes as

$$P_t(x) = \frac{|u|^2}{\|u\|_{L^2}^2} = \frac{|\sum_n u_n|^2}{\sum_n \|u_n\|_{L^2}^2 + \text{cross terms}}. \quad (9)$$

### III. SPECTRAL LOCALIZATION

The first ingredient is that at cascade time  $t_n$ , the velocity field is effectively supported at a single frequency scale.

**Lemma III.1** (Spectral localization at cascade times). *At time  $t_n$ , the velocity distribution satisfies*

$$P_{t_n}(x) = P_n^{(\text{loc})}(x) + \mathcal{O}(\epsilon_0^{n/2}), \quad (10)$$

where  $P_n^{(\text{loc})}(x) = |u_n(x, t_n)|^2 / \|u_n(\cdot, t_n)\|_{L^2}^2$  is the velocity distribution restricted to shell  $n$ , and the error is in the  $L^1$  norm.

*Proof.* By the energy concentration (5):

$$\|u_n(\cdot, t_n)\|_{L^2}^2 \geq (1 - C_0 \epsilon_0^n) \|u(\cdot, t_n)\|_{L^2}^2. \quad (11)$$

The cross terms in (9) satisfy  $|\langle u_n, u_k \rangle| \leq C \epsilon_0^{|n-k|/2} \|u_n\|_{L^2} \|u_k\|_{L^2}$  for  $k \neq n$ , using the frequency separation of the shells (the Schwartz decay of  $\hat{\psi}$  gives exponential suppression of overlap between shells separated by  $|n - k|$  steps).

The off-shell contribution to  $P_{t_n}$  is bounded by:

$$\left\| P_{t_n} - P_n^{(\text{loc})} \right\|_{L^1} \leq 2 \frac{\|u - u_n\|_{L^2}}{\|u\|_{L^2}} \leq 2\sqrt{C_0 \epsilon_0^n} = \mathcal{O}(\epsilon_0^{n/2}). \quad (12)$$

□

### IV. FIM AND SPECTRAL SUPPORT

We now connect the FIM to the spatial structure of the dominant frequency shell.

**Lemma IV.1** (FIM localization). *At time  $t_n$ , the FIM spectral gap satisfies*

$$\lambda_1(t_n) = \lambda_1^{(n)}(t_n) + \mathcal{O}(\epsilon_0^{n/2}), \quad (13)$$

where  $\lambda_1^{(n)}$  is the FIM spectral gap computed from  $P_n^{(\text{loc})}$  alone.

*Proof.* The score function  $s_i = \partial_{\theta_i} \ln P_t$  is a continuous functional of  $P_t$  (in the  $L^2(P_t)$  topology). By Lemma III.1,  $P_{t_n}$  is  $\mathcal{O}(\epsilon_0^{n/2})$ -close to  $P_n^{(\text{loc})}$  in  $L^1$ . The FIM is the  $L^2(P_t)$ -norm of the score, and the Lipschitz continuity of  $g_{ij}$  with respect to  $L^1$  perturbations of the distribution (a standard result in information geometry, see e.g. [8]) gives

$$\left| \lambda_1(t_n) - \lambda_1^{(n)}(t_n) \right| \leq C \left\| P_{t_n} - P_n^{(\text{loc})} \right\|_{L^1} \cdot \max(\lambda_1(t_n), \lambda_1^{(n)}(t_n)) = \mathcal{O}(\epsilon_0^{n/2}). \quad (14)$$

□

### V. ERASURE IMPLIES FIM CONTRACTION

This is the core section. We show that the replication step—which transfers energy from shell  $n$  to shell  $n + 1$  while erasing shell  $n$ —forces a definite decrease in  $\lambda_1$ .

### A. The Rescaling Structure

The key structural feature of the Tao cascade is that the replication step produces a *rescaled copy* of the velocity field. At time  $t_{n+1}$ , the velocity field at scale  $n+1$  is (up to the ODE evolution within the step):

$$u_{n+1}(x, t_{n+1}) \approx \lambda^{3/10} u_n(\lambda^{2/5}x, t_n), \quad (15)$$

where  $\lambda = (1 + \epsilon_0)^{5/2}$ . This is the spatial rescaling that concentrates the energy into a ball  $(1 + \epsilon_0)$  times smaller, at amplitudes  $(1 + \epsilon_0)^{3/4}$  times larger (the exponents follow from the  $L^2$ -normalization at the new scale).

The velocity distribution therefore transforms as:

$$P_{n+1}^{(\text{loc})}(x) = \frac{|u_{n+1}(x, t_{n+1})|^2}{\|u_{n+1}\|_{L^2}^2} \approx \frac{\lambda^{3/5} |u_n(\lambda^{2/5}x, t_n)|^2}{\lambda^{3/5-6/5} \|u_n(\cdot, t_n)\|_{L^2}^2} = \lambda^{6/5} P_n^{(\text{loc})}(\lambda^{2/5}x). \quad (16)$$

This is the standard rescaling of a probability density:  $P_{n+1}^{(\text{loc})}$  is  $P_n^{(\text{loc})}$  concentrated onto a smaller region by the dilation factor  $\lambda^{-2/5}$ .

### B. Score Function Under Rescaling

The score function transforms under the rescaling  $x \mapsto \lambda^{2/5}x$  as follows. For the localized distribution, the score with respect to parameter  $\theta_i$  is:

$$\begin{aligned} s_i^{(n+1)}(x) &= \partial_{\theta_i} \ln P_{n+1}^{(\text{loc})}(x) \\ &= \partial_{\theta_i} \ln P_n^{(\text{loc})}(\lambda^{2/5}x) + \partial_{\theta_i} \ln \lambda^{6/5} \\ &= s_i^{(n)}(\lambda^{2/5}x) + \partial_{\theta_i} \ln \lambda^{6/5}. \end{aligned} \quad (17)$$

The second term vanishes if the cascade parameter  $\lambda$  does not depend on  $\theta$  (i.e., the parameterization  $\theta \mapsto u_0^\theta$  does not change the cascade geometry, only the mode amplitudes). For a generic parameterization, this term is a constant (independent of  $x$ ) and is absorbed into the centering correction  $c_i$  of the score.

Therefore, the centered score satisfies:

$$s_i^{(n+1)}(x) - c_i^{(n+1)} = s_i^{(n)}(\lambda^{2/5}x) - c_i^{(n)}. \quad (18)$$

### C. FIM Contraction

**Proposition V.1** (FIM contraction per cascade step). *Under the rescaling (16), the FIM at step  $n+1$  relates to the FIM at step  $n$  by:*

$$g_{ij}^{(n+1)} = g_{ij}^{(n)} \cdot \frac{\mathbb{E}_{P_{n+1}}[(s_i^{(n+1)} - c_i^{(n+1)})(s_j^{(n+1)} - c_j^{(n+1)})]}{\mathbb{E}_{P_n}[(s_i^{(n)} - c_i^{(n)})(s_j^{(n)} - c_j^{(n)})]}. \quad (19)$$

More explicitly, using the change of variables  $y = \lambda^{2/5}x$ :

$$\begin{aligned} g_{ij}^{(n+1)} &= \int (s_i^{(n+1)}(x) - c_i^{(n+1)})(s_j^{(n+1)}(x) - c_j^{(n+1)}) P_{n+1}^{(\text{loc})}(x) dx \\ &= \int (s_i^{(n)}(y) - c_i^{(n)})(s_j^{(n)}(y) - c_j^{(n)}) \lambda^{6/5} P_n^{(\text{loc})}(y) \lambda^{-6/5} dy \\ &= g_{ij}^{(n)}. \end{aligned} \quad (20)$$

**Remark V.2** (Pure rescaling preserves the FIM). *If the replication step were an exact spatial rescaling with no additional distortion, the FIM would be exactly preserved—the rescaling is a diffeomorphism that maps the distribution to a smaller copy of itself, and the FIM is diffeomorphism-invariant. This is the key insight: the FIM does not collapse from the rescaling alone. The collapse comes from the imperfect nature of the replication and the erasure of the previous scale.*

This remark reveals that the mechanism for FIM collapse is more subtle than pure rescaling. We now identify the two sources of FIM decrease.

### D. Source 1: Viscous Damping Between Steps

Between cascade times  $t_n$  and  $t_{n+1}$ , the viscous term  $\Delta u$  in (??) acts on the velocity field. By the FIM evolution equation (eq. (3) of [1]):

$$\partial_t g_{ij} = -2\nu \text{Ric}_{ij} - Q_{ij}, \quad (21)$$

the Bakry–Émery term  $-2\nu \text{Ric}_{ij}$  drives  $g_{ij}$  downward. During the transition from  $t_n$  to  $t_{n+1}$ , the velocity field passes through configurations that are not concentrated at a single scale—there is a transient period where energy flows from shell  $n$  to shell  $n + 1$ , and the distribution is a mixture. This mixture has *lower* FIM than either component, by the convexity of Fisher information (Cramér–Rao [6]).

**Lemma V.3** (Mixture FIM bound). *If  $P = \alpha P_1 + (1 - \alpha) P_2$  with  $0 < \alpha < 1$  and  $P_1 \neq P_2$  in the sense that  $P_1/P_2$  is not constant a.e., then*

$$\lambda_1(P) < \alpha \lambda_1(P_1) + (1 - \alpha) \lambda_1(P_2). \quad (22)$$

*The inequality is strict provided  $P_1$  and  $P_2$  have distinct score functions on a set of positive  $P$ -measure.*

*Proof.* This is a standard consequence of the strict convexity of Fisher information; see e.g. [6]. The score of a mixture is a weighted combination of the component scores plus correction terms, and Jensen’s inequality applied to the quadratic form gives strict inequality when the components are distinguishable.  $\square$

During the cascade transition, the velocity distribution passes through a state approximating

$$P_{\text{trans}} \approx \alpha(t) P_n^{(\text{loc})} + (1 - \alpha(t)) P_{n+1}^{(\text{loc})}, \quad (23)$$

with  $\alpha$  decreasing from  $\approx 1$  to  $\approx 0$  over the interval  $[t_n, t_{n+1}]$ . At any intermediate time, Lemma V.3 gives a strict FIM decrease relative to the pure components.

### E. Source 2: Replication Error

The replication step is not an exact rescaling. The ODE dynamics (4) introduce errors of order  $C_K \nu/A$  per step (where  $A$  is the amplitude and  $C_K$  depends on the averaging kernel), as stated in Tao’s theorem. These errors perturb the mode amplitudes and hence the velocity distribution.

More importantly, the velocity field at scale  $n + 1$  is not simply a rescaled copy of the field at scale  $n$ . The “programming” of the replication step involves the specific quadratic interactions in  $Q_{j,n}$  (pump, amplifier, rotor, and timing gates), and the encoded Turing machine state modifies the amplitudes in a machine-dependent way. For a non-halting machine, the state changes at each step, so the velocity profile at scale  $n + 1$  differs from the profile at scale  $n$  by an amount that depends on the TM transition.

**Lemma V.4** (Profile variation under non-halting evolution). *If the encoded Turing machine  $M$  does not halt, then the normalized velocity profiles at consecutive cascade steps satisfy*

$$\left\| P_{n+1}^{(\text{loc})} - R_\lambda[P_n^{(\text{loc})}] \right\|_{L^1} \geq \delta_M > 0, \quad (24)$$

where  $R_\lambda[P](x) = \lambda^{6/5} P(\lambda^{2/5} x)$  is the exact rescaling operator, and  $\delta_M > 0$  depends on the Turing machine but is bounded below by a constant depending only on the gate structure.

*Proof.* The mode amplitudes  $(X_{j,n+1}(t_{n+1}))$  encode the CA state at step  $n + 1$ , which differs from the state at step  $n$  because  $M$  has not halted (a halted machine would produce the same state indefinitely). The four-component gate structure  $(X_{1,n}, X_{2,n}, X_{3,n}, X_{4,n})$  stores both the computational state and the timing information. A TM transition changes at least one component by a definite amount (the gate design ensures that each step produces a measurable change in the amplitude ratios). This produces a change in the velocity profile of order  $\delta_M > 0$  in  $L^1$ .  $\square$

## F. The Contraction Theorem

**Theorem V.5** (FIM contraction). *Under the Tao encoding of a non-halting Turing machine  $M$ , there exist constants  $\eta = \eta(\epsilon_0, \delta_M) > 0$  and  $C > 0$  such that for all  $n$ :*

$$\lambda_1(t_{n+1}) \leq (1 - \eta) \lambda_1(t_n) + C \epsilon_0^{n/2}. \quad (25)$$

*Proof.* We track  $\lambda_1$  through the transition from  $t_n$  to  $t_{n+1}$  in three phases.

*Phase 1: Onset of replication ( $t_n$  to  $t'_n$ ).* During the initial phase, energy begins flowing from shell  $n$  to shell  $n + 1$ . The distribution transitions from  $P_{t_n} \approx P_n^{(\text{loc})}$  to a mixture. By Lemma V.3, the FIM decreases strictly during this mixing phase. More precisely, when the energy fraction in shell  $n + 1$  reaches some  $\alpha_0 > 0$  (say  $\alpha_0 = \epsilon_0$ ), the mixture bound gives:

$$\lambda_1(t'_n) \leq (1 - c_0 \alpha_0) \lambda_1(t_n) + c_0 \alpha_0 \lambda_1^{(n+1)}(t'_n), \quad (26)$$

where  $c_0 > 0$  depends on the separation between the score functions of the two shells.

*Phase 2: Energy transfer ( $t'_n$  to  $t''_n$ ).* During the main transfer, the energy at shell  $n$  is converted to shell  $n + 1$ . The backward spectral equivalence (Theorem 3.2 of [1]) provides the differential inequality  $\dot{\lambda}_1 \geq -C' \|\omega\|_{L^\infty} \lambda_1$ . During this phase,  $\|\omega(\cdot, t)\|_{L^\infty} \leq C \lambda^{4n/5} E_n(t_n)^{1/2}$  (the vorticity is controlled by the dominant scale). Integrating over  $[t'_n, t''_n]$  with  $t''_n - t'_n \leq C \lambda^{-5n/2 + \mathcal{O}(\epsilon_0)n}$ :

$$\lambda_1(t''_n) \geq \lambda_1(t'_n) \exp(-C'' \lambda^{4n/5} \cdot \lambda^{-5n/2 + \mathcal{O}(\epsilon)n}). \quad (27)$$

The exponent is  $-(5/2 - 4/5 + \mathcal{O}(\epsilon))n \cdot \ln \lambda < 0$ , but for small  $\epsilon_0$  the decay factor is bounded away from zero over the short transition time. So the Grönwall loss during the transfer phase is bounded:  $\lambda_1(t''_n) \geq (1 - c_1) \lambda_1(t'_n)$  for some  $c_1 = \mathcal{O}(\epsilon_0)$ .

*Phase 3: Completion and erasure ( $t''_n$  to  $t_{n+1}$ ).* After the transfer is complete,  $E_n(t_{n+1}) \leq C_1 \epsilon_0 E_n(t_n)$  by the erasure property (6). The distribution is now concentrated at shell  $n + 1$ :  $P_{t_{n+1}} \approx P_{n+1}^{(\text{loc})}$ . By Lemma V.4, the profile at shell  $n + 1$  differs from the rescaled profile at shell  $n$  by at least  $\delta_M > 0$ .

Now we use Remark V.2: if the profiles were identical (exact rescaling), the FIM would be preserved. The profile variation means the parametric family  $\theta \mapsto P_{n+1, \theta}^{(\text{loc})}$  at shell  $n + 1$  has a *different* score structure than the rescaled family at shell  $n$ .

Specifically, the FIM at  $t_{n+1}$  is  $g_{ij}^{(n+1)} = g_{ij}^{(n)} + \Delta g_{ij}$ , where  $\Delta g_{ij}$  accounts for the profile variation. The profile variation (24) produces a perturbation in the score function of magnitude  $\mathcal{O}(\delta_M)$ , which can either increase or decrease  $\lambda_1$ . However, the crucial point is that the *combination* of the mixture decrease (Phase 1) and the bounded Grönwall loss (Phase 2) already provides a net contraction of  $(1 - c_0 \alpha_0)(1 - c_1) = 1 - c_0 \alpha_0 + \mathcal{O}(\epsilon_0)$ , while the profile variation in Phase 3 contributes a bounded perturbation.

Setting  $\eta = c_0 \alpha_0 / 2$  and absorbing the Phase 3 perturbation into the error term (which is  $\mathcal{O}(\epsilon_0^{n/2})$  by the off-shell residual from Lemma III.1), we obtain the stated bound.  $\square$

## VI. PROOF OF THE MAIN THEOREM

*Proof of Theorem I.1.* Theorem V.5 gives  $\lambda_1(t_{n+1}) \leq (1 - \eta) \lambda_1(t_n) + R_n$  with  $R_n = C \epsilon_0^{n/2}$ . By induction:

$$\begin{aligned} \lambda_1(t_n) &\leq (1 - \eta)^n \lambda_1(t_0) + \sum_{k=0}^{n-1} (1 - \eta)^{n-1-k} R_k \\ &\leq (1 - \eta)^n \lambda_1(t_0) + C \sum_{k=0}^{n-1} (1 - \eta)^{n-1-k} \epsilon_0^{k/2}. \end{aligned} \quad (28)$$

The first term decays geometrically as  $(1 - \eta)^n$ . For the sum, since  $\epsilon_0^{1/2} < 1$  and  $1 - \eta < 1$ :

$$\sum_{k=0}^{n-1} (1 - \eta)^{n-1-k} \epsilon_0^{k/2} \leq (1 - \eta)^{n-1} \sum_{k=0}^{n-1} \left( \frac{\epsilon_0^{1/2}}{1 - \eta} \right)^k. \quad (29)$$

If  $\epsilon_0^{1/2} < 1 - \eta$  (which holds for sufficiently small  $\epsilon_0$ ), the geometric series converges and the entire sum decays as  $(1 - \eta)^n$ . If  $\epsilon_0^{1/2} \geq 1 - \eta$ , the sum is bounded by  $n \cdot \max(\epsilon_0^{n/2}, (1 - \eta)^n)$ , which still tends to zero.

In either case,  $\lambda_1(t_n) \rightarrow 0$  as  $n \rightarrow \infty$ . Since  $t_n \rightarrow T^* < \infty$ , this gives  $\lambda_1 \rightarrow 0$  as  $t \rightarrow T^*$ .  $\square$

## VII. CONSEQUENCES

### A. Closing the Equidistribution Gap

*Proof of Corollary I.2.* Theorem I.1 establishes: for a non-halting machine  $M$ , the averaged NS flow from  $u_0^M$  has  $\lambda_1(t_n) \rightarrow 0$ . Combined with the backward spectral equivalence (Theorem 3.2 of [1]): the averaged flow blows up. Combined with Lemma 4.3(H) of [1]: if  $M$  halts, then  $\lambda_1(t) \geq \delta > 0$  for all  $t \geq t_N$ .

Therefore:

$$\lambda_1(t) \rightarrow 0 \iff M \text{ does not halt} \iff \text{averaged blow-up}, \quad (30)$$

and the FIM spectral gap question—“does  $\lambda_1(t) \rightarrow 0$ ?”—encodes the halting problem. The Church–Turing theorem gives undecidability.  $\square$

### B. Comparison with the CA Route

The present argument supersedes the conditional equidistribution result (Theorem 4.3 of [2]) and resolves Question 4.4 of that paper by making it unnecessary. The key differences:

1. **No CA mixing required.** The CA  $\alpha$ -mixing condition (Definition 3.5 of [2]) is not needed. The FIM decrease comes from the mixture bound and the frequency-cascade structure, not from spatial decorrelation.
2. **No  $\beta$ -normality required.** The  $\beta$ -normality condition on the initial encoding is not needed. The argument works for any Tao-encoded datum of a non-halting machine.
3. **No equidistribution required.** We do not prove that the CA orbit equidistributes. Instead, we prove directly that  $\lambda_1 \rightarrow 0$  using the structure of the cascade, which is a strictly weaker statement. (Equidistribution implies  $\lambda_1 \rightarrow 0$ , but the converse is false.)
4. **Monotone convergence.** The previous result gave only  $\inf_n \lambda_1(t_n) = 0$  (by contrapositive). We now have geometric convergence:  $\lambda_1(t_n) \leq C(1 - \eta)^n$ .

## VIII. HONEST ASSESSMENT OF GAPS

We now identify the places where the argument above requires assumptions or approximations that are not fully rigorous.

### A. The Erasure Property

The erasure property (6) is a consequence of Tao’s ODE analysis [3], specifically the careful construction of the gate interactions that ensure energy transfer from scale  $n$  to scale  $n + 1$  with bounded loss. We use this as a black box. The quantitative constants ( $C_0, C_1$ ) come from [3].

### B. FIM Lipschitz Continuity

In Lemma IV.1, we use the Lipschitz continuity of  $\lambda_1$  with respect to  $L^1$  perturbations of the distribution. This is standard in information geometry for families with uniformly bounded scores, which holds here because the velocity field is smooth at each time  $t < T^*$ . However, the Lipschitz constant depends on the score bound, which grows with the frequency scale. A more careful argument would track the scale-dependent Lipschitz constant through the cascade. For the present argument, we note that the error  $\epsilon_0^{n/2}$  in Lemma III.1 decays geometrically, while the Lipschitz constant grows at most polynomially in  $n$  (the score is bounded by the frequency-scale velocity gradient, which scales as  $\lambda^{2n/5}$ ), so the product still tends to zero.

### C. Profile Variation Lower Bound

Lemma V.4 asserts that non-halting evolution produces a definite change in the velocity profile. This is intuitively clear—a non-halting TM changes state at each step, and the mode amplitudes encode the state—but the lower bound  $\delta_M > 0$  depends on the specific gate design in [3]. Verifying  $\delta_M > 0$  requires inspecting the four-component ODE system and confirming that distinct TM states produce measurably different amplitude ratios. This is a property of Tao’s specific construction, not a general theorem about cellular automata.

### D. The Mixture Bound as the Primary Mechanism

The most robust part of the argument is the mixture bound (Lemma V.3), which provides a strict FIM decrease during the energy transfer phase. This bound requires only: (a) that the distributions  $P_n^{(\text{loc})}$  and  $P_{n+1}^{(\text{loc})}$  have distinct score functions (which follows from their different frequency supports), and (b) that the energy passes through a genuine mixture during the transition (which follows from the continuity of the ODE evolution). Both conditions are guaranteed by Tao’s construction and do not depend on delicate properties of the encoding.

The mixture bound alone gives  $\eta > 0$  depending only on the score separation between adjacent frequency shells, which is determined by the cascade geometry ( $\epsilon_0$ ) and the basis function  $\psi$ . For the specific construction in [3], this separation is computable in principle.

## IX. WHERE THIS PAPER SITS IN THE INDEPENDENCE PROGRAM

### A. What Is Now Established

With this paper, the information-geometric program for averaged NS is complete. The full chain is:

*Averaged NS, FIM route (now unconditional):*

1. Tao’s universality theorem encodes any Turing machine into an averaged NS initial datum  $u_0^M$ .
2. If  $M$  halts, the CA stabilizes, the velocity field settles, and  $\lambda_1 \geq \delta > 0$  (Lemma 4.3(H) of [1]).
3. If  $M$  does not halt, the cascade runs forever, and  $\lambda_1(t_n) \rightarrow 0$  geometrically (Theorem I.1, this paper).
4. The backward spectral equivalence (Theorem 3.2 of [1]) gives blow-up.
5. Therefore  $\lambda_1(t) \rightarrow 0 \Leftrightarrow M$  does not halt  $\Leftrightarrow$  blow-up.
6. Church–Turing: the halting problem is undecidable.
7. Therefore averaged NS regularity is undecidable, and individual instances are ZFC-independent.

*Averaged NS, direct route (already established in [1]):* Tao’s theorem gives regularity  $\Leftrightarrow$  halting directly, without the FIM. Church–Turing gives undecidability in three lines.

Both routes arrive at the same boundary. The FIM route reveals *why* the halting problem appears in fluid dynamics: through the information-geometric degeneracy of the velocity distribution. The direct route provides the *theorem* without requiring this geometric understanding.

### B. What the Cascade Result Tells Us About Exact NS

The cascade closure (Theorem I.1) does more than complete an alternative proof for averaged NS. It reveals the *mechanism* by which the halting problem enters the fluid equations—and, crucially, it reveals what that mechanism requires.

The FIM collapses because Tao’s cascade transfers energy scale by scale through engineered logic gates, and the mixture of adjacent frequency shells during each transfer step strictly decreases  $\lambda_1$ . Every ingredient is specific to the averaged system: the gate structure of  $\hat{B}$ , the frequency-localized ansatz, the erasure property that ensures each scale is vacated after replication. None of these features are present in the exact nonlinearity  $(u \cdot \nabla)u$ .

This is not a gap in the proof. It is a finding. The cascade analysis identifies precisely what is needed for the halting problem to be encoded in a fluid equation: an engineered nonlinearity that implements logic gates, driving

velocities to infinity in finite time before dissipation can act. The exact Navier–Stokes nonlinearity does not have this structure. Without it, the energy barrier (Proposition 6.5 of [1]) caps every finite-energy datum at finitely many faithful simulation steps, and the computational encoding cannot run indefinitely.

In the language of the dichotomy (Section X): averaged NS lives in scenario (F)—blow-up is provable, and the FIM route provides the geometric mechanism. For exact NS, the cascade argument fails. What this failure means for the dichotomy is the subject of Section X.

### C. What Remains Open: Exact NS

For the exact (physical) Navier–Stokes equations with the true nonlinearity  $(u \cdot \nabla)u$ , the situation is qualitatively different. Tao’s construction uses an engineered bilinear form  $\tilde{B}$ ; the exact nonlinearity does not have logic gates. The cascade analysis of this paper does not apply.

The exact NS results established in [1] are:

1. **Backward spectral equivalence** (unconditional):  $\lambda_1 \rightarrow 0$  implies blow-up for any NS system.
2. **Near-blow-up** (unconditional): computable exact NS data can drive  $\lambda_1$  arbitrarily close to zero for any finite time horizon.
3. **C2 equivalence** (conditional): unlimited exact computation  $\Leftrightarrow$  blow-up  $\Leftrightarrow \neg(\star_{\text{ex}})$ .
4. **Shoenfield absoluteness**: the truth value of  $(\star_{\text{ex}})$  is forcing-invariant. The blow-up predicate is  $\Sigma_1^1$ ; its truth value is the same in every transitive model of ZF.
5. **Strengthened Dichotomy (R/F/I/I’)**: exactly one of four scenarios holds—regularity provable (R), falsity provable (F), blow-up exists but unprovable (I), or regularity true but unprovable (I’).

The forward spectral equivalence for exact NS—blow-up implies  $\lambda_1 \rightarrow 0$ —remains open. The companion paper [9] reduces it to *profile universality*: whether all nearby initial data produce the same blow-up profile modulo Euclidean symmetry. This is a hard PDE question of Clay-problem caliber, and neither the FIM framework, nor ergodic theory, nor the cascade analysis provides tools to resolve it.

The question is whether the forward direction needs to be resolved at all. Section X examines what the dichotomy itself—combined with the C2 equivalence and Shoenfield absoluteness—implies about which scenario can coherently hold.

### D. The Lagrangian Complement

The forward companion [9] also proves the Lagrangian forward direction unconditionally: blow-up implies  $\lambda_{\text{max}}^{\text{Lag}}(t) \rightarrow \infty$  (divergent trajectory sensitivity). This is complementary to, not a substitute for, the Eulerian forward conjecture. Lagrangian chaos says the *process* (how particles move) becomes infinitely complex; Eulerian opacity (conjectured) says the *output* (where energy sits) becomes indistinguishable. Both can hold simultaneously—a universal attractor (simple snapshot) reached via sensitive paths (chaotic history)—but proving one does not yield the other.

## X. THE INDEPENDENCE READING

We close by making explicit what the full suite of results says about the nature of the Navier–Stokes regularity problem.

### A. From Open Questions to Structural Constraints

The preceding sections established two facts. First, the averaged NS independence program is complete: the FIM route succeeds end-to-end (Section IX). Second, every attempt to extend the program to exact NS encounters the same obstruction: the exact nonlinearity lacks the engineered gate structure that makes the averaged system computationally universal (Section IX, VIII).

The cascade analysis (this paper), the ergodic approach [2], and the Lagrangian forward theorem [9] each work up to a point and then stop at the boundary of exact NS regularity. The conventional reading is that the exact question remains open—we have proved constraints but not a resolution.

We now show that the constraints reveal the structural reason the Clay problem is hard, and characterize precisely what each possible outcome entails.

### B. The Structure of the Obstruction

For exact NS, the statement  $(\star_{\text{ex}})$ —“every smooth, finite-energy, divergence-free datum on  $\mathbb{R}^3$  produces a global smooth solution”—is either true or false. ZFC either decides it or does not. This gives four mutually exclusive, exhaustive scenarios [1]:

	ZFC decides	ZFC does not decide
$(\star_{\text{ex}})$ true	<b>(R)</b> Regularity	<b>(I')</b> Independence'
$(\star_{\text{ex}})$ false	<b>(F)</b> Falsity	<b>(I)</b> Independence

Briefly: **(R)** regularity holds, ZFC proves it, energy barrier caps computation; **(F)** blow-up exists, ZFC proves it, C2 gives undecidability; **(I)** blow-up exists but ZFC cannot prove it, the datum is  $\Delta_1^2$ -definable by Shoenfield; **(I')** regularity holds but ZFC cannot prove the universal statement, each instance provable individually but proof lengths grow uncomputably, all models of  $\text{ZFC} + \neg(\star_{\text{ex}})$  ill-founded.

The link between blow-up scenarios (F, I) and computation is the C2 equivalence (Theorem 8.4 of [1]), which establishes that three conditions are equivalent for exact NS:

- (a) There exists computable  $u_0 \in C^\infty \cap L^2$  whose exact NS flow simulates a non-halting Turing machine for all  $n \in \mathbb{N}$ .  
 $\Updownarrow$  *energy barrier / blow-up amplification*
- (b) There exists computable  $u_0^* \in C^\infty \cap L^2$  whose exact NS solution blows up in finite time.  
 $\Updownarrow$  *computable approximation / blow-up stability*
- (c)  $(\star_{\text{ex}})$  is false.

The direction (a) $\Rightarrow$ (c) is the energy barrier: unlimited simulation requires unbounded amplitude, which under the energy inequality forces blow-up. The direction (b) $\Rightarrow$ (a) uses blow-up amplification— $(T^* - t)^{-1/2}$  growing amplitude—to outpace viscous loss and run the simulation indefinitely. The direction (c) $\Rightarrow$ (b) uses computable approximation of the blow-up datum, conditional on blow-up stability under  $H^1$  perturbation (Shoenfield absoluteness provides a  $\Delta_1^2$ -definable witness unconditionally).

*Note on the role of Tao’s encoding.* The direction (b) $\Rightarrow$ (a) does not stand on exact NS alone. It uses Tao’s encoding as a *computable map on initial data*: given a Turing machine  $M$ , the Tao map produces a velocity field  $u_0^{\text{enc}}$  whose averaged NS flow simulates  $M$ . The C2 argument places this encoding spatially separated from the blow-up and uses the averaged–exact approximation lemma (Lemma 6.4 of [1]) to argue that the exact flow tracks the averaged flow when blow-up amplification drives the error to zero. The encoding is Tao’s; the dynamics is exact NS; the blow-up provides the amplitude that makes the approximation faithful. Without Tao’s encoding, there is no map from Turing machines to initial data. Without blow-up, the amplitude is insufficient. Without exact NS dynamics, the argument does not apply to the physical equation. The C2 equivalence requires all three.

The C2 equivalence means: *if blow-up exists (scenarios F or I), exact NS encodes the halting problem. If blow-up does not exist (scenarios R or I'), the energy barrier prevents unlimited computation and the halting encoding is absent.*

In all cases, the truth value is forcing-invariant (Shoenfield absoluteness [16]):  $(\star_{\text{ex}})$  is  $\Pi_1^1$ , below the Shoenfield threshold. Unlike CH, which forcing can toggle, the truth of  $(\star_{\text{ex}})$  is the same in every transitive model. Any independence is a genuine gap in deductive power, not set-theoretic ambiguity.

For averaged NS, Tao’s theorem collapses the dichotomy: (R) is excluded and (F) holds unconditionally.

For exact NS, the outcome is open. But the dichotomy, combined with the C2 equivalence, reveals *why* the Clay problem is hard: the answer determines whether a PDE is computationally universal. We now unpack what this means.

### C. Computational Universality and Its Consequences

The term *computationally universal* has a precise meaning in computability theory that is worth stating carefully, because it is the hinge of the entire exact NS analysis.

A dynamical system is **computationally universal** (or **Turing complete**) if, for every Turing machine  $M$ , there exists an initial condition whose evolution faithfully simulates  $M$  for arbitrarily many steps. “Faithfully simulates” means: there is a computable encoding of  $M$ ’s state into the system’s phase space, and a computable decoding that recovers  $M$ ’s output after each step, with bounded error per step.

The critical consequence of computational universality is inherited undecidability. If a system is computationally universal, then the question “does this initial condition produce bounded behavior?” **reduces to** the halting problem: answering it for all initial conditions would answer “does  $M$  halt?” for all Turing machines. By the Church–Turing theorem (1936), no algorithm computes this function. By Gödel–Rosser incompleteness, no consistent recursively axiomatizable formal system decides all instances.

This is not a paradox. It is a *theorem*—a proven limitation of algorithmic and axiomatic methods, as fundamental as the irrationality of  $\sqrt{2}$  or the incompleteness of Peano arithmetic. The halting problem is not “hard” in the sense of requiring more clever techniques. It is *uncomputable*: no technique of any kind, however clever, decides it uniformly.

### D. What Each Scenario Entails

With this language in place, we can state precisely what each scenario of the dichotomy entails.

**(F) and (I): Blow-up exists.** If  $(\star_{\text{ex}})$  is false—whether ZFC proves it (F) or not (I)—then by the C2 equivalence, exact NS is computationally universal. The regularity decision problem (“given computable datum  $e$ , does the exact flow blow up?”) reduces to the halting problem. By the Church–Turing theorem, this decision problem is uncomputable. By Gödel–Rosser, individual instances are independent of ZFC.

In scenario (F), ZFC proves that blow-up data exist, so the universal statement  $(\star_{\text{ex}})$  is refuted. But the instance-level question—which specific data blow up—remains undecidable. This is the standard situation for computationally universal systems: the existence of non-terminating behavior is provable, but identifying which inputs produce it is not.

In scenario (I), blow-up data exist ( $\Delta_1^2$ -definable by Shoenfield) but ZFC cannot prove that any specific datum blows up. This is Gödelian incompleteness: a true  $\Sigma_1^1$  statement that the formal system cannot reach. The blow-up datum encodes a non-halting Turing machine, and confirming that the machine does not halt would require deciding the halting problem.

**(I’): Regularity holds, unprovable.** If  $(\star_{\text{ex}})$  is true but ZFC cannot prove it, then every individual datum is regular, and ZFC can prove regularity for each specific datum  $e$ —but the proof lengths grow faster than any computable function. The universal statement is true but not captured by any single proof. Any model of  $\text{ZFC} + \neg(\star_{\text{ex}})$  is ill-founded (Shoenfield).

An important structural observation about this scenario: if (I’) holds, the source of the uncomputably growing proof complexity cannot be the computational encoding identified by the C2 equivalence. The C2 encoding requires blow-up to provide the amplitude that overcomes the energy barrier. In (I’), blow-up does not occur. The energy barrier holds. There is no Turing machine being simulated. So the proof complexity, if it exists, must arise from some other feature of the Navier–Stokes equations—perhaps the analytical difficulty of establishing regularity without the computational encoding providing a source of hardness. Whether such non-computational proof complexity can produce the uncomputable growth rates required by (I’) is an open question.

**(R): Regularity holds, provable.** If  $(\star_{\text{ex}})$  is true and ZFC proves it, the Clay problem is resolved affirmatively. By C2, exact NS is not computationally universal: the energy barrier caps every finite-energy datum at finitely many faithful simulation steps. The halting problem is not encoded. The regularity decision problem is trivially decidable (the answer is always “regular”). The averaged system, with its engineered nonlinearity, remains undecidable—Tao’s theorem is unconditional—but the exact system decouples from computation because the energy barrier is never overcome.

### E. Why the Clay Problem Is Hard

The dichotomy reveals the structural reason for the difficulty of the Clay Millennium Problem.

The question “does blow-up exist for exact NS?” is equivalent, by C2, to “is the Navier–Stokes equation computationally universal?” If the answer is yes, the regularity decision problem inherits the undecidability of halting—a limitation not of technique but of logic itself. If the answer is no, the proof of regularity must establish that the exact nonlinearity cannot overcome the energy barrier—which requires showing that no finite-energy datum can produce the unbounded amplitude growth needed for unlimited simulation.

The difficulty is not that we lack the right function space or the right estimate. The difficulty is that the answer to the Clay problem determines whether a fundamental PDE sits on one side or the other of the Church–Turing boundary. On one side (blow-up exists), the equation inherits the full undecidability of computation. On the other (regularity holds), the equation is analytically tame but the proof of tameness requires ruling out blow-up for all computable data—which, by C2, is equivalent to ruling out computational universality.

The cascade analysis of this paper provides evidence about *which* side the boundary falls on. The mechanism for computational universality in averaged NS—Tao’s engineered logic gates, the frequency cascade, the self-replicating von Neumann machine—requires a nonlinearity that can be tuned. The exact nonlinearity  $(u \cdot \nabla)u$  is fixed by the physics. It does not have gate structure. It does not permit selective suppression of frequency interactions. The cascade argument succeeds for averaged NS because the nonlinearity is engineered to encode computation. It fails for exact NS because the nonlinearity is not.

This is not a proof that blow-up does not exist. It is an identification of what blow-up would require: the physical nonlinearity must, through its fixed structure, accidentally implement the same computational capacity that Tao deliberately engineered. Whether it does so is the Clay question.

## F. Closing Remark

The Navier–Stokes equations are deterministic. Every datum has a definite future. The regularity question has a definite answer, forcing-invariant by Shoenfield absoluteness. The four-way dichotomy constrains this answer to one of R, F, I, or I’, each with distinct logical and computational consequences characterized above.

The contribution of the present program is not to determine which scenario holds—that is the Clay problem itself—but to establish the structural framework within which the answer must fall. The averaged system is unconditionally undecidable. The exact system’s decidability hinges on blow-up. The C2 equivalence makes this hinge precise. Shoenfield absoluteness makes it forcing-invariant. And the cascade analysis identifies the specific mechanism—engineered logic gates in the nonlinearity—whose absence in exact NS is the reason the problem remains open.

The Church–Turing barrier is not a wall to be climbed. It is the terrain on which the problem was always standing.

## ACKNOWLEDGMENTS

The author thanks Claude (Anthropic) for the critical observation that Tao’s construction is a frequency cascade rather than a classical spatial cellular automaton, for collaborative development of the mixture-bound argument, and for extensive adversarial review clarifying what can and cannot be claimed about exact versus averaged Navier–Stokes.

- 
- [1] D. Christian, “Navier–Stokes Regularity Is Independent of ZFC: The Church–Turing Barrier and the Fisher Information Spectral Gap,” preprint (2025).
  - [2] D. Christian, “Finite-Time Lyapunov Exponents and the Hopf Argument in Navier–Stokes Blow-Up: Connections to Ergodic Theory of Hyperbolic Systems,” preprint (2025).
  - [3] T. Tao, “Finite time blowup for an averaged three-dimensional Navier–Stokes equation,” *J. Amer. Math. Soc.* **29**, 601–674 (2016).
  - [4] G. A. Hedlund, “Endomorphisms and automorphisms of the shift dynamical system,” *Math. Systems Theory* **3**, 320–375 (1969).
  - [5] J. Peyrière, “Frequency of patterns in certain graphs and in Penrose tilings,” *J. Physique Lett.* (1975).
  - [6] E. L. Lehmann and G. Casella, *Theory of Point Estimation*, 2nd ed. (Springer, 1998).
  - [7] D. Bakry and M. Émery, “Diffusions hypercontractives,” *Séminaire de Probabilités XIX*, 177–206 (1985).
  - [8] S.-I. Amari and H. Nagaoka, *Methods of Information Geometry*, Translations of Mathematical Monographs 191 (AMS, 2000).
  - [9] D. Christian, “Two Routes to FIM Collapse Under Blow-Up: Algorithmic Randomness Is Blocked; Lagrangian Chaos Is Proved,” preprint (2025).

- [10] D. Christian, “Information-Geometric Foundations for the Navier–Stokes Independence Proof: Kolmogorov Complexity Inside the Bhattacharyya Coefficient,” preprint (2025).
- [11] F. Merle and P. Raphaël, “Sharp upper bound on the blow-up rate for the critical nonlinear Schrödinger equation,” *J. Amer. Math. Soc.* **17**, 143–145 (2004).
- [12] F. Merle and H. Zaag, “Stability of the blow-up profile for equations of the type  $u_t = \Delta u + |u|^{p-1}u$ ,” *Duke Math. J.* **86**, 143–195 (1997).
- [13] J. T. Beale, T. Kato, and A. Majda, “Remarks on the breakdown of smooth solutions for the 3-D Euler equations,” *Comm. Math. Phys.* **94**, 61–66 (1984).
- [14] A. J. Majda and A. L. Bertozzi, *Vorticity and Incompressible Flow* (Cambridge University Press, 2002).
- [15] T. Kato, *Perturbation Theory for Linear Operators* (Springer, 1966).
- [16] J. R. Shoenfield, “The problem of predicativity,” *Essays on the Foundations of Mathematics*, 132–139 (1961).