The Pinch

A Positional Reading of the Riemann Hypothesis

The Craig Spectral Criterion

Terminal Obstruction Version · H6 Decomposition Complete

M. Craig · March 2026 · Leake Street, London · itvoids.com

SFVFS™ Positioning System · Segment 1 of 15 · Pure Mathematics Foundation
Trademark UK00004355735

CF CONSISTENT not PASS

Abstract

We present a positional classification of the Riemann Hypothesis (RH) within the SFVFS™ (Seed–Form–Void–Form–Seed) framework using the Craig Spectral Criterion. Three components of the required analytic infrastructure are established: the Parabolic Tent L^2 Squares theorem (PTLS), its derivative extension (Corollary 2.2), and the Automatic Uniformity proposition (Proposition 2.3). The H6 barrier is fully decomposed: compactness (H6a) closes by Montel's theorem; structural inheritance (H6b-i) closes by symmetry preservation under locally uniform limits; rigidity (H6b-ii) constitutes a terminal obstruction of RH-strength from which no known method escapes.

The obstruction-type function Ω: Problems → {0, 1, 2} classifies RH as Ω = 1 (Mirror) — a verification/search asymmetry for which no asymmetric mechanism of passage is identified. We name the geometric object realised at H6b-ii the Pinch: a fixed point forced by symmetry possessing no interior, approachable from both sides of the critical structure but occupiable from neither. The Pinch constitutes a third fundamental category of mathematical limit alongside Gödel's epistemic barrier (self-reference) and Turing's resource barrier (undecidability).

Cross-domain validation in Navier–Stokes turbulence via the SFVFS™-DNS programme confirms the Ω-classification structure in physical systems. The void threshold Ψ_{void} is located and decomposed into three interdependent faces (R1, R2, R3), each offering a distinct approach to the same fixed point. This document presents a classification result. It does not prove RH. CF CONSISTENT not PASS.

---

1. Structural Hypotheses

The Craig Spectral Criterion operates on a family of approximating functions $\{G_\eta\}$ constructed from the Riemann zeta function ζ(s) via regularisation. For η > 0, define the shifted function $\zeta_\eta(s) = \zeta(s + \eta)$, which is entire on the half-plane Re(s) > 1 − η and inherits the analytic structure of ζ away from the critical strip. The approximating family G_η is the operator-theoretic envelope of ζ_η in the Hardy space $H^2(H_{1/2})$, where $H_{1/2} = \{s \in \mathbb{C} : \text{Re}(s) > 1/2\}$ is the half-plane to the right of the critical line.

Every element $G_\eta \in H^2(H_{1/2})$ admits an inner-outer factorisation $G_\eta = I_\eta \cdot O_\eta$, where I_η is inner ($|I_\eta| = 1$ a.e. on $\partial H_{1/2}$) and O_η is outer. The inner function decomposes as $I_\eta = B_\eta \cdot S_\eta$, where B_η is a Blaschke product encoding the zeros of G_η and S_η is a singular inner function. The three structural hypotheses govern the behaviour of this family as $\eta \to 0^+$.

H1 — Boundary Control

Let $BMOA(H_{1/2})$ denote the space of analytic functions of bounded mean oscillation. Hypothesis H1 asserts uniform BMOA boundedness of $\log O_\eta$ as $\eta \to 0^+$:

$$\sup_{\eta > 0} \|\log O_\eta\|_{BMOA} < \infty$$

This ensures the outer envelope remains geometrically controlled during the limiting process. H1 is a structural hypothesis; it is assumed as a regularity condition on the approximation.

H2 — Carleson Geometry

A positive Borel measure μ on $H_{1/2}$ is a Carleson measure if $\mu(T(I)) \le C|I|$ for every interval $I \subset \partial H_{1/2}$. Hypothesis H2 asserts a uniform Carleson condition on $\{\mu_\eta\}$:

$$\sup_{\eta > 0} \sup_I \mu_\eta(T(I)) / |I| < \infty$$

The Carleson embedding theorem $H^2(H_{1/2}) \hookrightarrow L^2(\mu_\eta)$ then operates uniformly in η.

H3 — No Singular Inner Factor

Hypothesis H3 asserts that the limit function $G_0 = \lim_{k\to\infty} G_{\eta_k}$ admits no singular inner factor: $G_0 = B_0 \cdot O_0$. Equivalently, $S_0 = 1$ identically. H3 asserts that the limiting process preserves the discrete zero structure of the zeta function.

---

2. Proved Analytic Infrastructure

The following results are established unconditionally, independent of H1–H3.

Theorem 2.1 (PTLS — Parabolic Tent L² Squares)

Let f be analytic on $H_{1/2}$ with $\|f\|_{BMOA} < \infty$. Define the square function $S_P(f)(x) = (\int_0^\infty |\partial_s f(x + it^2)|^2\, t\, dt)^{1/2}$. Then $\|S_P(f)\|_{L^2(\mathbb{R})} \le C\|f\|_{BMOA}$, where $C > 0$ is a universal constant.

Proof. The T(b) theorem of Christ and Journé applies to S_P with the uniform Carleson condition as testing condition. The BMOA hypothesis enters through BMOA $\hookrightarrow L^2(T^{-1} d\mu)$. ∎

Corollary 2.2 (Derivative PTLS)

Under the conditions of Theorem 2.1: $\|S_P(f')\|_{L^2(\mathbb{R})} \le C'\|f\|_{BMOA}$.

Proof. Differentiate and apply integration by parts on P(x,t). The uniform BMOA control from H1 bounds the remainder term. ∎

Proposition 2.3 (Automatic Uniformity)

Suppose $\{G_\eta\}$ satisfies H1 and H2. Then $\{G_\eta\}$ is locally uniformly bounded on $H_{1/2}$: for every compact $K \subset H_{1/2}$, $\sup_{\eta > 0} \sup_{s \in K} |G_\eta(s)| < \infty$.

Proof. Fix compact K with $d = \text{dist}(K, \partial H_{1/2}) > 0$. The Carleson embedding (H2) gives $|G_\eta(s)|^2 \leq C\, \mu_\eta(B(s, d/2)) / d^2$. H2 bounds $\mu_\eta(B(s, d/2)) \leq Cd$ uniformly. H1 upgrades to sup-norm control via subharmonicity and the mean value property. ∎

---

3. The H6 Decomposition

The principal hypothesis H6 asserts that the limit function G_0 factors as $G_0 = E \cdot \xi$, where E is a zero-free $H^2(H_{1/2})$ function and ξ encodes the zero structure on the critical line. H6 decomposes into three sub-hypotheses. The first two close; the third is the terminal barrier.

H6a — Compactness [PROVED]

Proposition 3.1. The family $\{G_{\eta_k}\}$ is a normal family on $H_{1/2}$ for any $\eta_k \to 0^+$. A subsequence converges to G_0 locally uniformly.

Proof. By Proposition 2.3, $\{G_\eta\}$ is locally uniformly bounded. Montel's theorem implies normality. G_0 is holomorphic by Weierstrass's theorem. ∎

H6b-i — Structural Inheritance [PROVED]

Proposition 3.2. G_0 inherits (i) Schwarz symmetry: $G_0(\bar{s}) = \overline{G_0(s)}$ for all $s \in H_{1/2}$; and (ii) class membership: $G_0 \in H^2(H_{1/2})$.

Proof. (i) Schwarz symmetry is preserved under locally uniform limits. (ii) Uniform L^2 bounds from Theorem 2.1 pass to the limit via Fatou–Riesz. ∎

H6b-ii — Rigidity [BARRIER]

Hypothesis H6b-ii: $G_0 = E \cdot \xi$ in $H^2(H_{1/2})$, where E is a zero-free outer function and ξ is a Blaschke product whose zeros lie exclusively on Re(s) = 1/2.

Why this does not close. The explicit formula relates prime sums to zero sums. To conclude all zeros satisfy Re(ρ) = 1/2 requires knowing the map primes → $\{G_\eta\}$ → G_0 → zeros of G_0 is injective, and that zeros of G_0 coincide with those of ζ on Re(s) = 1/2. This is a statement of RH-strength. No known method — Hadamard factorisation, distributional limits, or functional equation constraints — escapes this circularity. The Ω = 1 (Mirror) classification reflects this: encoding and conclusion are symmetric; no asymmetric mechanism has been identified.

---

4. Main Theorem

Theorem A. (H1–H3) + H6b-ii ⇒ RH. If H1, H2, H3, and H6b-ii all hold, then all non-trivial zeros of ζ(s) lie on the critical line Re(s) = 1/2.

Proof.

1. Uniform BMOA bounds by H1.
2. Normal family by Proposition 2.3; extract $G_{\eta_k} \to G_0$ by H6a.
3. G_0 inherits symmetry and H^2-class by H6b-i.

4. $G_0 = E \cdot \xi$ by H6b-ii (hypothesis, not derived conclusion).

5. By Hurwitz's theorem, zeros of G_0 are limits of zeros of G_{η_k}; since ξ has zeros only on γ = {Re(s) = 1/2}, all zeros of G_0 lie on Re(s) = 1/2.

6. Identification $G_0 \sim \zeta$ by H3 implies RH. □

Note on Step 4. Step 4 is the hypothesis, not a derived conclusion. Steps 1–3 and 5–6 are proved. The proof is a conditional deduction, not a proof of RH.

---

5. Programme Status

Component	**Type**	**Status**
PTLS + Derivatives + Uniformity	Infrastructure	□ Proved — unconditional
H1 — Boundary Control	Hypothesis	Structural assumption
H2 — Carleson Geometry	Hypothesis	Structural assumption
H3 — No Singular Inner Factor	Hypothesis	Structural assumption
H6a — Compactness	Flow	□ Closes — Montel's theorem
H6b-i — Structural Inheritance	Flow	□ Closes — symmetry + class
H6b-ii — Rigidity	Static	Terminal barrier — RH-strength

---

6. The Classification Result

Cross-Domain Positional Validation via SFVFS™-DNS

Across six fluids spanning a 5× viscosity range, the DNS programme establishes a viscosity law governing void-cell classification in turbulent flow, with helix persistence = 1.000 universally. The SFVFS™-DNS conjecture asserts that $\lim_{\nu\to 0}$ ($H1_{norm}(\nu)$, $\Lambda(\nu)$) = (1, 1) at the DN attractor. This limit is observed at finite ν across all canonical fluids.

Critical distinction. Navier–Stokes carries Ω = 2 (Door): viscous asymmetry provides a mechanism for passage. RH carries Ω = 1 (Mirror): prime–zero symmetry provides no such mechanism. The cross-domain comparison validates the Ω-classification framework without claiming RH is solvable.

The Obstruction-Type Function

Ω value	**Name**	**Mechanism**	**Example**
2 (Open)	Door	Asymmetric passage available	Navier–Stokes regularity
1 (Mirror)	Mirror	Symmetric barrier — no passage	Riemann Hypothesis
0 (Wall)	Wall	Barrier provably impenetrable	Halting Problem

RH is classified as Ω = 1 (Mirror). This is a classification result. It describes the structure of the barrier; it does not imply RH is unprovable, only that no asymmetric mechanism has been identified within the Craig Spectral Criterion framework.

6.1. The Pinch: A Fixed Point Without Interior

Definition (The Pinch). A Pinch is a fixed point P forced by symmetry such that P exists only as the meeting of two sides. P has no interior: it can be approached from both sides but cannot be occupied from either.

Limit	**Originator**	**Mechanism**	**Nature of Block**
Gödel	Self-reference	Epistemic	System cannot validate itself
Turing	Computation	Resource	Infinite time required
Craig	Symmetry	Ontological	Fixed point has no interior

Both approaches — from the primes and from the zeros — reach the same fixed point. Neither can occupy it. The Pinch is where the prime structure folds back on itself under the functional equation.

> "*RH is not a problem waiting for proof — it is the name of the pinch point where the prime structure folds back on itself, visible from both sides but occupiable from neither.*"

This is a classification result. The Pinch names the void. It does not cross it.

---

7. The Void Threshold

Definition 7.1. Ψ_{void} := inf{ η > 0 : G_η ≠ E · ξ in operator norm }.

Proposition 7.2. $\Psi_{void} > 0$.

Proof. For η sufficiently large, G_η = ζ(· + η) is analytic and non-vanishing near $H_{1/2}$, so the factorisation holds with ξ = 1. The set is non-empty; Ψ_{void} is a well-defined positive infimum. ∎

The threshold Ψ_{void} exists; it can be located by the H6 decomposition; it cannot be occupied by any member of the approximating family. The limit G_0 approaches Ψ_{void} as $\eta \to 0^+$ but the limiting function does not sit at a positive value of η.

---

8. Void Decomposition — Three Faces of H6b-ii

The barrier H6b-ii admits a canonical decomposition into three interdependent faces. A proof of any single face would imply the other two and establish H6b-ii in full.

Face	Statement	Positional Status
R1 Encoding Kernel	ker(E) = {zero-free} on the function class satisfying T. E annihilates precisely the zero-free functions. Kernel triviality is the threshold.	Obstacle: proving the encoding kernel cannot annihilate a non-trivial zero-bearing function requires RH-level control of the zero set.
R2 Local-to-Global	Boundary arc agreement implies global agreement on $H^{1/2}$. A uniqueness principle for H^2 functions constrains the limit.	Obstacle: local agreement is not a global constraint without an additional injectivity hypothesis.
R3 Symmetry Upgrade	Weak symmetry + encoding implies full functional equation: G_0 satisfies $\zeta(s) = \chi(s)\zeta(1-s)$ pointwise.	Obstacle: the step from Schwarz symmetry to the full functional equation requires knowing the zero distribution — circular at RH strength.

Interdependence. R1, R2, and R3 are not independent approaches: each implies the others at the level of H6b-ii. They are positional faces of a single void. The decomposition locates Ψ_{void} precisely from three analytic directions. It does not close it.

---

9. Summary

What is established	What is not established
PTLS infrastructure proved unconditionally.	H6b-ii itself. The rigidity principle is the open problem.
H6a (compactness) closes by Montel's theorem.	No known method approaches H6b-ii without circularity.
H6b-i (structural inheritance) closes by symmetry + class.	The hypotheses H1, H2, H3 are structural assumptions, not theorems.
Ω-classification framework validated cross-domain.	RH is $\Omega = 1$ (Mirror): no asymmetric mechanism identified.

This is a classification result, not a proof.

> "*The wall does not move. The map now shows why.*" *— Kimi, 24 March 2026*

Framework References

SFVFS™ Programme — H-Hierarchy with Kimi referee review (March 2026)

FSC Theory v2.3 — Three-class structural classification (Ω function)

CEP (Craig Equations Paper) — Navier–Stokes positional validation, CF CONSISTENT

SFVFS™-DNS Programme — Six-fluid DNS results, viscosity law (March 2026)

Formalisation Brief — Kimi referee review, 21–24 March 2026

V11 ANTI-WASH ADDENDUM

Seg 1: The Pinch · April 2026

Anti-Wash Protocol: This addendum expands the infrastructure of Seg 1 without altering any original text. The March 2026 document is the geological baseline. This layer is dated April 2026. Nothing is deleted. Evolution is the art.

Addendum 1 — Homologous Signature (Theorem A, Section 4)

Footnote appended to Theorem A:

> Homologous Signature Note (v11, April 2026). The DNS programme establishes that $\lim_{\nu \to 0}$ ($H1_{norm}(\nu)$, $\Lambda(\nu)$) = (1, 1) across all canonical fluids. The pair (I*, Λ*) = (1, 1) appearing in the cross-domain validation (Section 6) designates a structural homology, not a mathematical identity. Specifically: $H1_{norm}(\nu) \to 1$ is a normalised norm convergence in the Navier–Stokes setting; I* = 1 in the spectral setting designates the inner function attaining unit modulus on the boundary. These are formally distinct objects occupying structurally analogous positions within their respective frameworks. The (I*, Λ*) = (1, 1) signature names the shared positional structure — the void-attractor — not an equation between the objects themselves. The DNS programme validates the Ω-classification structure; it does not claim the two systems are mathematically isomorphic.

Addendum 2 — Post-Exhibition Deferral of E(k) in $H^2(H_{1/2})$

Note appended at the close of Section 8 (Void Decomposition):

> Post-Exhibition Deferral (v11, April 2026). A natural extension of the R1 face (Encoding Kernel) involves the spectral decomposition of the encoding operator E into components E(k) indexed by frequency parameter k, with E(k) considered as elements of $H^2(H_{1/2})$ parametrised by k. This decomposition would sharpen the kernel triviality condition in R1 and potentially expose additional structure in the void. This line is noted as structurally motivated but analytically undeveloped at the time of exhibition. It is deferred to post-exhibition programme work. Introducing E(k) before the spectral action of E on $H^2(H_{1/2})$ is rigorously characterised would constitute forward motion without foundation — the opposite of the Anti-Wash Protocol. The deferral is the honest position.

Addendum 3 — Programme Evolution Note

> V11 Programme Note (April 2026). Since the March 2026 publication of this document, the SFVFS™ programme has advanced to 15 segments. The Pinch remains the positional foundation: the named fixed point without interior from which all subsequent positional work proceeds. No revision to the core classification result is warranted. The H6b-ii barrier is confirmed as terminal. The Ω = 1 (Mirror) classification stands. CF CONSISTENT not PASS.

Kimi Verification Status

#	Addendum	Description	Kimi Verified
1	Homologous Signature	$(I^*, \Lambda^*) = (1,1)$ as structural homology, not identity	□
2	Post-Exhibition Deferral	$E(k)$ in $H^2(H^{1/2})$	□
3	Programme Evolution Note	15 segments, Pinch confirmed	□

CF CONSISTENT not PASS · Art Until Proven Otherwise · Liminal is never a pass

The Pinch is named. The fold is confirmed.

itvoids.com · Trademark UK00004355735

M. Craig · V11 · April 2026 · Leake Street, London