FLOW-STATIC COLLAPSE THEORY

A General Framework for Obstruction Classification

Marc Craig

February 2026

*"The value of a map is not that it crosses the mountains.
It is that it shows where the mountains are."*

CLASSIFICATION FRAMEWORK — NOT A PROOF METHOD

This document presents Flow-Static Collapse (FSC) Theory as a framework for classifying obstructions in limit-identification problems. The contribution is cartography, not conquest. FSC Theory does not prove hard theorems; it locates where the difficulty lives.

Abstract

We introduce **Flow-Static Collapse (FSC) Theory**, a framework for classifying obstructions in limit-identification problems. An FSC-scheme consists of a regularised family of analytic objects, an infrastructure package, an inheritance package, an encoding map, and a target class. The framework distinguishes **flow components** (which close under standard analysis) from **static barriers** (which do not).

The main theorems establish: (1) **Flow Saturation** — under sufficient infrastructure, extraction and inheritance close; (2) **Barrier Shape** — any terminal obstruction has the form "does encoding uniquely determine identification?"; (3) **Falsifiability** — rigidity fails if and only if two functions exist with identical encoding but different identification.

We present the Riemann Hypothesis as the base exemplar, where the terminal barrier reduces to "prime encoding determines zero locations" — which is RH-strength. The Split-to-Stable Test provides a termination criterion for obstruction decomposition. Named conjectures (R1, R2, R3) offer attackable targets for the static core.

Part I: Introduction

1.1 The Problem of Obstructions

Many problems in analysis take the form: a regularised family $\{G_\eta\}$ converges to a limit G_0 as $\eta \to 0$, and we wish to identify G_0 with a specific target object T. The difficulty typically divides into two parts:

> **(i) Flow:** Establishing convergence and showing the limit inherits properties from the approximants.
>
> **(ii) Static:** Proving the limit equals the target (identification/rigidity).

Flow components typically respond to analytical pressure — better bounds, tighter estimates, stronger compactness arguments. Static components do not. They represent irreducible content that no amount of analytical machinery can generate.

1.2 The Classification Goal

FSC Theory does not aim to solve hard problems. It aims to **classify** them: to systematically decompose a problem until all flow components close and only the static core remains. The contribution is **cartography**, not conquest. A successful FSC analysis produces:

- A clear inventory of what is proved (infrastructure)
- A decomposition of the limit hypothesis into flow and static parts
- Closure of all flow components
- Isolation of the terminal static barrier

- A stability test confirming the barrier is irreducible

1.3 What Kind of Equivalency Is This?

FSC Theory produces a **classification-equivalency**, not a reduction-equivalency. Compare:

Type	Example	Question Answered
Reduction	RH $\Leftrightarrow$ positivity of kernel	"Solve this instead"
Spectral	RH $\Leftrightarrow$ self-adjoint operator	"Find this object"
Statistical	RH $\Leftrightarrow$ GUE statistics	"Observe this behaviour"
FSC (this paper)	RH $\Leftrightarrow$ static rigidity	"What kind of obstruction is this?"

The FSC equivalency answers not "what else would imply RH?" but "**why does RH resist proof?**" — an orthogonal question.

Part II: Formal Definitions

2.1 FSC-Scheme

Definition. An **FSC-scheme** is a tuple (Ω, $\{G_\eta\}_{\eta>0}$, I, T, E, C) where:

- Ω is a domain (the ambient space)
- **$\{G_\eta\}$** is a family of holomorphic functions on Ω, parameterised by $\eta > 0$
- **I** is an **infrastructure package**: proved estimates (L^2 bounds, Carleson measures, BMOA bounds, etc.)
- **T** is an **inheritance package**: properties stable under locally uniform limits
- **E** is an **encoding map**: E: Hol(Ω) $\to$ D mapping functions to data
- **C** is a **target class**: the rigid class containing the desired identification

2.2 Flow and Static Components

Definition. A hypothesis H in an FSC-scheme is a **flow component** if it can be established using only the infrastructure I, standard limit theorems, and inheritance T.

Definition. A hypothesis H is a **static component** if no combination of infrastructure, limits, and inheritance suffices to establish it.

Remark. The flow/static distinction is relative to stated infrastructure. A component static today may become flow with new infrastructure. FSC Theory classifies relative to current tools.

2.3 The Canonical Decomposition

Every limit-identification problem admits a canonical decomposition:

FE (Flow-Extraction): $\{G_\eta\}$ is a normal family; extract convergent subsequence $G_{\eta_k} \to G_0$.
FI (Flow-Inheritance): G_0 inherits all properties in T.
SR (Static-Rigidity): $G_0 \in C$ (the limit is identified with the target).

FE and FI are flow. SR is static (in general).

2.4 Terminal Barrier

Definition. A **terminal barrier** is a static component that remains static under further decomposition. Formally: if SR decomposes into SR_1, SR_2, ..., and at least one is static, we recursively decompose until reaching components that do not further split (or whose splits relocate identical content).

Definition. A barrier is **stable** if further splitting relocates it without changing its logical content.

Part III: Main Theorems

3.1 Theorem (Flow Saturation)

Theorem 3.1 (Flow Saturation). Let (Ω, $\{G_\eta\}$, I, T, E, C) be an FSC-scheme. Suppose:

(i) I contains a **uniform bound**: $\exists M$ such that $\|G_\eta\|_{L^2(K)} \leq M$ for all $\eta > 0$ and all compact $K \subset \Omega$

(ii) Each property $P \in T$ is **limit-stable**: if $G_{\eta_k} \to G_0$ locally uniformly and each G_{η_k} satisfies P, then G_0 satisfies P

Then FE and FI close. Any remaining gap is static.

Proof. (i) Uniform L^2 bounds + Cauchy estimates $\to$ uniform L^∞ on compacta $\to$ Montel's theorem $\to$ subsequential convergence (FE closes). (ii) Limit-stability is the definition of T $\to$ inherited properties pass through (FI closes). What remains is SR. ■

3.2 Theorem (Barrier Shape)

Theorem 3.2 (Barrier Shape). In any FSC-scheme where flow saturates, the terminal barrier (if non-empty) has the form:

"Does $E(G_0) = E(\text{target})$ imply $G_0 \in C$?"

Proof. After flow saturation: FE gives existence of G_0. FI gives G_0 inherits T. The only remaining question is whether T + encoding data E suffices to force $G_0 \in C$. This is a uniqueness/rigidity question of the stated form. ■

3.3 Theorem (Falsifiability)

Theorem 3.3 (Falsifiability). The static rigidity SR fails if and only if there exist two distinct functions $F_1, F_2 \in \mathrm{Hol}(\Omega)$ such that:

(i) Both satisfy all inherited properties T

(ii) Both have identical encoding: $E(F_1) = E(F_2)$

(iii) $F_1 \in C$ but $F_2 \notin C$

Corollary. If no such pair exists, rigidity holds. The encoding uniquely determines membership in C (up to equivalence in C).

3.4 Sufficient Conditions for Closure

Theorem 3.4. SR closes if any of the following hold:

(i) Encoding injectivity: E is injective on the T-class (up to C-equivalence)

(ii) Rigidity from symmetry: The functional equation + growth bounds in T admit a unique solution

(iii) Local-to-global: Agreement on any open subset forces global equality

Part IV: The Split-to-Stable Test

4.1 Motivation

When a static component SR is identified, it may be productive to decompose it further. But decomposition can continue indefinitely, producing ever-finer hypotheses without progress. The **Split-to-Stable Test** provides a termination criterion.

4.2 The Protocol

Split-to-Stable Test: Starting at a static component SR, perform iterative binary splits:

> **Split 1:** SR → SR-a (structural/symmetry) + SR-b (rigidity/uniqueness)
>
> **Split 2:** If SR-b static, split: SR-b → SR-b-i (encoding) + SR-b-ii (identification)
>
> **Split 3:** If SR-b-ii static, split: SR-b-ii → (local uniqueness) + (global propagation)
>
> **Continue** until: (i) a component closes, or (ii) splits relocate identical content

Termination criterion: If after k splits the static component has the form "encoding determines identification" and further splits merely rephrase this, the barrier is **stable**.

Remark. The number of splits required is problem-dependent. The RH exemplar stabilises at depth 3. Other problems may stabilise earlier or later. The test is "split until stable," not "perform exactly n splits."

4.3 The x3 Split for RH (Detailed)

Applying the test to the Craig Spectral Criterion :

Split	Components	Status
1	SR-a (inheritance) / SR-b (rigidity)	Flow / Static
2	SR-b-i (encoding) / SR-b-ii (identification)	Flow / Static
3	(local uniqueness) / (global propagation)	Static / Flow
Terminal	**Encoding injectivity**	**BARRIER**

Terminal barrier: *"Do prime coefficients (von Mangoldt data) uniquely determine ξ on the inherited class?"* This is RH-strength — possibly equivalent to RH itself.

Why depth 3? At split 3, both branches ("local uniqueness" and "global propagation") reduce to the same question: whether knowing the encoding on part of the domain forces knowing the function everywhere. Further splits (e.g., "boundary uniqueness" vs "interior uniqueness") merely rephrase this. The content stabilises because there is only one hard question: *does the prime data determine the function?* — and that question is RH itself.

Part V: The Riemann Hypothesis as Base Exemplar

5.1 The FSC-Scheme for RH

The Craig Spectral Criterion instantiates an FSC-scheme:

- **Ω:** A horizontal strip S in the complex plane
- **$\{G_\eta\}$:** $G_\eta = Q - H_\eta$ where $H_\eta(z) = \Sigma\ \Lambda(n)\ n^{-1/2-\eta+2iz}$
- **I:** PTLS bounds, derivative PTLS, automatic uniformity ($\partial_\eta = (i/2)\partial_x$)
- **T:** Schwarz symmetry, BMOA class, Carleson bounds
- **E:** Von Mangoldt encoding (prime coefficients $\Lambda(n)$)
- **C:** $\{\Phi \cdot \xi$: E zero-free$\}$ where ξ is the completed zeta function

5.2 Status Summary

Component	Type	Status
Infrastructure (PTLS etc.)	Proved	✓ Complete
FE (compactness)	Flow	✓ Closes
FI (inheritance)	Flow	✓ Closes
SR (rigidity)	**Static**	**TERMINAL BARRIER**

5.3 The Terminal Obstruction Statement

All analytic/flow components close. The remaining obstruction is a rigidity principle: *"prime encoding determines zero locations."* This is the explicit formula of prime number theory turned into a uniqueness question. It cannot be closed by limit arguments, symmetry inheritance, or growth constraints alone.

The framework does not prove RH. It locates where RH lives.

Part VI: Named Conjectures for the Static Barrier

To enable focused attack, we decompose the terminal barrier into named conjectures:

Conjecture R1 (Encoding Kernel)

> Let E denote the prime encoding functional mapping entire functions to their von Mangoldt coefficients. On the class of functions satisfying T (inherited symmetry + growth), the kernel of E is trivial (up to zero-free factors).

In other words: two functions with the same prime encoding must be the same (mod zero-free).

Conjecture R2 (Local-to-Global)

> If two functions in the candidate class agree on their encoding over a test family (e.g., agreement of distributional identities on a boundary arc), then they agree globally (up to zero-free factor).

Conjecture R3 (Symmetry Upgrade)

> Schwarz/reflection symmetry + prime encoding forces the full functional equation symmetry of ξ.

In other words: you cannot have weak symmetry + the right encoding without getting strong symmetry.

Remark. R1, R2, R3 are interdependent. Proving any one likely implies progress on the others. They represent different angles of attack on the same static core. If any one is proved, SR closes and RH follows.

Part VII: Scope and Limitations

7.1 What FSC Theory Is

- A vocabulary for decomposing limit-identification problems
- A test for distinguishing flow from static components
- A termination criterion for obstruction decomposition
- A framework for generating named conjectures from terminal barriers

7.2 What FSC Theory Is Not

- A method for proving hard theorems
- A claim that static barriers can be overcome
- A reduction of hard problems to easy ones

7.3 The Classification Contribution

The value of FSC analysis is **classification**, not solution. Knowing that a problem's obstruction is a rigidity principle of a specific form is itself a contribution. It tells us:

- What kind of new ideas are needed (rigidity theorems, not better estimates)
- Why standard approaches fail (they saturate at flow)
- Where to focus effort (the named conjectures)
- How to recognise success (closing R1, R2, or R3)

7.4 Philosophical Position

FSC Theory answers not "what else would imply RH?" but "what kind of statement is RH?" This places it alongside:

- **Gödel:** Some truths are unprovable in F (Static: truth; Flow: proof)
- **Turing:** Halting is undecidable (Static: termination; Flow: computation)
- **FSC:** Some identifications require rigidity (Static: identification; Flow: analysis)

The pattern is general: systems have flow (what computation/proof/analysis can reach) and static cores (what requires principles beyond the system). FSC Theory makes this explicit for limit-identification problems.

Important disanalogy: Gödel and Turing proved *absolute* impossibility results within precisely defined formal systems. FSC barriers are *relative* to stated infrastructure — a component that is static today could become flow with new tools. This is a feature, not a bug: it means FSC classifies difficulty relative to current methods, not for all time.

Part VIII: Potential Further Exemplars

FSC Theory may apply to other limit-identification problems. Candidates include:

- **Generalised Riemann Hypotheses:** L-functions for Dirichlet characters, modular forms. The scheme structure is analogous; the encoding changes.
- **Invariant subspace problems:** Where a limit must be identified with a specific subspace. Flow = compactness; Static = rigidity.
- **Spectral inverse problems:** Where eigenvalue limits must be identified. Encoding = spectral data; Barrier = spectral rigidity.

If FSC survives even one non-RH exemplar with the same barrier shape, it becomes an independent theory. We leave development of these exemplars to future work.

Publication Statement

Document: Flow-Static Collapse Theory — Complete Framework

Author: Marc Craig

Date: 8 February 2026

Timestamp: 2026-02-08 00:33:08 UTC

Location: Leake Street, London

Base Exemplar: Craig Spectral Criterion (Riemann Hypothesis)

This is art until someone says otherwise.

Developed through collaboration between human intuition and AI systems (Claude, ChatGPT, Gemini, Kimi) over 13+ weeks. The framework emerged from systematic decomposition of the Riemann Hypothesis obstruction, yielding a general theory of flow-static collapse applicable beyond its origin.

— End of Document —

"All flow collapses. The static core remains. That's where the problem lives."