Conditional No-Upgrade for Explicit-Formula Positivity

A precise sufficient condition for the unreachability of the Riemann Hypothesis by limit-based methods

---

At a Glance

What this paper is. A conditional proof. It establishes that *if* a certain kind of mathematical workspace exists with three specific properties, *then* one well-known strategy for proving the Riemann Hypothesis cannot succeed within it.

What this paper shows. The most natural candidate for such a workspace does not have the required properties. We work out exactly why.

What this paper leaves open. Whether any workspace at all has the required properties. This is a single, precise question that captures the remaining analytic difficulty.

What this paper does not claim. It does not prove the Riemann Hypothesis. It does not disprove it. It does not foreclose all approaches — only one specific limit-based strategy, and only when the activating conditions are met.

Why it matters. The conditional result clarifies, in mathematically precise terms, why a particular natural approach to a 160-year-old problem has resisted completion — and reduces a vague intuition to a sharp open problem.

Summary for General Readers

The Riemann Hypothesis is one of the most famous unsolved problems in mathematics. It has resisted proof for over 160 years, and a million-dollar prize is offered for solving it. This paper does not solve it. What it does is explain, with precision, why one natural way of trying to solve it cannot work — and what would have to be true for that approach to have any chance.

To understand the idea, imagine you are trying to take a perfectly sharp photograph of something. You start with a slightly blurry picture, and you keep sharpening it. Each sharpened version is clearer than the last. The natural assumption is that if you keep going, you will eventually reach the perfectly sharp original.

This paper, together with its two companions, is about a situation where that assumption is wrong. There are mathematical objects where every "slightly blurry" version is well-behaved and easy to handle, but the "perfectly sharp" version is something completely different — not just sharper, but a different *kind* of object altogether. No matter how much you sharpen the blurry versions, you never reach the sharp one. The two live in different worlds.

The Riemann Hypothesis, when approached through one particular strategy, has exactly this structure. There is a mathematical inequality which, if true, would prove the Riemann Hypothesis. This inequality comes in two flavours: a "blurry" version (called the finite-resolution version) and a "sharp" version (the infinite-resolution version). Many mathematicians have hoped to prove the blurry version first and then take the sharp limit.

A "no-upgrade theorem" is a result that says: this kind of strategy cannot work. More precisely, a no-upgrade theorem identifies a situation where you can prove something at every blurry resolution — making the picture as sharp as you like, short of perfectly sharp — and yet none of those proofs combine to give you a proof at perfect sharpness. The reason is not that mathematicians have failed to find the right argument. It is that no such argument can exist within the chosen mathematical framework, because the perfectly sharp object is structurally different from all the blurry approximations. "Upgrading" from blurry truths to a sharp truth is forbidden by the structure itself, not by any lack of cleverness. This is what makes a no-upgrade theorem useful: it tells you to stop trying one kind of approach and look elsewhere.

The first paper in the trilogy shows that the blurry version can indeed be analysed, and identifies what would be needed to prove it. The second paper proves a general no-upgrade theorem: blurry versions cannot become sharp versions by any limiting process, when the two live in different mathematical worlds. This paper — the third — connects the two by asking: does the gap between blurry and sharp actually open up in the Riemann setting?

The answer turns out to be: yes, but only if you choose your mathematical workspace carefully. Specifically, there is a precise list of requirements (which we call hypotheses (a), (b), and (c)) that a workspace must meet for this proof strategy to even have a chance of working. We prove that *if* a workspace meeting these requirements exists, *then* the strategy is doomed by the no-upgrade theorem of the second paper.

This is what mathematicians call a *conditional proof*. A conditional proof is a real proof — every step is rigorous, every deduction is airtight — but the conclusion is reached only when a stated assumption is granted. Think of it as a sealed contract: if the condition holds, the conclusion follows with full mathematical force. If the condition fails, the contract is void and the conclusion is not delivered. Conditional proofs are not weaker than ordinary proofs in their internal logic; they are different in their scope. Many of the most important results in modern mathematics are conditional in this sense — true *given* a hypothesis whose own truth remains to be settled.

The condition in this paper is not arbitrary. It expresses a precise mathematical statement about *finite distinguishability* — the idea that any well-defined mathematical observation has a smallest scale below which structure cannot be told apart. At every finite level of distinguishability, the prime numbers can be sampled cleanly, and the inequality behind the Riemann Hypothesis behaves well. The hypotheses (a), (b), and (c) are the technical conditions that make this finite-distinguishability picture mathematically coherent: they specify what kind of workspace allows the regularised (finite-distinguishability) version of the problem to be well-posed. The infinite-distinguishability case — perfectly sharp resolution — falls outside any such workspace by its very nature, which is exactly why the no-upgrade theorem applies. The conditional proof, in other words, is the precise mathematical expression of the intuition that finite-distinguishability descriptions cannot extend to the infinite-distinguishability limit by ordinary refinement.

A musical analogy may help — though with one twist that needs flagging up front. In ordinary music, notes feel further apart the higher you go: each octave is twice the previous frequency, so the upper register feels stretched out. The "zeros" of the Riemann zeta function — the mysterious numbers whose location is what the Riemann Hypothesis is about — behave in exactly the opposite way. They can be thought of as a kind of infinite sequence of notes, each at a slightly different pitch, but the higher you climb up this scale, the *closer together* the notes become. The lowest zeros are well-spaced and easy to distinguish from one another. By the time you reach the millionth zero, the gap to its neighbour has shrunk by more than half compared to the first few. By the billionth, it has shrunk further still. The crowding is slow but unstoppable — and at some point the pitches become so densely packed that any real instrument, any physical or mathematical device with finite resolution, cannot tell one note from the next. The notes are still there, mathematically, but they have crossed below the threshold of distinguishability for the listener. Finite distinguishability says: you can only ever hear the notes that are far enough apart to be told apart. The Riemann Hypothesis, in its full form, is a statement about *all* the notes — including the ones too closely spaced to ever be distinguished by any finite-resolution instrument. Trying to prove the full hypothesis by reasoning only about distinguishable notes is like trying to describe a chord by listing only the pitches your ear can separate: you may capture the lower notes perfectly, but the upper register, where the notes blur together, is structurally inaccessible to that method. This is the gap the no-upgrade theorem identifies, and the conditional proof shows how that gap propagates from the musical intuition into rigorous mathematics.

But here is the twist: the most obvious candidate workspace — the one mathematicians would naturally reach for — does not meet the requirements. We work out the specific reason why. It comes down to a tension between two competing demands: you need the workspace to be big enough to capture what the Riemann Hypothesis actually says, but small enough to keep the mathematical objects under control. These two demands pull in opposite directions, and for the natural choice they cancel each other out exactly.

This leaves a clear open question: does any workspace exist that satisfies all the requirements simultaneously? If yes, this proof strategy for the Riemann Hypothesis is closed off, and anyone trying to prove it this way is wasting their time. If no — if the requirements are fundamentally incompatible — then that incompatibility is itself a kind of structural answer, explaining why this approach to the Riemann Hypothesis has resisted analysis for so long.

Either way, the result is informative. We have not solved the Riemann Hypothesis, and we have not proved that it cannot be solved. What we have done is identify precisely where the difficulty lies in one specific approach, and we have reduced a vague intuition ("this strategy seems to fail") to a precise mathematical question that can be attacked directly.

The broader lesson goes well beyond the Riemann Hypothesis. The same pattern — blurry versions behaving well but never converging to a sharp version — appears in many areas of mathematics and physics. It explains why certain calculations in quantum field theory require a procedure called renormalisation, why signal processing has to treat continuous and discrete information differently, and why distributions (mathematical objects like the "Dirac delta") are needed alongside ordinary functions. In each case, the lesson is the same: refinement is not always enough. Sometimes the sharp object lives in a different mathematical universe, and reaching it requires genuinely new ideas, not just better approximations.

What this paper is: a precise statement of one route to the Riemann Hypothesis, a specific obstruction showing why the natural attempt fails, and a sharp open question that captures what is left.

What this paper is not: a proof of the Riemann Hypothesis, a disproof of the Riemann Hypothesis, or a claim that the Riemann Hypothesis cannot be proved by some other method.

The deepest reading of the result is this:

The obstruction to the Riemann Hypothesis via explicit-formula positivity may not arise from insufficient analysis, but from attempting to recover an infinite-distinguishability object from a finite-distinguishability observable framework.

If that reading is correct, the difficulty of the Riemann Hypothesis along this route is not a missing technique. It is a structural mismatch between what the hypothesis is *about* and what finite-distinguishability mathematics can *see*. Whether that mismatch can be overcome — by working directly with infinite-distinguishability objects, by changing the framework, or by some entirely new approach — is the question this trilogy hands forward.

Technical Abstract

Let $Q(h) = \mathcal{A}(h * h^{\vee}) - \mathcal{P}(h)$ denote the explicit-formula quadratic form for even Schwartz h, where $\mathcal{A}$ is the archimedean contribution and $\mathcal{P}$ is the prime-power sampling functional. By the de Branges criterion, the Riemann Hypothesis is equivalent to $Q(h) \geq 0$ for all such h. The arithmetic functional $\mathcal{P}$ admits a natural regularisation $\mathcal{P}_\Delta$ at resolution $\Delta > 0$ via mollification of the prime-power point measure.

We prove a conditional no-upgrade theorem: if there exists a Banach space X containing the relevant test functions such that

(a) $\mathcal{P}_\Delta$ is bounded on X uniformly in the bandwidth parameter A, for each fixed $\Delta > 0$, (b) the archimedean form $\mathcal{A}$ is coercive on X modulo a fixed finite-dimensional baseline subspace, and (c) the coercivity and boundedness constants admit a uniform gap,

then for every $\Delta > 0$ the regularised positivity $Q_\infty \geq \mathcal{P}_\Delta$ holds on the appropriate subspace, and by the general no-upgrade theorem of the companion paper no admissible limit $\Delta \to 0$ within X recovers the atomic statement equivalent to RH.

We then prove that the natural candidate space — the weighted Sobolev space H^1_ω with $\omega(\xi) = e^{-|\xi|/2}$ — fails hypothesis (a). The explicit counterexample is the family f_A constructed to demonstrate the unboundedness of the atomic functional $\mathcal{P}_0$; the same construction shows that the smoothed functionals $\mathcal{P}_\Delta$ are also unbounded on H^1_ω uniformly in A. Hence the natural candidate is excluded.

We discuss alternative candidate spaces, identify the structural tension between hypotheses (a) and the RH-equivalence requirement, and state the resulting open problem precisely.

---

1. Introduction

1.1 The trilogy

This paper completes a three-part programme:

— **Paper I** (*A Spectral Framework for the Riemann Hypothesis: The VERSF Approach*) develops the de Branges / Weil explicit-formula route to RH, identifies the dilation generator as the correct spectral primitive, and analyses the regularised positivity $Q_\infty \geq \mathcal{P}_\Delta$ at each finite resolution $\Delta > 0$.

— **Paper II** (*Finite Resolution Does Not Converge to Infinite Resolution*) proves a general functional-analytic theorem: if $\{S_\Delta\}_{\Delta>0}$ is a family of bounded quadratic forms on a Banach space X and S_0 is unbounded on a dense subspace, then no admissible limit $S_\Delta \to S_0$ exists in any topology compatible with the norm structure of X.

— **Paper III** (this paper) bridges the two: it formulates the precise conditional under which Paper II's general theorem activates against the RH setup of Paper I, proves that the natural candidate space fails the activating hypotheses, and states the genuine open problem.

1.2 What is being claimed

The conditional theorem proved here is of the form:

If hypotheses (a)–(c) hold for some Banach space X large enough to make Q-positivity on X equivalent to RH, *then* RH cannot be obtained as a limit of finite-Δ positivity results within X.

This is a meta-theorem about proof strategies, not about RH itself. It rules out one specific approach — *regularise at finite Δ, prove positivity, take Δ → 0 within a fixed space* — provided one can identify an X for which the regularisation is meaningful.

1.3 What is not being claimed

This paper does not prove that hypotheses (a)–(c) hold for any specific space. The most natural candidate, H^1_ω, demonstrably fails (a). Whether any other candidate succeeds is open.

In particular, this paper does *not* claim:

— that RH has been proven or disproven; — that all spectral or explicit-formula approaches to RH are closed off; — that the regularise-and-limit strategy cannot work in some entirely different framework not considered here.

The contribution is to make the conditional precise, identify one concrete obstruction, and isolate the remaining question.

1.4 Logical status

The structure of the trilogy as a whole is:

Paper I: regularised positivity at each $\Delta > 0$ (conditional, modulo the same hypothesis identified here) Paper II: general no-upgrade theorem (unconditional) Paper III: conditional bridge, plus obstruction for the natural candidate (this paper)

Taken together, the trilogy provides:

— an unconditional general theorem (Paper II); — a precise conditional combining it with the RH explicit-formula route (Paper III, §3); — a precise obstruction showing the natural candidate fails (Paper III, §4); — a sharp open problem (Paper III, §5).

It does *not* provide an unconditional proof of regularised positivity in any specific space — that question is intertwined with the conditional, and the failure of the natural candidate is itself part of the result.

---

2. Setup

2.1 The explicit-formula quadratic form

Let h be an even Schwartz function on $\mathbb{R}$, set $h^\vee(t) = h(-t)$, and write $g = h * h^\vee$, so that $\hat{g}(\xi) = |\hat{h}(\xi)|^2 \geq 0$.

The Weil explicit formula, specialised to test functions of the form $g = h * h^\vee$, gives the identity

$$\Sigma_\gamma g(\gamma) = \mathcal{A}(g) - \mathcal{P}(h)$$

where the sum runs over ordinates γ of the nontrivial zeros of $\zeta(s)$, the archimedean contribution is

$$\mathcal{A}(g) = (1/2\pi) \int_{\mathbb{R}} g(t) \operatorname{Re} \psi(1/4 + it/2)\, dt + g(i/2) + g(-i/2) - (\log \pi / 2\pi) \int_{\mathbb{R}} g(t)\, dt$$

with $\psi = \Gamma'/\Gamma$ the digamma function, and the prime-power sampling functional is

$$\mathcal{P}(h) = (1/2\pi) \Sigma_p \Sigma_{k\geq 1} (\log p)/p^{k/2} \cdot |\hat{h}(k \log p)|^2.$$

Define the explicit-formula quadratic form

$$Q(h) := \mathcal{A}(h * h^\vee) - \mathcal{P}(h).$$

2.2 The de Branges criterion

Let $\Xi(t) := \xi(1/2 + it)$ where $\xi(s) = 1/2 s(s-1)\pi^{-s/2}\Gamma(s/2)\zeta(s)$ is the completed zeta function. Set $E(z) := \Xi(z) - i\Xi'(z)$.

Theorem 2.1 (de Branges; classical). The following are equivalent:

(i) All nontrivial zeros of ζ lie on the critical line $\operatorname{Re}(s) = 1/2$ (the Riemann Hypothesis). (ii) All zeros of Ξ are real. (iii) E is a Hermite–Biehler function: $|E(z)| > |E(\bar{z})|$ for $\operatorname{Im}(z) > 0$. (iv) $Q(h) \geq 0$ for every even Schwartz h.

Equivalence of (iv) with the others is the *positivity criterion* on which the explicit-formula route to RH is based.

2.3 Regularisation

The prime-power functional $\mathcal{P}$ corresponds in the frequency variable $\xi = k \log p$ to integration against the atomic measure

$$\mu := \Sigma_p \Sigma_{k\geq 1} (\log p)/p^{k/2} \cdot \delta_{k \log p}.$$

For each resolution $\Delta > 0$, fix a nonnegative mollifier $\eta \in C_c^\infty([-1, 1])$ with $\int \eta = 1$, and set $\eta_\Delta(x) := \Delta^{-1}\eta(x/\Delta)$. Define the smoothed measure

$$\mu_\Delta := \mu * \eta_\Delta$$

and the associated *Δ-regularised sampling functional*

$$\mathcal{P}_\Delta(h) := (1/2\pi) \int_{\mathbb{R}} |\hat{h}(\xi)|^2\, d\mu_\Delta(\xi).$$

As $\Delta \to 0$, $\mu_\Delta \to \mu$ in the distributional sense, and $\mathcal{P}_\Delta(h) \to \mathcal{P}(h)$ for h with sufficiently regular Fourier transform.

Define the *Δ-regularised quadratic form*

$$Q_\Delta(h) := \mathcal{A}(h * h^\vee) - \mathcal{P}_\Delta(h).$$

Pointwise on test functions, $Q_\Delta \to Q$ as $\Delta \to 0$.

2.4 The bandwidth parameter

For computational tractability, one often restricts to test functions h with $\mathrm{supp}(\hat{h}) \subset (-A, A)$ for some bandwidth $A > 0$. The sampling sum then truncates to prime powers with $k \log p \le A$. Throughout, we write $\mathcal{P}_{\Delta,A}$ and $Q_{\Delta,A}$ for the bandwidth-restricted versions.

The full statement of RH-equivalence requires that the limit $A \to \infty$ be taken; conversely, any uniform-in-A bound carries information about all of frequency space.

3. The Conditional No-Upgrade Theorem

3.1 Hypotheses

Let X be a Banach space whose elements (or whose Fourier transforms) include the even Schwartz functions on $\mathbb{R}$ — possibly after appropriate restriction.

Hypothesis (a) — Uniform boundedness of regularised sampling. For each fixed $\Delta > 0$, there exists $C_\Delta < \infty$ such that for every bandwidth A and every test function h with $f = \hat{h}$ supported in $(-A, A)$ and $f \in X$,

$$\mathcal{P}_{\Delta,A}(h) \le C_\Delta \|f\|_X^2. \text{ (a)}$$

The constant C_Δ may depend on Δ but must be independent of A.

Hypothesis (b) — Archimedean coercivity modulo finite codimension. There exists a fixed finite-dimensional subspace $B \subset X$ and a constant $c_\infty > 0$ such that for every bandwidth A and every test function h with $f = \hat{h}$ supported in $(-A, A)$, $f \in X$, and $f \perp B$,

$$\mathcal{A}(h * h^\vee) \ge c_\infty \|f\|_X^2. \text{ (b)}$$

Both B and c_∞ are independent of A.

Hypothesis (c) — Constants gap. The constants from (a) and (b) admit a strict ordering: there exists a choice of mollifier η and of c_∞ such that

$c_\infty > C_\Delta$. (c)

3.2 The conditional theorem

Theorem 3.1 (Conditional No-Upgrade). Suppose hypotheses (a), (b), (c) hold for some Banach space X that is *compatible with the explicit-formula framework* — meaning that X is a weighted Sobolev-type admissibility class in which the f_A construction of §4 yields elements of X with uniformly bounded X-norm. Then:

(i) *Finite-resolution positivity:* For every $\Delta > 0$ satisfying (c), every bandwidth A, and every test function h with $f = \hat{h} \in X$, $\mathrm{supp}(f) \subset (-A, A)$, and $f \perp B$,

$$Q_{\Delta,A}(h) \geq 0.$$

(ii) *No upgrade:* The atomic functional $\mathcal{P} = \mathcal{P}_0$ is unbounded on X. By Theorem 3.1 of Paper II, no admissible convergence $\mathcal{P}_\Delta \to \mathcal{P}_0$ exists in any topology on X compatible with its norm.

(iii) *Conditional obstruction:* If, in addition, the de Branges criterion (Theorem 2.1.iv) restricted to the class $\{h : \hat{h} \in X, \hat{h} \perp B\}$ is equivalent to the full criterion on all even Schwartz functions — equivalently, if X is "large enough" for the de Branges route — then RH cannot be obtained by admissible $\Delta \to 0$ limit procedures within X.

3.3 Proof

Part (i). By (b), $\mathcal{A}(h * h^\vee) \geq c_\infty \|f\|_X^2$. By (a), $\mathcal{P}_{\Delta,A}(h) \leq C_\Delta \|f\|_X^2$. Subtracting and using (c),

$$Q_{\Delta,A}(h) = \mathcal{A}(h * h^\vee) - \mathcal{P}_{\Delta,A}(h) \geq (c_\infty - C_\Delta) \|f\|_X^2 \geq 0.$$

Part (ii). We must show $\mathcal{P}_0$ is unbounded on X under the additional structural assumption that X is a weighted Sobolev-type admissibility class compatible with the explicit-formula framework — that is, a space whose norm allows the f_A construction of §4 (test functions constant on long intervals with uniformly bounded norm) to lie in X.

Under this structural assumption, the f_A family of Lemma 4.1 gives a sequence in X with bounded X-norm whose atomic sampling values $\mathcal{P}_0(h_{A_n})$ diverge. Hence $\mathcal{P}_0$ is unbounded on X.

The hypotheses of Theorem 3.1 of Paper II then apply with $S_\Delta := \mathcal{P}_\Delta$ (bounded on X by (a)) and $S_0 := \mathcal{P}_0$ (unbounded on a dense subspace of X). By that theorem, no admissible convergence exists.

Remark: The structural assumption on X is satisfied by all candidate spaces considered in §§4–5, and more generally by any space adapted to the explicit-formula structure in which test functions of the f_A form are admissible. If X is constructed to *exclude* such test functions (for instance, by an oscillation constraint forcing f to vary at every scale), then Part (ii) requires a separate argument, and clause (iii) of the theorem may fail because X may not satisfy hypothesis (d) of Problem 6.1. See §5.3 for discussion.

Part (iii). If the de Branges criterion on the test class associated to X is equivalent to RH, then RH is equivalent to the $\Delta = 0$ statement $Q_0(h) \geq 0$ for all h in that class. By (ii), this $\Delta = 0$ statement cannot be obtained as an admissible limit of the $\Delta > 0$ statements proved in (i). Hence no admissible $\Delta \to 0$ limit procedure within X carries finite-Δ positivity to RH.

This completes the proof of Theorem 3.1.

3.4 What the theorem does and does not say

Theorem 3.1 does *not* prove that RH is unreachable. It proves that one specific proof strategy — *establish positivity at each $\Delta > 0$ in a fixed space X, then take $\Delta \to 0$ within X* — cannot succeed when X satisfies the hypotheses.

There are at least three escape routes that the theorem does not foreclose:

(1) *Direct attack at $\Delta = 0$.* Prove $Q_0(h) \geq 0$ by methods that do not rely on finite-Δ approximation — for instance, by working in a distributional setting on a different space $\tilde{X} \supset X$ where $\mathcal{P}_0$ *is* bounded.

(2) *Renormalisation.* Modify the regularisation scheme itself so that the renormalised limit $\lim_{\Delta \to 0} (\mathcal{P}_\Delta - \text{counterterms})$ is bounded on X.

(3) *Alternative characterisation.* Find a reformulation of RH that does not pass through $\Delta = 0$ at all.

Each of these constitutes a genuine mathematical move, not a refinement of the limit-based strategy.

3.5 The deeper question

Theorem 3.1 is genuinely useful only if some Banach space X satisfies all three hypotheses *and* is large enough for clause (iii) — for the de Branges criterion on X to be equivalent to RH.

The remainder of this paper investigates this question. The answer for the most natural candidate is negative, and the structural reason is informative.

---

4. The Natural Candidate Fails

4.1 The weighted Sobolev space H^1_ω

The prime-power weights $(\log p)/p^{k/2} = (\log p) \cdot e^{-(k \log p)/2}$ suggest the weight

$$\omega(\xi) := e^{-|\xi|/2}$$

as the natural way to absorb the arithmetic decay into a function-space norm. Define

$$\|f\|_{H^1_\omega(0,A)}^2 := \int_0^A \left(|f(\xi)|^2 + |f'(\xi)|^2\right) \omega(\xi)\, d\xi$$

and let $H^1_\omega(0, A)$ be the completion of $C^\infty_c([0, A])$ in this norm.

This is the candidate space adopted throughout Paper I. The motivation is sound: the weight ω matches the exponential decay of the arithmetic weights, and the H^1 regularity is the minimal Sobolev requirement for point evaluation to make sense.

4.2 The f_A counterexample

Lemma 4.1 (Bounded H^1_ω norm with unbounded sampling). Fix $\xi_0 > 0$. For each $A > \xi_0 + 2$, let $f_A \in C^\infty([0, A])$ satisfy:

— $f_A(\xi) = 1$ for $\xi \in [\xi_0, A - 1]$; — $f_A(\xi) = 0$ for $\xi \in [0, \xi_0 - 1] \cup [A, \infty)$; — f_A interpolates smoothly on the two taper regions $[\xi_0 - 1, \xi_0]$ and $[A - 1, A]$, with derivatives bounded by an absolute constant.

Then:

(i) *H^1_ω norm uniformly bounded:*

$$\sup_A \|f_A\|_{H^1_\omega(0,A)}^2 \le 2\, e^{-(\xi_0-1)/2} + C_0$$

for an absolute constant C_0 depending only on the taper.

(ii) *Atomic sampling diverges:*

$$\begin{aligned}\mathcal{P}_{0,A}(h_A) &= (1/2\pi) \sum_{p^k \le e^A} (\log p)/p^{k/2} \cdot |f_A(k \log p)|^2 \\ &\ge (1/2\pi) \sum_{e^{\xi_0} \le n \le e^{A-1}} \Lambda(n)/\sqrt{n} \\ &\sim (1/\pi)\, e^{(A-1)/2}\end{aligned}$$

as $A \to \infty$, where the asymptotic follows from partial summation against $\psi(x) \sim x$ (PNT).

(iii) *Smoothed sampling also diverges:*

$$\begin{aligned}\mathcal{P}_{\Delta,A}(h_A) &= (1/2\pi) \int_0^A |f_A(\xi)|^2\, d\mu_\Delta(\xi) \\ &\ge (1/2\pi)\, \mu_\Delta([\xi_0, A - 2]) \\ &\sim (1/\pi)\, e^{(A-2)/2}\end{aligned}$$

as $A \to \infty$, for any fixed $\Delta \in (0, 1]$.

(Here h_A denotes the test function whose Fourier transform is f_A.)

Proof.

(i) Direct estimate:

$\int_0^A |f_A(\xi)|^2 e^{-\xi/2}\, d\xi = \int_{\xi_0}^{A-1} e^{-\xi/2}\, d\xi + \text{(taper contributions)} \le 2(e^{-\xi_0/2} - e^{-(A-1)/2}) + C_0/2 \le 2 e^{-\xi_0/2} + C_0/2$.

The derivative is supported on the two taper regions of width 1 each, with $|f'_A| \le \text{const}$:

$\int_0^A |f'_A(\xi)|^2 e^{-\xi/2}\, d\xi \le \text{const} \cdot (e^{-(\xi_0-1)/2} + e^{-(A-1)/2}) \le C_0/2$.

Summing gives the uniform bound.

(ii) On the bulk region $[\xi_0, A - 1]$ we have $f_A \equiv 1$, so

$\mathcal{P}_{0,A}(h_A) \ge (1/2\pi) \Sigma_{k \log p \in [\xi_0, A-1]} (\log p)/p^{k/2} = (1/2\pi) \Sigma_{e^{\xi_0} \le p^k \le e^{A-1}} \Lambda(p^k)/\sqrt{p^k}$.

By PNT, $\Sigma_{n \le x} \Lambda(n)/\sqrt{n} \sim 2\sqrt{x}$ (partial summation from $\psi(x) \sim x$). Hence

$\Sigma_{e^{\xi_0} \le n \le e^{A-1}} \Lambda(n)/\sqrt{n} \sim 2 e^{(A-1)/2} - 2 e^{\xi_0/2}$.

(iii) The smoothed measure μ_Δ has total mass on $[\xi_0, A-2]$ satisfying

$\mu_\Delta([\xi_0, A - 2]) \ge \mu([\xi_0 + \Delta, A - 2 - \Delta]) \sim 2 e^{(A-2-\Delta)/2}$

by the same PNT estimate. Since $|f_A| = 1$ on this interval, the bound follows.

4.3 Consequence for hypothesis (a)

Corollary 4.2. The space $X = H^1\_\omega(0, A)$ does not satisfy hypothesis (a).

Proof. From Lemma 4.1, for every fixed $\Delta \in (0, 1]$,

$\mathcal{P}_{\Delta,A}(h_A) / \|f_A\|_{H^1\_\omega(0,A)}^2 \ge (1/\pi) e^{(A-2)/2} / (2 e^{-(\xi_0-1)/2} + C_0) \to \infty$

as $A \to \infty$. So no constant C_Δ independent of A can satisfy (a).

4.4 Why the natural candidate fails: a structural diagnosis

The failure is not an artefact of the construction. It reflects a structural tension between the two demands placed on X:

— To make the limit $\Delta \to 0$ *meaningful* (so that the de Branges criterion on X is RH-equivalent), X must be large enough to contain test functions whose Fourier transforms are constant on long intervals. Such test functions are necessary for distinguishing positivity violations at small frequencies from violations at large frequencies.

— To make the regularised sampling *uniformly bounded* (hypothesis (a)), the norm on X must penalise test functions that are large at many prime-power frequencies. Functions constant on long intervals are large at *all* prime-power frequencies in those intervals — and there are exponentially many of them, by PNT.

The weight $\omega(\xi) = e^{-|\xi|/2}$ resolves this tension in the wrong direction for hypothesis (a): it allows constant functions to have finite norm (because $\int e^{-|\xi|/2}\, d\xi < \infty$), but the arithmetic density of prime powers in $[\xi_0, A]$ grows as $e^{A/2}$, which exactly matches the decay rate of the weight. The two cancel, and the sampling is *not* bounded.

Any heavier weight would suppress the sampling — but it would also exclude constant test functions, and with them the test-function class needed for de Branges equivalence.

This is the deep tension: the very feature of H^1_ω that makes the atomic functional unbounded (allowing constant functions) is what allows the regularised functionals to be unbounded too. The space cannot separate them.

---

5. Alternative Candidates

5.1 Restricted bandwidth classes

One could restrict to test functions with $\operatorname{supp}(\hat{h}) \subset (-A_0, A_0)$ for some fixed *finite* A_0. On such a class, both $\mathcal{P}_{\Delta,A_0}$ and $\mathcal{P}_{0,A_0}$ are bounded (they are finite-rank operators). Hypothesis (a) becomes trivial.

However, the de Branges criterion on a fixed-bandwidth class is *not* RH-equivalent: positivity on bandlimited h does not control Q on Schwartz h whose Fourier transforms have unbounded support. So clause (iii) of Theorem 3.1 fails. The conditional theorem activates, but trivially — the "RH" it forecloses is not the actual RH.

5.2 Heavier-weight Sobolev spaces

Take a heavier weight $\tilde{\omega}(\xi) = e^{-\alpha|\xi|}$ with $\alpha > 1/2$. Then constant functions on $[\xi_0, A]$ still have bounded $\tilde{\omega}$-norm. The sampling functional's growth rate is exponential, governed by the prime-power density, which is controlled by PNT to be $\sim e^{\xi/2}$. The mismatch becomes

$\mathcal{P}_{\Delta,A}(h_A) \sim e^{A/2}$, $\|f_A\|_{\tilde{\omega}}^2 \sim$ bounded, so hypothesis (a) still fails. Heavier weights do not help because they do not change the asymptotic — they only shift the constants.

Lighter weights ($\alpha < 1/2$) cause the constant function to have infinite norm, which excludes the f_A family from X — but then the resulting class is too small for de Branges equivalence, by an argument analogous to §5.1.

The exact match $\alpha = 1/2$ (the H^1_ω of §4) is special: it is the unique weight at which the f_A family lies on the boundary of integrability, and at which the prime-power density and weight decay are exactly equal. The failure of (a) at this weight is therefore unavoidable.

5.3 Spaces with oscillation constraints

One could try to define X by imposing constraints beyond H^s smoothness — for instance, a "no-flatness" condition forbidding f from being constant on long intervals. Paper I's TPB-Adm condition is one such attempt.

The technical obstacle: any such condition is either (i) not preserved under the limit $\Delta \to 0$ (so the limit object leaves X), or (ii) sufficiently restrictive that the test class is no longer large enough for de Branges equivalence.

The TPB-Adm condition specifically is a weighted Poincaré-Hardy inequality. On the class of test functions satisfying it, $\mathcal{P}_\Delta$ becomes bounded — but at the cost of excluding precisely the f_A-type families that are needed to detect positivity violations at all frequencies. We do not have a proof that TPB-admissible positivity is RH-equivalent, and we suspect it is not.

5.4 Distributional extensions

Embedding X into a distributional space $\tilde{X} \supset X$ on which $\mathcal{P}_0$ is bounded would constitute escape route (1) of §3.4 — a direct attack at $\Delta = 0$ rather than a limit-based argument. The natural candidate is H^{-s} for $s > 1/2$, in which the atomic measure μ is a bounded distribution. But the de Branges criterion in H^{-s} requires its own analysis, and the relationship to $Q(h) \geq 0$ in the original variable is not standard. This is open territory.

5.5 The space of strict admissibility

It is possible that no Banach space simultaneously satisfies hypotheses (a)–(c) and is large enough for clause (iii) of Theorem 3.1. If this is provably the case, the conditional theorem becomes vacuous in its intended application, and the result must be reinterpreted.

We do not know whether such a space exists.

---

6. The Open Problem

6.1 Precise statement

Problem 6.1 (Existence of an Activating Space). Does there exist a Banach space $X \subset \mathcal{S}'(\mathbb{R})$ of tempered distributions, containing the even Schwartz functions, such that:

(a) for each $\Delta > 0$ there is $C_\Delta < \infty$ with $\mathcal{P}_{\Delta,A}(h) \le C_\Delta \|\hat{h}\|X^2$ *for all A and all even Schwartz h; (b) there exists* $c\infty > 0$ and a fixed finite-dimensional $B \subset X$ such that $\mathcal{A}(h * h^\vee) \ge c_\infty \|\hat{h}\|X^2$ *on* $f \perp B$*; (c)* $c\infty > C_\Delta$ for some $\Delta > 0$; (d) the de Branges criterion $Q(h) \ge 0$ restricted to $\{h : \hat{h} \in X, \hat{h} \perp B\}$ is equivalent to the full criterion on even Schwartz h?

A positive resolution to Problem 6.1, combined with Theorem 3.1, would establish unconditionally that RH cannot be proved by the regularise-and-limit strategy within X.

A negative resolution — a proof that no such X exists — would also be a substantive result, in a different direction: it would show that the regularise-and-limit strategy is incompatible with RH-equivalence on structural grounds.

6.2 What is known

— The natural candidate $X = H^1_\omega$ with $\omega(\xi) = e^{-|\xi|/2}$ fails (a). (§4, Lemma 4.1.) — Heavier-weight variants fail (a) for the same structural reason. (§5.2.) — Lighter-weight or bandwidth-restricted variants fail (d). (§5.1, §5.2.) — Oscillation-constrained variants such as TPB-Adm may satisfy (a) but their satisfaction of (d) is not known and is suspected to fail. (§5.3.) — Distributional extensions may work via escape route (1) but constitute a different proof strategy, not a resolution of Problem 6.1. (§5.4.)

6.3 A reformulation

Problem 6.1 can be recast as a question about the compatibility of two requirements:

Smallness: The norm on X must be strong enough to control $\mathcal{P}_{\Delta,A}$ uniformly in A.
Largeness: The class $\{h : \hat{h} \in X\}$ must be wide enough for de Branges equivalence to RH.

The f_A counterexample shows that for the natural weighted-Sobolev family, these two requirements are in direct conflict: the weight that makes X "large enough" for de Branges is exactly the weight that allows the sampling to be unbounded.

It is plausible — though not proven — that this conflict is universal: that any Banach space satisfying (d) automatically fails (a), and vice versa. If so, Theorem 3.1 in its conditional form is unactivatable, and the meta-result reduces to a statement about the impossibility of finding a space where the activating hypotheses hold.

That, in itself, would be a coherent characterisation of why the explicit-formula route to RH has resisted analytic completion.

7. Status and Conclusion

7.1 What this paper establishes

Theorem 3.1 is a clean conditional: if hypotheses (a)–(d) hold for some Banach space X, then RH cannot be proved by the regularise-and-limit strategy within X.

Lemma 4.1 is an unconditional functional-analytic computation: the f_A family demonstrates that the natural candidate $X = H^1\_\omega(0, A)$ fails hypothesis (a).

Problem 6.1 is the precise open question into which the entire trilogy's analytic content reduces.

7.2 What the trilogy as a whole establishes

— A correct identification of the spectral primitive (Paper I, §§3–5): the dilation generator, not Schrödinger.

— A de Branges / Weil reformulation of RH as positivity $Q(h) \geq 0$ (Paper I, §F).

— The Carleson interval bound for prime-power density (Paper I, §7L), provable from PNT.

— A general no-upgrade theorem in functional analysis (Paper II): bounded forms cannot admissibly converge to unbounded ones.

— A precise conditional bridge (Paper III, Theorem 3.1) showing what would be needed to apply Paper II's theorem to the RH setup.

— A structural obstruction (Paper III, Lemma 4.1) showing the natural candidate space fails one of the bridge's hypotheses.

— A sharp open problem (Paper III, Problem 6.1) characterising the remaining analytic question.

7.3 What the trilogy does not establish

— Neither RH itself nor its negation.

— Unconditional finite-resolution positivity in any Banach space large enough for RH-equivalence.

— A no-go theorem foreclosing the explicit-formula route entirely.

The honest reading is that the trilogy identifies a specific structural feature of the RH problem — a tension between two requirements on the test-function space — and isolates the question of whether that tension is fatal or merely apparent.

7.4 Closing remark

The classical difficulty of RH may be encoded, at least in part, in the tension identified by Problem 6.1: that the very smoothing which makes prime-power sampling tractable also prevents the resulting positivity from being upgraded to the atomic statement equivalent to RH. The two requirements — *boundedness of the regularised observable* and *largeness of the test class* — pull in opposite directions, and the question is whether any framework reconciles them.

If yes: the regularise-and-limit strategy is foreclosed by Theorem 3.1, and a proof of RH via this route requires a genuinely different approach.

If no: the inability to find a reconciling framework is itself a structural feature of the problem, and the obstruction it represents may help explain why explicit-formula approaches have not closed.

Either resolution would be informative. Both leave RH itself untouched.

A single sentence captures the underlying interpretation:

The obstruction to the Riemann Hypothesis via explicit-formula positivity may not arise from insufficient analysis, but from attempting to recover an infinite-distinguishability object from a finite-distinguishability observable framework.

This is the philosophical statement the trilogy points at. Whether it is provable, refutable, or simply the right framing for an open problem remains to be settled.

References

[1] Berry, M. V. and Keating, J. P. "The Riemann zeros and eigenvalue asymptotics." *SIAM Review* 41 (1999), 236–266.

[2] de Branges, L. *Hilbert Spaces of Entire Functions.* Prentice-Hall, 1968.

[3] Connes, A. "Trace formula in noncommutative geometry and the zeros of the Riemann zeta function." *Selecta Mathematica* 5 (1999), 29–106.

[4] Kato, T. *Perturbation Theory for Linear Operators,* 2nd ed. Springer, 1980.

[5] Montgomery, H. L. "The pair correlation of zeros of the zeta function." *Proc. Symp. Pure Math.* 24 (1973), 181–193.

[6] Reed, M. and Simon, B. *Methods of Modern Mathematical Physics,* Vol. I–IV. Academic Press, 1972–1979.

[7] Rudin, W. *Functional Analysis,* 2nd ed. McGraw-Hill, 1991.

[8] Titchmarsh, E. C. *The Theory of the Riemann Zeta-Function,* 2nd ed. Oxford University Press, 1986.

[9] Weil, A. "Sur les 'formules explicites' de la théorie des nombres premiers." *Comm. Sém. Math. Lund* (1952), 252–265.

[10] Taylor, K. *A Spectral Framework for the Riemann Hypothesis: The VERSF Approach.* AIDA Institute, 2025. [Paper I of this trilogy.]

[11] Taylor, K. *Finite Resolution Does Not Converge to Infinite Resolution: A No-Upgrade Theorem for Singular Observables.* AIDA Institute, 2025. [Paper II of this trilogy.]