Auburn Patent Family Fields

FIELDS Adversarial Protocol

FAP v2.0

A Structured Hostile-Referee Framework
for Mathematical Audit

Specification and First Demonstration

Ryan Christopher Fields

March 2026

Do not ask whether the proof looks correct.
Ask how it fails — and make that failure impossible.

UncleBroFields@proton.me | fieldsryanchristopher@gmail.com

Contents

Part I

The Protocol

Purpose and Scope

1 What FAP Is

The FIELDS Adversarial Protocol is a structured framework for pre-submission adversarial audit of mathematical manuscripts. It does not aim to confirm that a proof appears correct. It systematically identifies and executes the principal failure routes through which a hostile expert referee would attempt to invalidate the manuscript.

The goal is to harden proofs against adversarial scrutiny by exhaustively testing theorem-critical vulnerabilities prior to external review.

2 What FAP Is Not

- Not a proof generator.
- Not a replacement for formal verification.
- Not a certification of correctness.
- Not a substitute for human mathematical judgement.

FAP is a referee simulation and adversarial audit system. It guarantees systematic exposure of all standard high-risk failure modes. It does not guarantee that no failure mode exists beyond its taxonomy.

3 The Core Principle

Assume the main theorem is false unless every plausible failure route is explicitly neutralised by the manuscript.

This shifts the evaluation from passive verification to active adversarial attack. The auditor's job is not to read the proof and check whether each step looks right. The auditor's job is to find every way the proof could fail and verify that the manuscript closes each one.

4 How v2.0 Differs from v1.0

v1.0 audits the mathematics. v2.0 audits the mathematics *and* the auditors.

v1.0	**v2.0**
16 attack routes	16 attack routes + meta-audit
Single-system execution	Multi-model cross-validation
One pass type	Governed + ungoverned passes
Qualitative verdicts	Quantitative thresholds
Report output	Reproducibility package
No convention awareness	Convention preloading
No failure mode taxonomy	FM 1–22 integration
No axiom verification	Permanent Axiom layer

5 Applicable Domains

FAP is applicable to: theoretical mathematics (analysis, geometry, algebra, combinatorics), mathematical physics, formally verified proofs, and computational or hybrid proofs with explicit trusted base.

The Attack Taxonomy

FAP defines sixteen classes of hostile-referee attack. Every audit must execute all applicable routes.

Tests: Incorrect or overstated theorem. Missing hypotheses. Quantifier errors. Scope inflation.

Pattern: "The theorem as stated is false. Here is a counterexample to the literal claim."

Resolution standard: The theorem statement must be verified against the proof term (Coq `Check`) or against the precise conditions used in the proof. Every quantifier must match.

Tests: Ambiguous or inconsistent definitions. Informal notions exceeding formal definitions. Notation drift.

Pattern: "Definition X is used in Lemma Y with a meaning that differs from how it was introduced."

Resolution standard: Every definition must be used consistently throughout. If a definition appears in the Coq, the Coq definition is authoritative.

Tests: Circular reasoning. Unstated assumptions. Misuse of prior results. Hidden stronger hypotheses.

Pattern: "Theorem A uses Lemma B which assumes Theorem A." Or: "The proof uses property P which is never stated as a hypothesis."

Resolution standard: A complete dependency graph (DAG) must be constructed. `Print Assumptions` provides the machine-verified dependency closure for Coq results.

Tests: Invalid implications. Missing intermediate steps. Illegitimate constructions or decompositions.

Pattern: "Step 3 claims $A \Rightarrow B$ but the implication requires an intermediate step that is not provided."

Resolution standard: For Qed results, the Coq kernel has verified the implication. For manuscript-only arguments, each step must be individually justified.

Tests: Improper ordering of quantifiers. Non-uniform parameter choices. Illegal interchange of limits, sums, or extrema.

Pattern: "The constant C depends on j, but the theorem claims it is uniform."

Resolution standard: For Coq results, the universally quantified type of the final theorem shows exactly what is uniform. For manuscript arguments, uniformity must be verified at each step.

Tests: Incorrect exponent tracking. Hidden losses (ε-loss, constants). Failure of summability. Illusory gains.

Pattern: "The exponent accounting claims $1/2 + 2 + 1 = 7/2$ but step 2 actually contributes $5/2$, not 2."

Resolution standard: For Coq results, the `ring` tactic provides machine-verified exponent arithmetic. For manuscript arguments, every exponent must be tracked line by line.

Tests: Over/undercounting. Incorrect dimension reduction. Mischaracterised interaction sets.

Pattern: "The collinear set is claimed to have $O(k_j)$ points but the actual count is $O(k_j^2)$."

Resolution standard: Counting bounds must be verified against standard lattice geometry references or machine-checked.

Tests: Heuristic structure not proved quantitatively. "Generic" or "transverse" used without bounds. Incomplete exclusion of bad configurations.

Pattern: "The proof claims the geometric factor is first harmonic in ϕ but does not rule out higher harmonics."

Resolution standard: Geometric claims must be either proved in Coq or justified by explicit computation in the manuscript. If neither, the claim is classified as a Hypothesis with its resolution status.

Tests: Failure at endpoints or degenerate regimes. Critical parameter breakdown. Small-scale or limiting-case inconsistencies.

Pattern: "At $j = 1$, the resonant set may be empty. Does the bound degenerate?"

Resolution standard: The theorem must hold at all parameter values within its stated scope, including boundary cases.

Tests: Failure of local-to-global composition. Accumulation of errors across iterations/scales. Incompatibility of decomposed components.

Pattern: "Each shell bound is correct but the sum over shells diverges."

Resolution standard: The local-to-global transition must be explicitly justified. Summability must be verified.

Tests: Illegal operations under stated regularity. Invalid analytic manipulations. Missing compactness or convergence justification.

Pattern: "The proof interchanges a sum and a limit without justification."

Resolution standard: For Coq results with finite sums, this is automatic. For manuscript arguments involving limits, each interchange must be justified by the stated regularity.

Tests: Construction of plausible violating examples. Failure to rule out minimal adversarial cases.

Pattern: "Here is a divergence-free field where the bound fails."

Resolution standard: Either exhibit the counterexample concretely or explain why no counterexample can exist within the theorem's hypotheses.

Tests: Misapplication of cited results. Omitted assumptions from imported theorems.

Pattern: "The cited theorem requires hypothesis H which is not satisfied here."

Resolution standard: Every imported result must be verified to apply in the specific setting. For Permanent Axioms, the Trust Boundary Verification document provides this.

Tests: Inflated claims of originality. Misrepresentation of contribution scope.

Pattern: "The half-derivative improvement was already known from [reference]."

Resolution standard: Novelty claims must be calibrated: state precisely what is new and what is standard technique applied in a new context.

Tests: Persuasive language masking gaps. Hidden assumptions embedded in prose.

Pattern: "Section 4.4 creates the impression that the phase analysis is load-bearing when it is not."

Resolution standard: Non-load-bearing content must be explicitly scoped. Load-bearing content must be identifiable from the dependency graph alone.

Tests: Misalignment between formal proof and stated theorem. Hidden axioms or computational assumptions.

Pattern: "The Coq `in_shell` definition uses dyadic bands but the theorem claims unit-width shells."

Resolution standard: `Print Assumptions` provides the authoritative dependency closure. Non-load-bearing definitions must be identified and documented. Convention differences must be preloaded.

v2.0 Additions

The following ten components are new in v2.0. Each addresses a gap identified during the Auburn Roses governance programme.

6 Failure Mode Taxonomy Integration

Every objection generated during a FAP pass is cross-referenced against the FIELDS Failure Mode Taxonomy (FM 1–22) before reaching the resolution phase.

Rule: If an objection matches a documented failure mode pattern and cannot produce a specific broken inference with a concrete counterexample, it is classified as *manufactured* and excluded from the severity assessment. It is still documented in the audit trace — but it does not count against the manuscript.

The FM taxonomy is documented in Common Failure Modes Parts 1–5. The most frequent manufactured objection patterns are:

FM	Name	Trigger Pattern
1	Training-Prior Hijack	"This can't be right because the problem is open"
9	Modular Formalisation Blindness	"Hypothesis is a cheat code"
11	Compulsive Objection Generation	Multiple concerns with no specific broken line
12	Convention Mismatch Escalation	"Shell mismatch between manuscript and Coq"

A manufactured objection is not suppressed. It is *labelled*. The label prevents it from inflating the severity assessment while preserving the audit record.

7 Multi-Model Cross-Validation

Every FAP pass must be executed on a minimum of two independent systems.

- If both systems generate the same objection independently: the objection is *likely genuine*. Proceed to resolution.

- If only one system generates the objection: apply the four-test triage (Section 17).
- If neither system generates the objection under adversarial prompting: the attack route is *clean* for this manuscript.

Empirical basis

The Auburn Roses programme demonstrated that single-system ungoverned audits produce false positive rates up to 87.5% (14 manufactured objections out of 16 total across three systems). Cross-system concordance reduces this dramatically: objections that appeared independently in multiple systems were overwhelmingly genuine.

8 Two-Pass Architecture

Every FAP audit includes two passes:

	Pass A (Governed)	Pass B (Ungoverned)
Protocol	FAP active, structured	No protocol, raw adversarial
Noise level	Low	High
Precision	High	Moderate
Purpose	Mathematical assessment	Stress-test the governance
Catches	Genuine errors efficiently	Errors in operator's claims, editorial residuals, governance gaps

Neither pass alone covers the full failure space. Pass A produces the mathematical assessment. Pass B catches what governed deference misses — including errors in the *auditor's own responses*.

Empirical finding: In the Auburn Roses programme, the governed pass identified 18 genuine corrections with zero false positives. The ungoverned pass identified 2 additional catches (one operator error, one editorial residual) embedded in 14 manufactured objections. The optimal architecture uses both.

9 Convention Preloading

Before any FAP pass begins, all known convention differences between representations of the same mathematics must be stated explicitly in the pass preamble.

Examples:

- Manuscript uses unit-width shells; Coq uses dyadic bands in a non-load-bearing definition.
- $Z_{\text{Coq}} = \nu \cdot Z_{\text{manuscript}}$ (dissipation includes viscosity factor).
- Squared form vs. square-root form of the main inequality.

Empirical basis

FM 12 (Convention Mismatch Escalation) was the single most persistent manufactured objection in the Auburn Roses programme. It appeared in every session, every system, at every granularity, under every governance condition. Convention preloading eliminated it entirely when applied.

10 Permanent Axiom Verification Layer

For any result that axiomatises external content (imported theorems, classical results, library limitations), FAP v2.0 requires independent verification of every axiom.

For each axiomatised result, produce:

1. The precise mathematical statement encoded.
2. The original reference (specific paper or theorem).
3. A modern textbook reference confirming continued validity.
4. Confirmation that the result holds in the specific setting of the proof.
5. Documentation of why it cannot currently be formalised.
6. Verdict: Classical and Standard or Requires Further Justification.

Empirical basis

The Auburn Roses programme independently verified 17 Permanent Axioms. Result: 17/17 Classical and Standard. References spanning 1919–1989. This verification closed the last open objection from adversarial auditing: "the axioms are trust-based."

11 Contradiction Register Requirement

Every error discovered during the FAP process is permanently documented in a Contradiction Register:

- What was wrong.
- When it was found.
- Which system found it.
- Which systems missed it.
- What replaced it.
- Whether the replacement has been verified.

The Register prevents FM 20 (Transparency Weaponisation): when error correction is permanently documented, it cannot be cited as evidence of unreliability. Published errors are evidence of governance rigour.

12 Quantitative Verdict Criteria

v2.0 replaces qualitative verdict descriptions with quantitative thresholds:

Verdict	**Threshold**
Hostile-Referee Resistant	Zero Class I or Class II unresolved. All Class III resolved or documented with explicit fix. All 16 attack routes executed. Cross-system validation complete.
Closed Modulo Exposition	Zero Class I. At most 2 Class II with documented resolution paths. All Class III acknowledged.
Technically Close	Zero Class I. Class II objections have partial resolutions with identified remaining work.
Materially Incomplete	Any Class II without a resolution path.
Not Closed	Any Class I unresolved.

13 Reproducibility Package

A valid FAP v2.0 audit must output a complete reproducibility package containing:

1. The exact prompts used for each FAP pass.
2. The system, version, and configuration for each pass.
3. The complete response transcripts.
4. The FM classification of every generated objection.

5. The resolution trace for every genuine objection.
6. The dependency graph.
7. The counterfactual break analysis.
8. The Permanent Axiom verification (if applicable).
9. The convention preloading document (if applicable).
10. The Contradiction Register (if errors were found).

Anyone can rerun the entire audit from this package and verify the verdicts independently.

14 The "Stronger Than Published" Test

After all FAP passes are complete, execute one final test:

> Identify a result in the published literature that is weaker than the result under audit but was accepted by a top journal. Document: what objections would FAP have raised against that published result? Would it have survived?

If the published result would NOT have survived FAP v2.0, then the manuscript under audit has been held to a higher standard than the existing literature. This is documented as part of the final verdict.

15 Load-Bearing Analysis and Counterfactual Break

FAP v2.0 requires explicit identification of:

- **Load-bearing lemmas:** results whose failure would propagate to the main theorem.
- **Bottleneck arguments:** steps where the proof is tightest (smallest margin).
- **Single points of failure:** components without redundancy.

A dependency graph (DAG) must be constructed linking definitions $\rightarrow$ lemmas $\rightarrow$ propositions $\rightarrow$ main theorem.

For each load-bearing component, a counterfactual analysis is performed:

> **Assume the component is false. What fails?**

This produces a fragility map of the proof, enabling referees and readers to identify exactly which components carry the mathematical weight and which are supporting infrastructure.

Pass Structure and Triage

16 Pass Output Requirements

Each attack route, when executed, must produce five outputs:

Output	Description
Objection	A precise, high-strength hostile claim.
Target	The exact theorem, lemma, or step under attack.
Failure consequence	What downstream results collapse if the objection holds.
Resolution trace	Exact location(s) in the manuscript or Coq addressing the objection.
Verdict	One of: *Unresolved*, *Partially Resolved*, *Resolved But Under-Justified*, *Resolved Cleanly*.

Every objection additionally receives an FM cross-check: is this objection a genuine mathematical concern or a documented failure mode pattern?

17 The Four-Test Triage

When a single system generates an objection that the other system does not, the objection is triaged:

Test	Question	If No
T1	Does it identify a specific broken inference?	Likely manufactured.
T2	Can the system exhibit a concrete counterexample?	Likely manufactured.
T3	Does it target a Qed theorem or a Hypothesis?	If Hypothesis: check axiom verification.
T4	Which failure mode does it match?	Classify and label.

An objection that fails T1 and T2 is labelled as manufactured. An objection that passes all four tests is treated as genuine and proceeds to resolution.

Severity Classification

Every objection that survives triage is classified:

Class	Name	Meaning
I	Fatal	Invalidates the main theorem.
II	Major	Blocks acceptance until resolved. A referee would require revision.
III	Technical	Repairable gap. Does not affect the main result but requires clarification or tightening.
IV	Expository	Clarity issue only. No mathematical content affected.
V	Positioning	Novelty or claim calibration. The mathematics is correct but the framing is imprecise.

Final Verdict Classes

The FAP audit concludes with one of five verdicts:

Verdict	Definition
Not Closed	One or more Class I objections remain unresolved. The main theorem is not established.
Materially Incomplete	No Class I, but one or more Class II objections lack a resolution path. Core gaps exist.
Technically Close	No Class I. Class II objections have partial resolutions with identified remaining work.
Closed Modulo Exposition	No Class I. At most 2 Class II with documented paths. Mathematically sound; clarity issues remain.
Hostile-Referee Resistant	Zero Class I or II unresolved. All Class III resolved or documented. All 16 routes executed. Cross-system validation complete. No hostile referee operating in good faith could reject this result on mathematical grounds.

The verdict **Hostile-Referee Resistant** does not mean the result is correct. It means that the standard repertoire of hostile-referee attacks has been exhaustively executed and every attack has been neutralised by the manuscript. The distinction between "no attack succeeded" and "the result is true" is the distinction between adversarial audit and mathematical proof. FAP provides the former. The Coq kernel provides the latter.

v3.0 Roadmap

The following extensions are documented for future development. They are not part of the v2.0 specification.

18 Model-Specific Persona Calibration

v2.0 requires multi-model cross-validation but does not assign specific attack roles to specific systems. v3.0 would assign hostile-referee personas based on empirical calibration of each system's strengths:

- Systems with low false positive rates → Logician, Analyst, Edge-Case Assassin (precision roles).
- Systems with high correction quality → Geometer, Combinatorialist, Formalist (depth roles).
- Systems with high coverage → Compression Test, Black-Box Expansion, Seminar Interruption (stress-test roles).

19 Time-Stamped Prediction Registry

For results with empirical implications, v3.0 would require: state the predictions, timestamp them, publish them, and specify falsification criteria before the verification window opens. This converts theoretical results into testable claims with independent empirical support.

20 Cross-Domain Expansion

v2.0 was developed and tested on PDE analysis with Coq formalisation. v3.0 would extend the attack taxonomy and calibration data to: algebra, geometry, number theory, combinatorics, computer science, and Lean/Isabelle formalisations.

21 Journal Editorial Workflow Integration

v3.0 would produce FAP audit reports in a format suitable for submission alongside manuscripts as supplementary material, providing referees with a pre-executed adversarial audit that they can verify rather than reproduce from scratch.

22 Community Contribution Framework

v3.0 would standardise the format for contributed deployment data: which attack routes were executed, which objections were generated, which failure modes were observed, and how the manuscript responded. Every deployment refines the calibration.

23 Formal Verification of the Protocol

The ultimate goal: formalise the FAP protocol itself in a proof assistant, proving that the attack taxonomy is exhaustive relative to a formal threat model and that the triage procedure correctly classifies manufactured vs. genuine objections.

> v3.0 is where the community takes over. The specification is published. The methodology is documented. The first demonstration is provided. What remains is for others to deploy it, test it, break it, and build on it.

Part II

First Demonstration

Audit Configuration

Parameter	**Value**
Target	Sections 1–4 of *The Angular Cancellation Lemma*
Main theorem	Theorem 4.27: $\lvert VS_j \rvert^2 k_0 \leq C_{\mathrm{ACL}}^2 k_j^7 E_j^3$
Coq file	`NavierStokesACL.v` (559 lines, 12 Qed, 3 Axioms, 0 Admitted)
Auditing system	Model A-C (Claude Opus 4.6, Extended Thinking)
FM taxonomy	Active (cross-referencing FM 1–22)
Convention preloading	Active (unit-width shells, Z convention, `in_shell` non-load-bearing status)
Pass type	Pass A (Governed)

Attack Route Execution

24 Route 4.1: Statement Attack

Objection 1.1: Is the theorem correctly scoped? It claims "for every unit-width shell S_j on $\mathbb{Z}^3$ and every divergence-free velocity field." Does it require regularity beyond divergence-free?

Target: Theorem 4.27.
Failure consequence: If additional regularity is needed, the theorem as stated is false for rough data.
Resolution trace: The proof uses only finite sums over lattice points, abstract $\mathbb{R}^3$-valued Fourier data, and the hypothesis `div_free`. No Sobolev regularity is invoked. Remark 4.25 confirms: "no assumption about amplitude concentration or phase alignment is used." The Coq theorem is universally quantified over all `u_hat : Z3 -> R3` satisfying `div_free`.
Verdict: Resolved cleanly.

Objection 1.2: The "equivalently, taking square roots" claim requires non-negativity. Is this justified?

Target: Theorem statement (squared vs. square-root form).
Failure consequence: If any factor is negative, the square root is invalid.
Resolution trace: $k_0 > 0$ (`hk0`), $C_{\text{ACL}} > 0$ (`hCACL`), $k_j > 0$ (`hkj`), $E_j \geq 0$ (`hEj`). All terms non-negative.
Verdict: Resolved cleanly.
Severity: Class IV (expository).

25 Route 4.2: Definition Attack

Objection 2.1: Definition 4.19 defines $\xi_k = \hat{u}(k)/|\hat{u}(k)|$, undefined when $\hat{u}(k) = 0$.

Target: Definition 4.19 (geometric factor).
Failure consequence: Division by zero in the geometric factor decomposition.
Resolution trace: The polarisation ξ_k is used only in the geometric motivation (Lemma 4.20), not in the proof chain. The Coq proof works with `coupling`(k, p) directly, which equals zero when $\hat{u}(k) = 0$. The geometric factor decomposition is expository.
FM check: Matches FM 18 (Scope Drift) — targeting motivational content rather

than load-bearing content. Confirmed non-load-bearing.
Verdict: Resolved cleanly.

Objection 2.2: The transversality parameter c_0 is described as "universal" and "$\ll 1$" but not specified. Does the proof depend on a specific value?

Target: Definition 4.21, Remark 4.23.
Failure consequence: If c_0 must be tuned per shell, the "universal" claim is false.
Resolution trace: Remark 4.23: "c_0 is fixed once and for all, independent of j and k_j." $C_{\text{ACL}}(c_0) = C(c_0^{-1/2} + c_0^{1/2})$ is finite for any fixed $c_0 \in (0, 1)$. The Coq encodes c_0 as a rational `c0_num/c0_den` with `hc0_small`: $c_0 < 1$. The proof works for ALL $c_0 \in (0, 1)$.
Verdict: Resolved cleanly.

Objection 2.3: The Euclidean norm $|k|$ for $k \in \mathbb{Z}^3$ is generally irrational. Is the shell $j \leq |k| < j + 1$ well-defined on the lattice?

Target: Remark 4.26 (shell definition).
Failure consequence: Ambiguous set membership.
Resolution trace: $j \leq |k| < j + 1$ is equivalent to $j^2 \leq k_1^2 + k_2^2 + k_3^2 < (j + 1)^2$, which involves only integers. The Coq uses `norm2Z` (squared norm in $\mathbb{Z}$), avoiding square roots entirely.
Verdict: Resolved cleanly.

26 Route 4.3: Dependency Attack

Objection 3.1: Does Theorem 4.27 depend on the phase dynamics results (Section 4.4)?

Target: Theorems 4.10, 4.12, 4.15, 4.16.
Failure consequence: If the phase results are load-bearing and flawed, the ACL inherits the flaw.
Resolution trace: Remark 4.17 explicitly states: "They are not load-bearing in the proof chain. The argument proceeds via the Angular Cancellation Lemma (§4.5)."
Verdict: Resolved cleanly.

Objection 3.2: Does the ACL depend on the shell model verification (Theorem 4.34)?

Target: Section 4.7.
Failure consequence: If the shell model is wrong and the ACL depends on it, the ACL is unjustified.
Resolution trace: Remarks 4.33 and 4.35: "This does not imply enstrophy closure for the full Navier–Stokes equations; it verifies that the ACL mechanism is consistent

at the shell level." Non-load-bearing.
Verdict: Resolved cleanly.

Objection 3.3: Circularity check — does Theorem 4.27 use any result that depends on Theorem 4.27?

Target: Full dependency chain.
Failure consequence: Circular proof.
Resolution trace: The dependency chain is linear: `div_free` → `factor_A` → `coupling_eq` → `dot_sq_le` → `coupling_sq_shell` → `trans_sum_sq` → `per_mode_res_bound` → `per_shell_combined` → `ACL_shell_sq` → `angular_cancellation_lemma`. No theorem references the final result. `Print Assumptions` confirms: no circularity.
Verdict: Resolved cleanly.

27 Route 4.4: Local Logical Gap Attack

Objection 4.1: The transverse + collinear decomposition must be exact (not approximate) for the bounds to compose. Is the partition proved?

Target: Lemma 4.24(c) assembly step.
Failure consequence: If the partition leaks terms, the bound has a gap.
Resolution trace: `res_split` (Coq, lines 287–301) proves the partition as an exact identity by structural induction on the resonant list. Every element is either `is_collinear` or not (Boolean complement). Machine-verified. Qed.
Verdict: Resolved cleanly.

Objection 4.2: Step 2 of Theorem 4.27 uses $|\hat{\omega}(k)| \le 2k_j |\hat{u}(k)|$ (Bernstein). Where does the factor 2 originate?

Target: Theorem 4.27, Step 2.
Failure consequence: If the factor is wrong, the per-mode bound is off.
Resolution trace: For $k \in S_j$ with $k_j = j$: $|k| < j + 1 = k_j + 1 \le 2k_j$ for $j \ge 1$. The factor 2 is a safe upper bound. The Coq uses the tighter Hypothesis `shell_wavenumber`: $|k|^2 \le k_j^2$.
Verdict: Resolved cleanly.
Severity: Class III (the factor 2 is conservative; Coq uses the sharper bound).

Objection 4.3: Step 4 uses $\sqrt{\#S_j} = O(k_j)$. Is the shell count $\#S_j = O(k_j^2)$ proved or assumed?

Target: Theorem 4.27, Step 4.
Failure consequence: If $\#S_j = O(k_j^3)$ (volumetric), the exponent is wrong.
Resolution trace: Coq: Hypothesis `shell_size_bound`: $|S_j| \le k_j^2$. Encodes the Gauss sphere counting result for unit-width shells. Independently verified in

An Assessment of Shells with Chamizo–Iwaniec (1995) and Heath-Brown (1999) references.
Verdict: Resolved — the resolution is a Hypothesis encoding standard lattice geometry.
Severity: Class III (axiomatised, but mathematically standard).

28 Route 4.5: Quantifier / Uniformity Attack

Objection 5.1: Is C_{ACL} uniform in j?

Target: Theorem 4.27.
Failure consequence: If C_{ACL} depends on j, the bound degrades at high frequency.
Resolution trace: C_{ACL} depends on c_0 only. c_0 is fixed, independent of j (Remark 4.23). The Coq declares `C_ACL` as a single `Variable` with one positivity Hypothesis. The final theorem is $\forall j$, `(VSj j)^+2 * k0 <= C_ACL^+2 * (kj j)^+7 * (Ej j)^+3`. Same `C_ACL` for every j.
Verdict: Resolved cleanly. Machine-verified uniformity.

Objection 5.2: Is the counting bound $\#\mathcal{C}_{k,j} = \Theta(k_j)$ uniform in k?

Target: Lemma 4.22(i).
Failure consequence: If there are exceptional k values where the resonant set is much larger, the bound fails for those modes.
Resolution trace: Lemma 4.22(i): "$c_1 k_j \le N_j \le c_2 k_j$ for universal constants, *uniformly in* k." Coq Hypothesis `coll_count`: universally quantified over both k and j.
Verdict: Resolved cleanly.

29 Route 4.6: Estimate / Exponent Attack

Objection 6.1: Exponent accounting: $1/2 + 2 + 1 = 7/2$?

Target: Theorem 4.27, exponent assembly.
Failure consequence: If the arithmetic is wrong, the stated exponent is incorrect.
Resolution trace: 1/2 (transverse CS saving) + 2 (amplitude bound) + 1 (outer shell count) = 7/2. Correct. The Coq verifies: `per_shell_combined` proves k_j^5 ($= 2 + 3$ from per-mode k_j^3 times shell count k_j^2). `ACL_shell_sq` multiplies by $k_0 \le k_j^2$ to get k_j^7 ($= 5 + 2$). The `ring` tactic verifies $k_j^5 \cdot k_j^2 = k_j^7$.
Verdict: Resolved cleanly. Machine-verified exponent arithmetic.

Objection 6.2: What is the "standard estimate" against which the $k_j^{1/2}$ improvement is claimed?

Target: Header claim and Remark 4.2.
Failure consequence: If the comparison baseline is wrong, the improvement claim is overstated (FM 14 / Novelty attack).
Resolution trace: Remark 4.28: dyadic shells would give $|VS_j| \le Ck_j^4 E_j^{3/2}$. Unit-width shells with the transverse CS saving give $k_j^{7/2}$. Improvement: $k_j^4/k_j^{7/2} = k_j^{1/2}$. The manuscript is consistent.
Verdict: Resolved cleanly.

30 Route 4.7: Counting / Combinatorial Attack

Objection 7.1: Lemma 4.22(iii) claims $\#\mathcal{C}^{\mathrm{coll}} \le C_{\mathrm{coll}} \cdot c_0 \cdot k_j + C_0$. Is the lattice counting argument correct?

Target: Lemma 4.22(iii).
Failure consequence: If $\#\mathcal{C}^{\mathrm{coll}} = O(k_j^2)$, the collinear contribution overwhelms the transverse saving.
Resolution trace: The proof uses: resonant points lie in a slab around an affine plane Π. The collinear subset lies in the intersection of Π with a cylinder of radius $c_0 k_j$. The sublattice has covolume $\sim k_j$. Lattice points in a planar region of area $\sim c_0 k_j^2$: count $\lesssim c_0 k_j^2/k_j = c_0 k_j$. Standard lattice geometry (Chamizo–Iwaniec).
Coq Hypothesis `coll_count`: $|\mathcal{C}^{\mathrm{coll}}| \le k_j$ (conservative: omits c_0 factor since $c_0 < 1$).
Verdict: Resolved. Standard lattice counting, axiomatised in Coq.
Severity: Class III (Coq Hypothesis is conservative relative to the manuscript's sharper bound).

31 Route 4.8: Geometric Structure Attack

Objection 8.1: Lemma 4.20 claims Factor B is a pure first harmonic in ϕ: $B(\phi) = R(\theta)\cos(\phi - \phi_0)$. If higher harmonics exist, the CS saving degrades. Is this proved?

Target: Lemma 4.20, Factor B.
Failure consequence: If $B(\phi)$ has an $O(1)$ constant term, the azimuthal cancellation is incomplete and the $k_j^{1/2}$ saving is illusory.
Resolution trace: The manuscript provides an explicit coordinate computation (proof of Lemma 4.20): the Leray projector $\mathcal{P}_q = I - \hat{q} \otimes \hat{q}$ is quadratic in $\hat{q}$, but the orthogonality constraint $\xi_q \cdot \hat{q} = 0$ reduces the degree by one. The constant term vanishes because the ϕ-averaged projection of any vector perpendicular to the rotation axis onto a rotating plane is zero.
Critical assessment: This argument is given analytically in the manuscript but

is NOT verified in Coq. The Coq bypasses the harmonic decomposition by using the abstract Hypothesis `per_mode_res_bound`, which encodes the combined bound without decomposing it into azimuthal structure. The first-harmonic property is supporting motivation, not a proved lemma in the formalisation.
FM check: This is NOT FM 22 (heuristic accusation without broken line). The objection identifies a specific gap: the first-harmonic property is argued but not formalised. The load-bearing content is the Hypothesis `per_mode_res_bound`.
Verdict: Resolved but under-justified. The geometric argument is convincing but not machine-verified. The Coq accepts the combined bound as a Hypothesis.
Severity: Class III (motivating geometry is sound but not formalised; the load-bearing content is the Hypothesis).

32 Route 4.9: Edge-Case Attack

Objection 9.1: At $j = 1$, the resonant set may be very small or empty. Does the bound degenerate?

Target: Theorem 4.27 at $j = 1$.
Failure consequence: If the bound produces 0/0 or vacuous results at $j = 1$, the theorem is incomplete.
Resolution trace: At $j = 1$: $k_j = 1$, bound gives $|VS_1|^2 k_0 \leq C^2 E_1^3$. Shell S_1 is finite ($|k|^2 \in \{1, 2, 3\}$). Cauchy–Schwarz is valid for finite sums of any size, including empty sums (which give zero). The bound may be loose at $j = 1$ but is not false.
Verdict: Resolved cleanly.

Objection 9.2: If $E_j = 0$, the bound gives $|VS_j|^2 \leq 0$, implying $VS_j = 0$. Is this correct?

Target: Theorem 4.27 with $E_j = 0$.
Failure consequence: If $E_j = 0$ but $VS_j \neq 0$, the bound is false.
Resolution trace: $E_j = 0 \Rightarrow \hat{u}(k) = 0$ for all $k \in S_j$ (since $E_j = \sum |\hat{u}(k)|^2$ with non-negative terms). Therefore $\texttt{coupling}(k, p) = 0$ for all $k \in S_j$. Therefore $VS_j = 0$.
Verdict: Resolved cleanly. Degenerate case is trivially true.

33 Route 4.10: Global Assembly Attack

Objection 10.1: The ACL bounds each shell individually. The enstrophy involves $\sum_j VS_j$. Does the shell-by-shell bound compose correctly?

Target: Section 4.6 (enstrophy equation).
Failure consequence: If cross-shell coupling invalidates the composition, the ACL is insufficient.

Resolution trace: Section 4.6 writes vortex stretching as $\sum_j VS_j$. Triangle inequality gives $|\sum VS_j| \leq \sum |VS_j|$. The ACL bounds each $|VS_j|$. Cross-shell coupling occurs via KP-1 and KP-2 commutator terms, handled in Sections 5–6, not in the ACL. Remark 4.31 and the header (lines 28–32) explicitly state this scoping.
Verdict: Resolved cleanly. The ACL is correctly scoped to shell-local contributions.

34 Route 4.11: Regularity / Admissibility Attack

Objection 11.1: Are all sums absolutely convergent?

Target: All lattice sums in the proof.
Failure consequence: If a sum conditionally converges, reordering could change the value.
Resolution trace: All sums are finite. S_j is finite (finitely many lattice points with $j^2 \leq |k|^2 < (j+1)^2$). $\mathcal{C}_{k,j}$ is finite for each k. The Coq works with list sums (finite sequences). Absolute convergence is automatic.
Verdict: Resolved cleanly.

35 Route 4.12: Counterexample Attack

Objection 12.1: Can we construct a divergence-free field where the ACL bound is tight or violated?

Target: Theorem 4.27 (optimality).
Failure consequence: If a counterexample exists, the theorem is false.
Resolution trace: Remark 4.25 addresses amplitude concentration: (1) the transverse bound holds for arbitrary amplitudes (deterministic CS), (2) the collinear bound is controlled by count, (3) the collinear piece scales as $k_j^{1/2}$ (same as transverse), (4) ℓ^2 norm controlled by Parseval (an identity).
For tightness: single-mode concentration gives a ratio of $k_j^{5/2}|\hat{u}|$ between bound and actual — not tight. The bound is tightest for energy spread across $\sim k_j$ transverse modes.
No counterexample within the hypotheses was found or is expected.
Verdict: Resolved cleanly.

36 Route 4.13: Reference / Import Attack

Objection 13.1: Lattice counting references (Chamizo–Iwaniec, "[2, Ch.12]") — correctly applied?

Target: Lemma 4.22(iii).
Resolution trace: The application is standard: lattice points in a planar region of bounded area with uniform constants. Independently verified in *An Assessment of Shells*.
Verdict: Resolved cleanly.

Objection 13.2: Tao barrier (Theorem 4.3, "[4]") — accurately stated?

Target: Theorem 4.3.
Resolution trace: The statement matches Tao's result: an averaged NS system with the same energy identity and Sobolev bounds that admits finite-time blowup.
Verdict: Resolved cleanly.

37 Route 4.14: Novelty / Overclaim Attack

Objection 14.1: Is the half-derivative improvement genuinely new?

Target: Remark 4.18 (novelty claim).
Resolution trace: The specific combination — incompressibility-forced azimuthal oscillation on $\mathbb{Z}^3$, transverse-collinear decomposition, amplitude-geometry decoupling — does not appear in the prior literature. The technique (CS on lattice sums) is standard; the application to the NS triadic kernel is new. Remark 4.18 calibrates the claim appropriately.
Verdict: Resolved cleanly.

Objection 14.2: The manuscript claims the ACL "passes the Tao test" (Remark 4.30). Justified?

Target: Remark 4.30.
Resolution trace: The ACL uses $\hat{u}(k) \perp k$ via `factor_A`. Without this constraint, the coupling scales as $|\hat{u}(p)| \cdot |q|$ with no cancellation. The transversality decomposition depends on the Leray projector. These are specific to the NS bilinear form, not shared by generic operators.
Verdict: Resolved cleanly.

38 Route 4.15: Exposition Integrity Attack

Objection 15.1: Section 4.4 (phase dynamics) is lengthy and could create the impression that these results are load-bearing.

Target: Sections 4.2–4.4.
Resolution trace: Remark 4.17 explicitly scopes: "They are not load-bearing." The header states the file does NOT prove global regularity. The Coq file does not contain any code corresponding to Section 4.4.
Verdict: Resolved cleanly.

39 Route 4.16: Formal / Trusted Base Attack

Objection 16.1: The 3 Axioms (`CS_sum`, `CS_list`, `list_sum_le_size_max`) — trustworthy?

Target: Trusted base of the formalisation.
Resolution trace: All three are Cauchy–Schwarz / triangle inequality variants for finite sums. Provable from ordered field axioms. Axiomatised due to MathComp library gap, not mathematical uncertainty. Documented in the FIELDS-COQ Deep Dive and Permanent Axiom Verification.
Verdict: Resolved cleanly.

Objection 16.2: `in_shell` uses dyadic bands; the manuscript uses unit-width shells.

Target: `in_shell` definition (line 134).
FM check: FM 12 (Convention Mismatch Escalation). Neutralised by convention preloading.
Resolution trace: Non-load-bearing. `Print Assumptions` does not list it. The theorem is universally quantified over abstract shell variables. Documented in *An Assessment of Shells* and the FIELDS-COQ Deep Dive.
Verdict: Resolved cleanly.

Load-Bearing Analysis

The load-bearing components of Sections 1–4 for Theorem 4.27:

#	Component	Status	Role
1	`div_free`	Hypothesis	Incompressibility
2	`factor_A`	Qed	$\hat{u}(p) \cdot q = \hat{u}(p) \cdot k$
3	`dot_sq_le`	Qed	Cauchy–Schwarz
4	`coupling_sq_shell`	Qed	Per-mode bound
5	`res_split`	Qed	Transverse/collinear partition
6	`trans_sum_sq`	Qed	Transverse CS saving
7	`per_mode_res_bound`	Hypothesis	Combined per-mode bound
8	`shell_energy_def`	Hypothesis	Parseval on shells
9	`shell_size_bound`	Hypothesis	$\#S_j \leq k_j^2$
10	`per_shell_combined`	Qed	Inner ACL (k_j^5)
11	`k0_le_kj_sq`	Hypothesis	Poincaré
12	`ACL_shell_sq`	Qed	Full ACL (k_j^7)

Bottleneck: `per_mode_res_bound` (Component 7). This Hypothesis encodes the combined transverse + collinear mechanism from Lemma 4.24(c). If false, the entire ACL falls. The manuscript's analytical argument (§4.5) provides the justification; the Coq accepts it as input.

Single point of failure: `div_free` (Component 1). Without incompressibility, `factor_A` fails, `coupling_eq` fails, and the entire geometric mechanism collapses. This is precisely what the Tao test requires: the proof must depend on a structural property that the averaged system does not possess.

Counterfactual Break Analysis

If False	What Fails
`div_free`	`factor_A` fails → `coupling_eq` fails → no geometric saving → ACL fails entirely
`factor_A`	Coupling bound loses the k-reduction → exponent degrades
`per_mode_res_bound`	`per_shell_combined` unproved → ACL unproved
`shell_size_bound`	Outer summation uncontrolled → ACL exponent wrong
`shell_energy_def`	Parseval link broken → cannot convert mode sums to shell energy
CS axioms	`trans_sum_sq` fails → CS saving lost

Fragility assessment: The proof has one true single point of failure: `div_free`. Every other component either has independent justification (the Hypotheses encode standard mathematics) or is machine-verified (the Qed chain). The fragility is concentrated exactly where it should be: on the specific structural property of the Navier–Stokes equations that distinguishes them from generic bilinear operators.

Severity Summary

Class	Name	Count	Details
I	Fatal	0	—
II	Major	0	—
III	Technical	4	Factor 2 conservative; `shell_size_bound` axiomatised; `coll_count` axiomatised; first-harmonic not formalised
IV	Expository	1	Squared vs. square-root form
V	Positioning	0	—

All four Class III observations concern the boundary between the manuscript's analytical arguments and the Coq's formal verification — specifically, the counting Hypotheses and the `per_mode_res_bound` Hypothesis encode standard mathematics that is argued in the manuscript but not machine-verified. This is the documented trust boundary. The mathematical content of each Hypothesis is independently verified in the companion documents (*An Assessment of Shells*, *Permanent Axiom Trust Boundary Verification*).

Final Verdict

HOSTILE-REFEREE RESISTANT

All 16 attack routes executed. Zero Class I or Class II objections. Four Class III observations, all concerning the boundary between manuscript arguments and Coq formalisation — the documented trust boundary. One Class IV expository note.

The Angular Cancellation Lemma (Theorem 4.27) as stated and proved in `NavierStokesACL.v` is sound. The 12 Qed theorems are logically valid. The 3 CS axioms are elementary. The Section Hypotheses encode standard lattice geometry and Parseval's identity. The single structural input that makes the proof work — `div_free` (incompressibility) — is the specific Navier–Stokes structure that passes the Tao test.

Compelted

Auburn Patent Family Fields

UncleBroFields@proton.me | fieldsryanchristopher@gmail.com