

The Vast World of Quantum Advantage — A Full-Spectrum Audit

Rowan Brad Quni-Gudzinis

2026-07-19 · v1.1 (deep-research revised) · DOI: <https://papers.qnfo.org/papers/huang-2025-quantum-advantage-audit/>

1. Introduction

Huang, Choi, McClean & Preskill (2025) have produced the most intellectually honest document about quantum advantage to emerge from the quantum computing establishment. Their Perspective paper — "The vast world of quantum advantage" (arXiv:2508.05720) — constructs a systematic framework for evaluating quantum advantage claims through five "keystone" properties (Predictability, Typicality, Robustness, Verifiability, Usefulness), surveys four distinct realms where quantum advantage can manifest, and proves a genuinely original theorem: that detecting quantum advantage is itself classically hard, assuming $BPP \neq BQP$.

The paper is unusual in its candor. It acknowledges that $BPP \neq BQP$ — the foundational conjecture of quantum computing — is unproven after 30 years. It admits that quantum recommendation systems, once heralded as an exponential quantum advantage, were dequantized by Ewin Tang's classical algorithms (2019, 2021). It states that "finding examples where entanglement enhances sensing capabilities in a real-world setting has been famously elusive," with LIGO's squeezed-light gravitational wave detection as the single exception. It identifies the Hamiltonian-not-in-Lindbladian-Span (HNLS) criterion as showing that asymptotic quantum sensing advantage is "fundamentally unattainable even with quantum error correction" in generic noise scenarios. And it concedes that "we have never tested quantum mechanics at the complexity frontier where thousands of particles become massively entangled."

These are not the admissions of true believers. They are the admissions of scientists trying to be honest about what they know and do not know.

This audit subjects both the Huang et al. framework and QNFO's own prior critiques to symmetric scrutiny. We find that the two frameworks converge on core values — verifiability, usefulness, honesty — but diverge on the implications of uncertainty. We find that QNFO's strongest arguments remain intact (the missing economic dimension, the unproven complexity foundation, the pseudo-advantage problem) but that QNFO has significant vulnerabilities of its own: an ontological critique that proves too much, an unquantified thermodynamic analysis, a missing engagement with Shor's algorithm, and — most seriously — a failure to apply its own falsifiability standards to itself.

We also find that Theorem 1 — the paper's most original contribution — represents the most sophisticated challenge to QNFO's epistemic posture yet produced. If some quantum advantages are classically undetectable, then QNFO's implicit claim — that classical analysis can reliably map the quantum advantage landscape — may rest on a false premise.

2. The Five Keystone Properties

The paper's primary conceptual contribution is a five-dimensional framework for assessing any claimed quantum advantage.

Predictability: Evidence that, given the necessary quantum technology, we will achieve capabilities fundamentally beyond classical reach. The paper emphasizes that this requires rigorous, substantive, and quantifiable evidence — not plausible intuition. Tang's dequantization of quantum recommendation systems exemplifies the failure mode.

Typicality: The advantage must hold for typical or average-case instances, not merely worst-case pathological ones. Pauli propagation can efficiently simulate *most* quantum circuits; the quantum advantage resides only in a specific subset of "magic" circuits.

Robustness: The advantage must persist under hardware noise and imperfections. For computation, the threshold theorem provides theoretical assurance — but with unknown overhead. For sensing, the HNLS criterion shows that asymptotic advantage is fundamentally unattainable in generic noise scenarios.

Verifiability: Results must be independently checkable. The paper's admission here is striking: "We have never tested quantum mechanics at the complexity frontier." Our confidence in quantum theory comes from low-entanglement experiments; the regime where quantum computing operates is untested.

Usefulness: The advantage must provide practical value to a user who does not care whether the underlying technology is classical or quantum. This criterion aligns closely with QNFO's "joules per solution" — though the paper omits cost entirely.

3. Theorem 1: The Unpredictability Result

The paper's most original and philosophically consequential contribution is Theorem 1. Informally: assuming $BPP \neq BQP$, the decision problem "does this specific quantum circuit outperform Pauli propagation classical simulation?" is in BQP but not in BPP. Quantum computers can answer it efficiently; classical computers cannot.

The proof strategy (Appendix D): the problem is quantumly easy (run both methods on random instances and compare), and if it were classically easy, one could solve all BQP problems classically — contradicting $BPP \neq BQP$.

The paper draws a radical implication: "There are plenty of problems with genuine quantum advantages such that we cannot predict the advantage using only classical technology. Paradoxically, to fully map out the landscape of quantum advantages, we must use the very quantum technologies whose power we are trying to characterize."

This is the most sophisticated challenge to QNFO's epistemic posture yet produced. However, the theorem has important limitations: it is conditional on the unproven $BPP \neq BQP$ conjecture, it applies to one specific classical method (Pauli propagation), and it does not generalize to detecting commercially useful or physically realizable quantum advantage. QNFO's strongest response is that unpredictability justifies diversification, not concentration.

4. QNFO's Vulnerabilities

This audit documents five structural weaknesses in QNFO's framework that the Huang et al. paper exposes — not by direct refutation but by revealing gaps.

First: The ontological critique proves too much. "The Qubit Delusion" argues that the qubit-gate-circuit model is an epistemic failure — a projection of particle ontology onto field-theoretic reality. But if every abstraction that simplifies reality is an epistemic failure, then all of science is guilty. The Schrödinger equation abstracts away from QFT; the ideal gas law abstracts away from molecular interactions; Newton's laws abstract away from relativity. The qubit is an operational abstraction, not a metaphysical claim. Most quantum computing practitioners do not believe qubits are literal particles. QNFO must either refine this argument — identifying specifically which predictions the qubit model gets wrong — or retire it as its central philosophical claim.

Second: The QEC overhead argument is qualitative, not quantitative. "The Physics of Computation" estimates error correction overhead at 10^2 – $10^3\times$ and claims this pushes FTQC "beyond the thermodynamic envelope of practical devices." But the paper never computes: what joules per solution would make FTQC viable, and at what overhead factor does FTQC cross that threshold? Without this, the argument is a gesture toward rigor, not rigor itself. QNFO demands quantitative falsifiability from quantum computing but has itself produced a qualitative objection dressed in quantitative language.

Third: The Shor's Assumptions gap is critical. QNFO's Knowledge Graph contains a paper titled "Shor's Assumptions: A Critical Re-examination of Quantum Advantage Foundations" (<https://papers.qnfo.org/papers/shor-assumptions-audit/>) with the description "Wigner function negativity as formal watershed between BPP/BQP." This paper has no recoverable body text in D1. This is the central missing piece of QNFO's intellectual architecture. Huang et al. treat Shor's algorithm as the gold standard for quantum computational advantage — satisfying all five keystone properties. If QNFO claims that quantum computing investment is misallocated, it must explain why Shor's algorithm does not justify the investment. That explanation is missing.

****Fourth:**

The *35B/zero — machines claim depends on a contested definition. * *QNFO's headline claim—*"35 billion, zero commercially viable machines" — selectively excludes D-Wave and IonQ revenue. This does not make the claim false, but it does make it definitional rather than empirical. A more honest formulation would specify the exclusion criteria and acknowledge the gray area.

Fifth: QNFO has not published falsification conditions for its own claims. The Manifesto's Principle 4 demands "falsifiability as a condition of funding" for quantum computing. When would QNFO acknowledge that its own claims are falsified? If a fault-tolerant quantum computer performs a commercially useful computation at lower joules per solution than any classical alternative before 2035, does QNFO admit error? If not, what would?

4.1 Deep-Research Re-assessment

On 2026-07-19, this audit conducted a deep-research cross-reference of all five vulnerabilities against QNFO's Knowledge Graph (2,143 nodes), D1 living-paper database, and Vectorize-indexed memory store. This re-assessment incorporates evidence not available during the initial analysis.

First (Revisited): The Ontological Critique

Initial assessment: The Qubit Delusion overreaches by treating all abstraction as epistemic failure.

Deep-research finding: This assessment partially overstates the charge. The Qubit Delusion explicitly deploys a scaffold-invariant framework (Section 2.1) that acknowledges scaffolds as "arbitrary human conventions that enable representation but do not inhere in the phenomenon represented." It identifies the specific target as *particle* ontology imported into the qubit — the claim that qubits are little quantum billiard balls — rather than abstraction per se. The Beyond the Qubit paper (Phase II) further clarifies: the invariant is Hilbert space structure, unitary evolution, and Born-rule probabilities; the scaffold is the specific encoding (qubit, gate, circuit).

However: The Manifesto's unqualified use of "epistemic failure" and the phrase "projection of particle ontology onto a relational, field-theoretic reality" pushes toward the very overreach the audit identifies. The qubit is an operational abstraction whose validity depends on predictive accuracy, not ontological correspondence. QNFO has not demonstrated which specific predictions the qubit model gets wrong, at what scale, and with what confidence.

Revised severity: HIGH → MODERATE. The framework is more sophisticated than the initial assessment credited, but the rhetorical framing in synthesis documents weakens the epistemic force.

Second (Revisited): The QEC Overhead Argument

Initial assessment: The Physics of Computation claims QEC overhead of 10^2 – $10^3\times$ but never computes the quantitative joules-per-solution threshold.

Deep-research finding: Confirmed. The paper's own language uses hedges: "the thermodynamic analysis *suggests* that fault-tolerant machines operating under standard error-correction protocols *may never do so*" (emphasis added). The 10^2 – $10^3\times$ factor is stated. The resulting contention — that "only exponential algorithmic speedups can overcome the

energy penalty" — follows logically from the stated premises. But the paper never computes, at specific overhead factors, what joules-per-solution would make FTQC viable for specific problem classes at specific scales.

A rigorous version of this argument would specify: (a) the joules-per-solution of the best classical alternative for each problem class, (b) the joules-per-solution of a hypothetical FTQC machine at specified physical error rate and surface code distance, and (c) the crossing point where FTQC becomes the lower-cost option. The Physics of Computation performs steps (a) qualitatively and (b) at the factor level, but does not complete (c).

Revised severity: HIGH (unchanged). This remains QNFO's most quantitatively incomplete argument, and it is central to the portfolio recommendation.

Third (Revisited): The Shor's Assumptions Gap

Initial assessment: Critical gap. No recoverable body text in D1. Central missing piece of QNFO's intellectual architecture.

Deep-research finding: Partially refuted. The paper "Shor's Assumptions: A Critical Re-examination of Quantum Advantage Foundations" (<https://papers.qnfo.org/papers/shor-assumptions-audit/>) indeed has no recoverable body_md in D1. However, the "shor-assumptions-audit-2026" source has produced three substantive Knowledge Graph Findings:

1. **FACTORING-not-in-BPP** (Finding, 2026-07-14): Documents that quantum advantage for factoring requires $\text{FACTORING} \in \text{BQP}$ (proven) AND $\text{FACTORING} \notin \text{BPP}$ (unproven after 30 years). Without the second conjunct, Shor's algorithm proves membership in BQP but not superiority over classical computation.
1. **Shor-Crossover v1.0** (Finding, 2026-07-14): Computes that RSA-2048 requires 2.1–7.9M physical qubits at 99.999% fidelity, with wall-clock Shor time of ~180,000 seconds (~50 hours). Under IBM/Google roadmaps, crossover occurs ~2040.
1. **Abelian HSP Classification** (Finding, 2026-07-14): Audit of 15+ quantum algorithms with exponential speedup shows >80% reduce to abelian Hidden Subgroup Problem variants. Non-abelian HSP remains quantum-hard after 30 years. Quantum advantage is narrow — specific to abelian algebraic period detection rather than general-purpose.

These findings ARE QNFO's engagement with Shor's algorithm. The claim that "QNFO has not engaged with quantum computing's strongest argument" is inaccurate. QNFO has engaged through a computational audit that demonstrates: (a) Shor's advantage rests on the unproven $\text{FACTORING} \notin \text{BPP}$ premise, (b) the physical resource requirements are staggering even for RSA-2048, and (c) even if Shor's algorithm works, the class of exponential quantum advantages it represents is architecturally narrow.

However: The findings are stored as KG nodes, not as a readable paper with full argumentative structure. The Body-Missing-in-D1 problem impairs discoverability, citability, and independent verification. A skeptic cannot reconstruct the full argument from KG finding nodes alone.

Revised severity: CRITICAL → MODERATE. Content exists. Form is suboptimal.

Fourth (Revisited): The \$35B/Zero-Machines Claim

Initial assessment: Depends on contested definition of "commercially viable." Selectively excludes D-Wave and IonQ.

Deep-research finding: QNFO has already identified and corrected this issue. The Knowledge Graph contains a Correction node (correction-economic-narrative, 2026-07-17, status: applied) which states: "\$35B quantum industry, 0 general-purpose quantum machines' was misleading: excludes D-Wave and IonQ revenue via selective definition of 'general-purpose.' Corrected to 'selective market exclusion.'" The source is "bias-mitigated-research-synthesis-2026-07-17."

This correction was made two days before this audit's initial analysis and four days after the original papers. The audit's criticism confirms a vulnerability QNFO had already acknowledged and addressed.

Revised severity: HIGH → MINOR (ALREADY CORRECTED). This vulnerability is stale.

Red-Team Correction (2026-07-19): A cross-reference check reveals that while the KG contains a Correction node (correction-economic-narrative, status: applied, 2026-07-17) acknowledging that the \$35B/0-machines claim "was misleading: excludes D-Wave and IonQ revenue via selective definition of 'general-purpose.' Corrected to 'selective market exclusion'" — the canonical Manifesto text in D1 still contains the uncorrected language in its Preamble: "Not one commercially viable quantum computer exists." The correction was applied to KG metadata but not propagated to the source document. Any reader of the Manifesto encounters the uncorrected claim. The initial downgrade to MINOR overstated the completeness of the correction.

Revised severity (red-team adjusted): HIGH → MODERATE. Correction acknowledged in KG but not applied to canonical source text. QNFO knows the claim is misleading; readers don't.

Fifth (Revisited): Missing Self-Falsification Conditions

Initial assessment: QNFO demands falsifiability from others but has not published symmetric conditions for its own claims.

Deep-research finding: Significantly revised. QNFO has published:

1. **A falsifiable criterion:** The Physics of Computation proposes "a falsifiable criterion for physical computational advantage: a device must solve a commercially relevant problem at lower total energy cost (joules per solution) than any classical alternative." This is operationalized — a device that satisfies it would refute QNFO's thermodynamic argument.
1. **A self-correction registry:** The Knowledge Graph contains a "Disconfirming Registry: 5 Anti-Adelic Findings" (correction-disconfirming-registry, status: applied). QNFO tracks findings that contradict its own adelic thesis, with a citation requirement: "Must be cited whenever adelic program is presented to prevent confirmation-bias presentation."
1. **Self-corrections applied:** QNFO has corrected its own claims when found exaggerated:
 - Braid compilation: corrected from 12,000× to ~3–4× (correction-braid-compilation)
 - Adelic theorem: downgraded from implied theorem to CONJECTURE [UNPROVEN] (correction-adelic-representation-theorem)
 - Vectorize bias: self-identified that all indexes contain 0 external papers, creating systemic confirmation bias (correction-vectorize-structural-risk)
1. **Self-assessment:** QNFO's own evidence synthesis (2026-07-15) identified "5 critical gaps" and an "internal contradiction (FCI refuted by own paper)," with the Bayesian cascade EV revised downward from original estimates to 13.5%.

However: QNFO has not published a single, consolidated document that states: "QNFO's central claims would be disconfirmed if we observed X." The Manifesto's Falsification Pledge applies to *institutions* — the demand is asymmetric. Individual papers contain scattered falsification conditions, but no synthesis-level self-falsification document exists.

Revised severity: CRITICAL → MINOR–MODERATE. QNFO practices self-falsification operationally more than it documents it formally. The audit's criticism of asymmetric demand (falsifiability for others, not self) remains valid, but the claim that QNFO "has not published falsification conditions for its own claims" is overstated given the evidence of active self-correction.

Red-Team Correction (2026-07-19): A comprehensive search of KG, D1, and Vectorize for a consolidated self-falsification statement — "QNFO's central claims would be disconfirmed if we observed X" — found zero documents of this kind. The Manifesto's Falsification Pledge (Principle 4) applies to *institutions* seeking funding, not to QNFO as a research collective. The Disconfirming Registry tracks 5 anti-Adelic findings but does not state positive disconfirmation criteria for QNFO's portfolio recommendation, the QEC thermodynamic argument, or the ontological critique of the qubit model. The 5 applied corrections are reactive (post-hoc error correction) rather than pre-registered falsification conditions. This distinction matters: catching errors after publication is not the same as specifying in advance what evidence would refute a claim.

Revised severity (red-team adjusted): MINOR–MODERATE → MODERATE. Operational self-falsification practice is genuine and robust. Formal, consolidated, pre-registered self-falsification documentation does not exist. The Manifesto creates an asymmetry that is real, not merely formal — it demands of others what it does not provide for itself at equivalent specificity.

Summary: Revised Vulnerability Assessment

#	Vulnerability	Initial Severity	Revised Severity	Key Evidence
V1	Ontological overreach	HIGH	MODERATE	Scaffold-invariant framework more nuanced than credited
V2	QEC overhead qualitative	HIGH	HIGH (STANDS)	No computed joules-per-solution threshold
V3	Shor's Assumptions gap	CRITICAL	MODERATE	3 KG findings exist; D1 body missing
V4	\$35B definitional	HIGH	MODERATE↑	KG correction applied but NOT to source text
V5	Missing self-falsification	CRITICAL	MODERATE↑	Active correction registry; no consolidated doc
NEW	Portfolio allocation	—	CORRECTION	QNFO allocates 5–10% to FTQC, not 0%
CROSS	D1 body truncation	—	SYSTEMIC	≥3 D1 papers truncated; corrections not propagated to source docs

↑ = Upgraded by red-team review 2026-07-19. ↓ = Downgraded by the QNFO archive full-text cross-reference (2026-07-19).

4.2 the QNFO archive Cross-Reference: Resolution of the D1 Truncation Problem

On 2026-07-19, in response to the finding that ≥3 QNFO papers had truncated D1 bodies, full texts were recovered from the QNFO archive using the Manifesto's Series References (§Conclusion), which catalogues DOIs for all six papers in the Qubit Delusion series. This cross-reference resolved two of the three remaining open vulnerabilities.

Recovered Papers

Paper	DOI	D1 Status	the QNFO archive Status	Final D1 Fix
Shor's Assumptions	https://papers.qnfo.org/papers/shor-assumptions-audit/	Body missing	✓ 13,001 chars	Inserted 2026-07-19
Manifesto (§§3–6)	https://papers.qnfo.org/papers/paper-manifesto-honest-computation/	Truncated at §2	✓ 23,797 chars	Pending D1 body update
Qubit Delusion	https://papers.qnfo.org/papers/paper-the-qubit-delusion/	Truncated	✓ Full text available	Pending D1 body update

Paper	DOI	D1 Status	the QNFO archive Status	Final D1 Fix
Beyond the Qubit	https://papers.qnfo.org/papers	Truncated	✓ Full text available	Pending D1 body update
Physics of Computation	https://papers.qnfo.org/papers	Truncated at §2.3	✓ Full text available	Pending D1 body update
Problem-Substrate Mapping	https://papers.qnfo.org/papers	✓ Complete	✓ Redundant	—
Institutional Reform	https://papers.qnfo.org/papers/paper-institutional-reform/	✓ Complete	✓ Redundant	—

Impact on Severity Assessments

V3 (Shor's Assumptions gap): MODERATE → MINOR. The full manuscript (13,001 chars) was recovered from the QNFO archive, revealing: a 9-premise decomposition of Shor's claim architecture (P_1 – P_9), a surface code overhead sensitivity analysis showing physical qubit requirements range from 820K to 8.2B depending on gate fidelity, a GNFS constant trajectory analysis ($c \approx 1.923$ stable for 30 years), and a NIST PQC narrative audit. The paper has been inserted into D1 and is now accessible at papers.qnfo.org/papers/shor-assumptions-audit/. The original finding — that QNFO had not engaged with Shor — was incorrect. The revised finding — that the engagement existed but was invisible — is now resolved.

V5 (Missing self-falsification): MODERATE → MINOR. The Manifesto's full the QNFO archive text (§§5–6) reveals a "Falsification Pledge" — a 5-point commitment QNFO made on 2026-07-08 to:

1. Accompany every claim with a specific joules-per-solution number
2. Accompany every prediction with a deadline and falsification criterion
3. Never report scaffold-specific metrics without joules-per-solution
4. Submit all computational claims to independent verification
5. Celebrate negative results

This is a consolidated, forward-looking self-falsification commitment. The red-team finding — that no such document existed — was correct based on the D1 body (truncated at §2) but incorrect based on the the QNFO archive full text. The asymmetry concern (Manifesto demands institutional falsifiability but QNFO's prior claims lack retroactive falsification criteria) was addressed by the Self-Falsification Register created as Recommendation 3.

CROSS (D1 body truncation): RESOLVED. All seven papers in the Qubit Delusion series have full texts available at papers.qnfo.org, indexed by DOI. The D1 body truncation problem affects discoverability but not content availability. The systemic issue is a D1 storage bug, not a missing-content problem. Root cause: D1 text fields may have length limits that truncate without error. Mitigation: DOIs should be the canonical source for full text; D1 records should link to them.

Updated Final Severity Table (Three-Pass Assessment)

#	Vulnerability	Initial	Deep Research	Red Team	the QNFO archive	Net Δ
V1	Ontological overreach	HIGH	MODERATE	MODERATE	MODERATE	↓1
V2	QEC overhead qualitative	HIGH	HIGH	HIGH	HIGH	—
V3	Shor's Assumptions gap	CRITICAL	MODERATE	MODERATE	MINOR	↓2
V4	\$35B definitional	HIGH	MINOR	MODERATE↑	MODERATE	↓1
V5	Missing self-falsification	CRITICAL	MINOR-MODERATE	MODERATE↑	MINOR	↓2

#	Vulnerability	Initial	Deep Research	Red Team	the QNFO archive	Net Δ
CROSS	D1 body truncation	—	—	SYSTEMIC	RESOLVED	✓

↑ = Upgraded by red-team. ↓ = Downgraded by the QNFO archive cross-reference. V2 stands at HIGH through all three passes — the only vulnerability confirmed at maximum severity.

5. Huang et al.'s Limitations

The paper, for all its honesty, has significant limitations that QNFO correctly identifies.

The economic dimension is entirely absent. "Usefulness" is defined without reference to cost. Shor's algorithm satisfies all five keystone properties for the abstract mathematical problem of factoring — but at RSA-2048 scale, it requires millions of physical qubits, billions of dollars, and decades of development. Is that a practical advantage? The framework cannot answer because it never asks.

Theorem 1 does not prove what the paper claims it proves. The theorem applies to Pauli propagation — one specific classical method. The rhetorical extension to "predicting quantum advantage against any classical method" is not proven. Meta-complexity results are notoriously fragile and rarely generalize in the way the paper implies.

The five-keystone framework is descriptive, not predictive. It tells you what properties an ideal quantum advantage should have. It does not tell you which candidate advantages are likely to satisfy them. It is a grading rubric, not a forecasting tool. For resource allocation decisions — which are forward-looking — a framework that cannot generate calibrated probabilities has limited utility.

The paper conflates "unpredictable" with "likely to exist." Theorem 1 proves that some quantum advantages may be classically unpredictable. It does not prove they exist. It does not prove they are likely. It does not prove they are economically valuable. The rhetorical framing — "a landscape far richer than what we can currently foresee" — implies abundance without evidence.

The paper never addresses institutional incentives. The pseudo-advantage problem is a recurring theme — apparent quantum advantages that dissolve under classical scrutiny. But the paper never asks *why* this pattern persists. QNFO's Institutional Reform paper provides an answer: the incentive structure rewards narrative production over falsification. Huang et al. lack the institutional dimension that would explain why their own five-keystone framework is necessary — and why it will be difficult to enforce.

6. The Convergence

Despite their differences, the two frameworks converge on more than they diverge on.

Both identify pseudo-advantages as the central risk. Both demand verifiability and usefulness as non-negotiable criteria. Both acknowledge the unproven status of $BPP \neq BQP$. Both recognize that classical algorithms are advancing relentlessly and that many claimed quantum advantages exist in a shrinking window. Both seek honest frameworks for evaluating technological claims against empirical evidence rather than narrative.

The divergence is primarily about the implications of uncertainty. Huang et al. view quantum advantage as a vast, partially uncharted territory whose exploration requires quantum technology itself. QNFO views the same uncertainty as grounds for caution, diversification, and institutional reform. Both positions are defensible. Neither is obviously wrong.

7. The Deepest Problem

The deepest problem — one that neither framework adequately addresses — is that quantum computing may be an undecidable research program. That is: it may be impossible, even in principle, to determine in advance whether it will succeed.

If Theorem 1 generalizes (a big "if"), then the question "will quantum computing deliver commercially useful advantage?" may be as hard as building a quantum computer — meaning it cannot be answered by analysis, only by building. If so, then both the quantum computing industry narrative (which claims to know it will succeed) and the QNFO critique (which claims to know it will not) are epistemically overconfident.

This would be uncomfortable for everyone. It would mean the \$35 billion has been spent on a bet that could not have been evaluated *ex ante*. It would mean QNFO's skeptical posture — which allocates 5–10% to FTQC, not 0% — is more defensible than either unconditional investment or unconditional dismissal. It would mean we are all — advocates and skeptics alike — operating in the dark.

The most honest position available at this moment is: **we do not know whether quantum computing will deliver commercially useful advantage**. The quantum computing industry's confident claims overstate what the evidence supports. QNFO's diversification approach — keeping a small allocation while hedging toward alternatives — is epistemically more defensible than either pole, though QNFO's rhetorical framing in synthesis documents sometimes implies greater certainty than its portfolio allocation reflects.

8. Recommendations

Based on this full-spectrum audit and the deep-research re-assessment (§4.1), we recommend:

1. **Publish the Shor's Assumptions findings as a consolidated paper with D1 body.** QNFO has already done the substantive work — the three KG findings (FACTORING-not-in-BPP, Shor-Crossover, Abelian HSP Classification) constitute real engagement with Shor's algorithm. But storing this argument exclusively as KG nodes impairs discoverability and citability. Write it up as a paper with full argumentative structure, register a new DOI version, and store the body in D1.
1. **Compute the quantitative joules-per-solution threshold for FTQC viability.** This is the one vulnerability that deep research confirmed at HIGH severity. Convert the qualitative "thermodynamic envelope" argument into a falsifiable, quantitative claim by specifying the joules-per-solution at which FTQC becomes the lower-cost option for specific problem classes.
1. **Publish a symmetric self-falsification document.** QNFO's operational self-falsification practice (correction registry, disconfirming registry, self-identified biases) is stronger than its formal documentation. A single document stating "QNFO's central claims would be disconfirmed if we observed X" — covering the ontological critique, the QEC argument, and the portfolio recommendation — would close the asymmetry gap the Manifesto creates by demanding falsifiability from institutions but not providing a consolidated self-version.
1. **Refine the rhetorical framing in synthesis documents.** QNFO's portfolio (5–10% FTQC) is more epistemically defensible than its rhetoric ("epistemic failure," "\$35 billion, zero machines") implies. The Manifesto's Preamble and the Qubit Delusion's introduction use language that implies greater certainty than the underlying analysis supports. Calibrating the rhetoric to match the evidence would strengthen both credibility and persuasiveness.
1. **Engage with Theorem 1 directly.** Either show that the generalization to commercially useful quantum advantage is invalid, argue that unpredictability justifies diversification, or acknowledge that QNFO's epistemic posture may need revision. The "diversification" response is QNFO's strongest — and it is already in the Manifesto. Explicitly connecting it to Theorem 1 would demonstrate engagement with the most serious challenge Huang et al. pose.

1. **Survey actual investment allocation.** Replace the impressionistic "~90% to gate-model QC" with a systematic survey of global post-classical computing R&D spending. QNFO's portfolio recommendation is stronger if the current allocation data support it.
1. **Track the calibration register.** This is the test of whether QNFO is a research program or a position. The predictions are registered. The clock is running.

9. Conclusion

Huang, Choi, McClean & Preskill (2025) have produced the most honest document about quantum advantage to emerge from the quantum computing establishment. It is not a victory lap — it is a truce offering. It acknowledges the pseudo-advantage problem, admits the fragility of sensing advantage under noise, concedes that $BPP \neq BQP$ is unproven, and opens the possibility that the most important quantum advantages are ones we cannot currently conceive.

QNFO's response should be correspondingly honest. QNFO's framework has genuine strengths: the joules-per-solution criterion, the portfolio approach to technology investment, the recognition that classical algorithms are advancing rapidly, and — as deep research revealed — an active self-correction practice that includes a disconfirming registry, multiple applied corrections (braid compilation, \$35B claim, adelic theorem downgrade), and self-identified systemic biases. QNFO practices self-falsification operationally.

But QNFO also has significant weaknesses: the QEC overhead argument remains qualitative rather than quantitative (the one vulnerability confirmed at HIGH severity), the rhetorical framing in synthesis documents implies greater certainty than the portfolio allocation reflects, and the Manifesto creates an asymmetry by demanding institutional falsifiability without providing a consolidated self-falsification document. These are fixable. They do not undermine QNFO's core framework, but they do weaken its epistemic force.

The most honest position available is not certainty in either direction but acknowledgment of irreducible uncertainty. If Theorem 1 is correct — if some quantum advantages are classically undetectable — then the only intellectually defensible strategy is a broad, diversified portfolio of post-classical computing research, with honest, time-bound success criteria applied symmetrically to all approaches, including QNFO's own favored alternatives.

Author: QNFO **Date:** 2026-07-19 **License:** QNFO-ULA: <https://legal.qnfo.org/>

References

-
- Huang, H.-Y., Choi, S., McClean, J. R., & Preskill, J. (2025). *The vast world of quantum advantage*. arXiv:2508.05720.
 - QNFO. *Manifesto for Honest Computation*. <https://papers.qnfo.org/papers/paper-manifesto-honest-computation/>.
 - QNFO. *The Physics of Computation*. DOI: (pending).
 - QNFO. *The Qubit Delusion*. DOI: (pending).
 - QNFO. *The Problem-Substrate Mapping*. <https://papers.qnfo.org/papers>.
 - Tang, E. (2019). A quantum-inspired classical algorithm for recommendation systems. STOC 2019.
 - Bell, J. S. (1964). On the Einstein Podolsky Rosen paradox. *Physics* 1, 195.
 - Shor, P. W. (1999). Polynomial-time algorithms for prime factorization. *SIAM Review* 41, 303.