The Scope of Quantum Computing: A Clarification

Why Most "Quantum" Problems Don't Require Quantum Mechanics

Quantum computing is to computation what particle accelerators are to physics: essential for a narrow class of problems, irrelevant for most.

Plain Language Summary

When Is Quantum Computing Actually Necessary?

Quantum computing is often presented as a solution to problems that are "hard" for classical computers—problems with enormous search spaces, complex constraints, or exponentially many possibilities. From this perspective, it is natural to assume that many such problems must therefore require quantum mechanics.

Our results suggest that this intuition, while understandable, is incomplete.

To clarify when quantum computing is genuinely necessary, it is helpful to separate three distinct stages that are frequently conflated in discussions of quantum advantage:

- **Marking** — creating any reliable signal (a tag, score, constraint, or evidence) that distinguishes correct possibilities from incorrect ones[1]
- **Amplification** — concentrating probability or weight onto those marked possibilities (can be achieved via dynamics that preserve or exploit the asymmetry)
- **Commitment** — producing a definite outcome or answer (measurement, inherently irreversible)

[1] *In technical terms, a "mark" is any physically instantiated asymmetry correlated with correctness—often created by measurement, constraint enforcement, or energy biasing—and typically involves irreversible record formation.*

This decomposition is not arbitrary; it corresponds to three physically distinct operations, each with different resource requirements.

The central conclusion of this work is that **only the first of these stages sometimes requires quantum mechanics.**

Marking vs Amplification

Once a possibility has been physically marked—even weakly—the remaining steps do not intrinsically require quantum effects. Probability mass can be concentrated onto marked regions using entirely classical mechanisms such as dissipation, feedback, structured propagation of information, or constraint enforcement. A final answer can then be obtained by an irreversible commitment to one outcome.

In other words:

Quantum mechanics is not required to amplify answers or to read them out.
It is required only when the act of marking correctness is itself quantum.

This distinction resolves much of the confusion surrounding claims of quantum computational advantage.

Problems That Do Not Require Quantum Computing

Many problems that are widely described as "quantum computing problems" do not, in fact, require quantum mechanics in a foundational sense.

These include:

- **Constraint satisfaction problems**, such as scheduling, routing, and many SAT-like problems, where correctness is defined by classical rules
- **Error correction and diagnosis**, where measurements or syndromes already provide physical marking
- **Bayesian inference and probabilistic reasoning**, where evidence naturally marks hypotheses
- **Optimization problems**, where energy, cost, or loss functions provide classical marking signals

In all of these cases, the notion of correctness can be expressed and physically marked using classical processes. While the search space may be very large, the difficulty does not lie in amplifying or selecting a marked answer once it exists. Instead, it lies in how efficiently marking signals are produced or propagated—something that can often be done classically when the problem has structure.

The size of the search space alone does not determine whether quantum computing is necessary.

Problems That May Benefit from Quantum Computing

There exists an intermediate class of problems where quantum computing may offer advantages, even though it is not strictly required.

Examples include:

- Certain **cryptographic or algebraic problems** where relevant structure is difficult to extract classically
- Tasks where **evaluating correctness efficiently requires simulating quantum evolution**
- Situations where **interference or phase information** assists in constructing marking signals

In these cases, quantum mechanics can help create the mark more efficiently. However, once a mark exists, amplification and answer recovery still do not intrinsically depend on quantum hardware.

Problems That Truly Require Quantum Computing

A smaller class of problems genuinely requires quantum computing because correctness itself is quantum-defined.

These include:

- **Simulating quantum systems** whose properties have no classical description
- **Determining quantum features** such as entanglement or phase relationships
- Tasks where the **"answer" cannot be expressed as a classical predicate**

Here, quantum computing is not merely advantageous—it is essential. No classical marking mechanism exists for these problems.

A Practical Rule of Thumb

A useful way to assess whether a problem truly requires quantum computing is to ask:

Can the notion of correctness be physically marked using classical processes such as constraints, energy, measurements, or evidence?

- If **yes**, then quantum computing is not intrinsically required.
- If **no**, and correctness itself depends on quantum properties, then quantum computing is essential.

Implications

When this distinction is applied across the range of problems currently described as quantum computing applications, it suggests that a minority—plausibly on the order of 5–15% under the admissibility criterion used here—require quantum mechanics in a foundational sense.[1] The majority either do not require quantum computing at all or may benefit from it only as an optional or auxiliary tool.

This clarification does not weaken the case for quantum computing. On the contrary, it strengthens it by placing quantum advantage on clear physical ground. Quantum computers are not general-purpose replacements for classical computers; they are specialized tools for problems where the definition of correctness is inherently quantum.

Understanding this boundary helps guide research, investment, and application development toward domains where quantum computing is genuinely indispensable.

Abstract

When can a single measurement recover a fact hidden among exponentially many possibilities? We present a unified framework based on information geometry that answers this question for both classical and quantum systems.

The core mechanism is that **marking** a correct answer induces geometric asymmetry in probability space, and **natural gradient flow** under the Fisher–Rao metric concentrates probability mass into the marked region. This flow is realized classically via belief propagation on structured constraint graphs, and quantum-mechanically via Lindblad dynamics with engineered dissipation.

We demonstrate three key results:

1. **Classical sufficiency**: When problem structure compresses the hypothesis space (low treewidth, sparse constraints), single-shot recovery succeeds *without quantum resources*. We validate this on a 200-bit system (10^{60} nominal possibilities) and show that LDPC decoding—deployed in every 5G phone—is an instance of this mechanism.
2. **Structural boundary**: Single-shot recovery fails when treewidth is high or constraints are dense. The framework precisely characterizes *when* classical methods break down.
3. **Quantum regime**: For unstructured problems where classical propagation fails, we derive a Lindblad-based amplifier achieving O(1) resource scaling versus Grover's $O(\sqrt{N})$. The open question is whether this can be realized with polynomial physical resources.

When we separate marking from amplification, we find that a minority—plausibly on the order of 5–15% under the admissibility criterion used here—of problems commonly labeled as "quantum" actually require quantum mechanics in a foundational sense.[1] In most cases, the difficulty lies in structure or optimization, not in quantum physics itself. The contribution is not a quantum speedup claim, but a **decision framework**: given a problem's structure, does it require quantum resources or not?

Contents

Plain Language Summary 1

When Is Quantum Computing Actually Necessary? 1

Marking vs Amplification 2

Problems That Do Not Require Quantum Computing 2

Problems That May Benefit from Quantum Computing 3

Problems That Truly Require Quantum Computing 3

A Practical Rule of Thumb 3

Implications 4

Abstract 4

1. Introduction: The Real Question 12

1.1 The Mechanism in One Sentence 12

1.2 Scope and Limitations 13

1.2.1 A Taxonomy of Difficulty: Marking Cost vs Structural Propagation 13

1.2.2 A Third Filter: Taylor Admissibility and Finite Distinguishability 14

1.3 Document Structure 16

1.4 The Central Theorem (Full Formalization) 16

1.4.1 Formal Definitions 17

1.4.2 The Marking Sufficiency Theorem 18

1.4.3 Complexity Class Connections 19

1.4.4 The Marking Hierarchy 20

2. Probability Geometry and the Fisher–Rao Metric 21

3. Natural Gradient Flow on the Simplex 21
4. Stage A: Amplification of the Marked Set 22
5. Stage B: Tie-Breaking Within the Marked Set 22
5.1 Conditional Dynamics Inside the Marked Set 23
6. Lyapunov Function and Convergence Guarantee 23
7. Classical Realization: Factor Graphs and Belief Propagation 24
7.1 Factor-Graph Representation 24
7.2 XOR (Parity) Chain Constraints 24
7.3 Marking as Local Probabilistic Evidence 24
7.4 Exact Belief Propagation on Trees 25
7.5 Single-Shot Readout and Global Recovery 25
7.6 Structural Requirements for Single-Shot Recovery 26
8. Worked Example: Four-Variable XOR Chain 26
8.1 Setup 26
8.2 Admissible Global Assignments 26
8.3 Marking via Local Evidence 26
8.4 Posterior Calculation 27
8.5 Propagated Marginals 27
8.6 Single-Shot Readout 27
9. Stress Test: 200-Bit System (10^{60} Nominal Possibilities) 27
9.1 Problem Setup 28
9.2 Marking and Propagation 28
9.3 Empirical Results 28
9.4 What This Shows 28
9.5 What This Does Not Show 29
9.6 Supplementary Bounded-Treewidth Experiments (w = 2–4) 29
10. Real-World Application: LDPC Decoding 30
10.1 The Problem: Noisy Channel Decoding 30
10.2 Mapping to the Framework 30

10.3 Runnable Implementation 31
10.4 Expected Behavior 33
10.5 Why This Matters 34
11. Empirical Validation Across Parameters 34
11.1 Test Protocol 34
11.2 Results 34
11.3 Interpretation 35
12. The Structural Boundary: When Classical Fails 35
12.1 The Treewidth Criterion 35
12.2 Examples of Structural Failure 35
12.3 What Quantum Might Provide 35
12.3.1 A Crucial Non-Implication 36
12.4 The Honest Tradeoff 36
12.5 A Practical Classification Heuristic 36
12.6 Common Demo Types That Are Not QC-Required 37
12.7 What Is in the QC-Required Bucket 38
12.8 Extended Classification: Ten Common "Quantum" Applications 39
13. Case Study: MaxCut—A Canonical QC Benchmark That Doesn't Need QC 40
13.1 What MaxCut Actually Is 40
13.2 Why QC Companies Use It 40
13.3 Applying the Framework 41
13.4 Classical Solution: Same Problem, No Quantum Computer 41
13.5 Results 43
13.6 What This Demonstrates 43
13.7 The Honest Interpretation 44
13.8 Summary 44
13.9 Case Study: Protein Design (10^130 Possibilities, Probability Geometry in Action) 44
13.9.1 The Problem 44
13.9.2 The Probability Simplex Over Sequence Space 45

13.9.3 Marking via Energy Functions 45
13.9.4 The Factor-Graph Structure 45
13.9.5 Belief Propagation Demonstration 45
13.9.6 Interpretation Through Probability Geometry 47
13.9.7 Why Quantum Computing Doesn't Help 48
13.9.8 Summary 48
14. Case Study: Google's Willow Chip—A Genuinely Quantum-Required Task 48
14.1 What Willow Actually Demonstrated 48
14.2 Applying the Framework 49
14.3 Why the "10 Septillion Years" Comparison Is Accurate but Narrow 49
14.4 What Willow Does and Does Not Demonstrate 50
14.5 How Willow Validates the Framework 50
14.6 The Clarification This Enables 50
15. Quantum Implementation: Lindblad Dynamics 51
15.1 Single-Shot Criterion 51
15.2 Marked-State Pump Channel 51
15.3 Population Dynamics 52
16. Closed-Form Single-Shot Success Law 52
17. Comparison with Grover Search 52
17.1 Relation to Query Complexity 53
17.2 What Is New Relative to Known Dissipative Search 53
18. Binary Special Case and EF Connection 54
19. Experimental Protocol 55
19.1 Hardware Requirements 55
19.2 Measurements 55
19.3 Candidate Platforms 55
20. Compressed Ancilla-Mediated Amplifier 55
20.1 Compressed Pump Construction 55
20.2 Effective Two-Dimensional Dynamics 56

20.3 Scaling Result 56
20.4 Interpretation and Caveats 56
20.5 Realizability Roadmap 57
20.6 Positive Result: The Symmetric Subspace Theorem 58
20.7 Analysis: The Coherence Requirement and Realizability Landscape 59
20.8 The Realizability Landscape: Structure-Dependent Classification 60
20.9 Unification: Classical and Quantum Perspectives Converge 61
21. An Algebra of Physical Distinguishability and Irreversible Inference 62
21.1 Motivation 62
21.2 Physical Distinguishability as an Equivalence Relation 63
21.3 Algebraic Operations 63
21.3.1 Multiplicative Combination (Constraint Intersection) 63
21.3.2 Convex Mixing (Coarse-Graining) 64
21.4 Facts as Absorbing Idempotents 64
21.5 Dynamics as Algebra Endomorphisms 65
21.6 Relation to Classical and Quantum Formalisms 65
21.7 Why This Algebra Matters 66
21.8 Summary 66
22. Discussion: The Proper Role of Quantum Computing 67
22.1 What This Framework Does Not Say 67
22.2 The Cost and Complexity Make Sense Once Scope Is Correct 67
22.3 The Reframing That Resolves the Tension 68
22.4 Example: Drug Discovery 68
22.5 Example: Cryptography (Genuinely QC-Required) 69
22.6 Why Narrow Scope Strengthens Rather Than Weakens QC 71
22.7 The Real Problem Was Expectation Mismatch 71
22.8 Relationship to Adiabatic Quantum Computing 71
22.9 Addressing the Strongest Counterarguments 72
22.10 Historical Precedent: The Pattern of Scope Clarification 74

23. Falsification Criteria and Outlook 76
23.1 Sharp Falsification Test 76
23.1.1 Experimental Falsification Protocol 76
23.2 What Success Would Mean 77
23.3 Summary of Key Results 77
24. Conclusion: A Decision Framework 77
24.1 The Core Mechanism 78
24.2 The Decision Tree 78
24.3 What We Proved 78
24.4 What Remains Open 78
24.5 The Takeaway 79
25. Why This Matters: The Stakes of Scope Clarification 80
25.1 For Investors and Policymakers 80
25.2 For Researchers 80
25.3 For Engineers 81
25.4 For the Field's Credibility 81
25.5 For Scientific Understanding 81
25.6 A Note on Intellectual Honesty 81
Notes 82
References 83
Appendix A: Physical Admissibility and the Collapse of Quantum Necessity 84
A.1 Recap: The Physical Admissibility Framework 84
A.2 Where Quantum Hardness Typically Lives 85
A.3 The Taylor Limit on Hilbert Space 85
A.4 The "Patternlessness" Criterion 85
A.5 Successive Collapse of the Quantum-Necessary Class 86
A.6 Revised Estimate 86
A.7 Summary 87
Appendix B — Interpreting Figure B: Marking, Amplification, and Commitment 87

Marking (left panel) 89
Amplification (middle panels) 89
Commitment (right panel) 89
On “Quantum-Required” Problems (bottom panel) 89
Key takeaway 89
Appendix C — Scope, Limits, and Falsifiability of the Amplification Framework 90
C.1 What the Lindblad Amplifier Analysis Does — and Does Not — Claim 90
C.2 Why the Fisher–Rao Geometry Is Not Redundant with Treewidth 91
C.3 What Would Falsify the 5–15% Estimate 91
Falsifier 1 — Existence of a Counterexample Domain 91
Falsifier 2 — Annotation Instability 92
Falsifier 3 — Provable Quantum Advantage on Classically Markable Objectives 92
C.4 Summary 92
Appendix D — Physical Realizability of Lindblad Amplification 92
D.1 Model and Assumptions 93
D.2 No-Go Theorem for Unstructured Marks 93
D.3 Constructive Realizability for Structured Marks 94
D.4 Resolution of the Lindblad Realizability Question 94
D.5 Implication for the Framework 94

1. Introduction: The Real Question

The question "do we need a quantum computer for this?" is usually framed in terms of computational complexity: P vs BQP, oracle separations, query lower bounds. This framing, while rigorous, often obscures the practical question engineers actually face:

Given the structure of my problem, can I solve it classically, or do I need quantum resources?

This paper provides a geometric answer. The key insight is:

Single-shot fact recovery succeeds when probability mass can be concentrated via structured propagation. The question is whether that structure exists classically or must be supplied quantum-mechanically.

Several components of the present framework resemble ideas that appear separately in existing literatures: oracles in query complexity implicitly "mark" solutions, local evidence propagates in constraint satisfaction problems, and engineered dissipation can pump probability mass in open quantum systems. What is new here is not the introduction of another amplification mechanism, but the recognition that these disparate constructions are instances of a single physical process governed by information geometry. Query complexity treats marking as an abstract oracle; constraint satisfaction treats propagation as a combinatorial procedure; dissipative quantum computation treats pumping channels as given generators. By instead defining marking as a physically instantiated asymmetry correlated with correctness, and analyzing amplification as natural-gradient flow under the Fisher–Rao metric, we identify a substrate-independent mechanism that applies equally to classical and quantum inference. This reframing shifts the classical–quantum boundary: quantum computing is not required to amplify or commit to marked solutions, but only in those cases where no classical physical marking process exists at all. The resulting criterion is neither a restatement of treewidth conditions nor a reinterpretation of oracle models; it yields a new, falsifiable classification of when quantum resources are foundationally necessary.

1.1 The Mechanism in One Sentence

Marking creates local asymmetry. **Structure** (constraints, geometry) propagates that asymmetry globally. **Readout** extracts the answer in one shot.

- If structure exists classically (sparse constraints, low treewidth) → **classical BP suffices**
- If structure is absent classically but can be created via coherence/entanglement → **quantum helps**
- If no structure exists at all → **nothing helps; you must search**

1.2 Scope and Limitations

This framework does not address the cost of identifying or computing the mark. It addresses the dynamics of probability concentration *after* a physical marking exists. This is analogous to:

- **Post-selection amplification**: Given that an event occurred, amplify its signal
- **Error-syndrome processing**: Given syndrome bits, decode the logical state
- **Constraint propagation**: Given local evidence, infer global structure

The relevant question is not "can we find the needle?" but rather: *given that the needle has been marked, what is the minimal cost to extract it?*

1.2.1 A Taxonomy of Difficulty: Marking Cost vs Structural Propagation

The statement "this framework does not address the cost of computing the mark" is not a disclaimer of weakness; it is a deliberate separation of concerns. Many debates about quantum advantage implicitly merge three distinct costs into one:

1. The cost of **creating** a physical mark
2. The cost of **propagating** that mark into global concentration
3. The cost of **committing** to a fact

This paper isolates (2) and (3), and treats (1) as problem-dependent.

A useful classification treats two axes as independent:

- **Marking cost**: how expensive is it to produce a physically instantiated asymmetry correlated with correctness?
- **Structural propagatability**: does the problem admit low-complexity global propagation of local marks (e.g., low treewidth, sparse factorization, symmetry reduction)?

This yields three regimes:

Regime	Marking	Structure	Outcome
Classical wins	Cheap	Present	Local marks propagate efficiently (XOR chains, LDPC, tree inference). Single-shot recovery is classical.
Quantum might help	Cheap	Absent	Mark exists but classical propagation blocked by high treewidth. Engineered quantum dynamics may concentrate weight—but physical realizability is open.
Outside scope	Expensive (dominant)	Any	Marking cost dominates; framework does not claim advantage.

Clarification: "Outside scope" refers to cases where the cost of producing the mark dominates the total computation. If marking is expensive but structure is present, the overall task may still be tractable—one pays the marking cost once and then propagates efficiently. The excluded regime is specifically: *marking is as hard as solving the problem and no structural leverage is available to reduce that cost.*

This taxonomy preempts the "relocation" objection by making it explicit: **the framework is not a universal solver; it is a decision framework that tells you whether the remaining task—fact recovery after marking—is classical or potentially quantum.**

Note on the continuum: In practice, both axes admit gradations. Marking may be partially informative (weak likelihood ratios), and structure may be approximate (nearly tree-like graphs where loopy BP converges). The framework's predictions degrade gracefully: weaker marks require longer propagation or more measurements; approximate structure yields approximate single-shot recovery. The binary presentation is for clarity; the underlying mathematics is continuous.

1.2.2 A Third Filter: Taylor Admissibility and Finite Distinguishability

The taxonomy above separates difficulty into (i) the cost of producing a mark and (ii) the availability of structure that propagates marks into global concentration. A further refinement becomes available if physical reality imposes a hard bound on operational distinguishability—i.e., a maximum number of physically distinguishable states within a causal patch, together with a minimum meaningful resolution for measurement and computation. The BCB framework proposes a dimensionless bound L_T (Taylor's Number) that limits the total number of distinguishable states accessible to any observer within a single causal horizon, implying that attempts to resolve distinctions beyond this bound become physically meaningless.

We call this additional constraint the **Taylor admissibility filter**:

Taylor admissibility (conditional). A computation is Taylor-admissible if it does not require more than L_T mutually distinguishable states (or, equivalently, does not require physical precision finer than the implied minimum meaningful resolution). If an operation requires distinctions beyond this bound, the operation may remain mathematically definable but loses physical meaning as an executable process.

Why this extends the category of solvable physical problems

Complexity theory classifies difficulty in abstract models that permit unlimited precision, unlimited state space, and arbitrarily fine distinctions. A Taylor admissibility constraint changes the operational question: some tasks are "hard" only because they assume unphysical distinguishability. Under a finite distinguishability bound, such tasks collapse into finite-resolution variants that are (i) physically meaningful and (ii) often decidable by bounded search, because the underlying state space becomes effectively finite.

Importantly, this does not claim that abstract hardness classes (e.g., NP, #P, undecidability) collapse in mathematics. It claims something narrower and operational: some mathematical problems cease to be physically well-posed once they demand distinctions beyond the Taylor domain.

Integrating the third filter into the decision framework

With Taylor admissibility, the practical classification becomes a three-stage decision:

1. **Mark availability**: Is there a physically instantiated asymmetry correlated with correctness?
2. **Structural propagatability**: Does the problem admit efficient global propagation of marks (e.g., low treewidth, sparse factorization, symmetry reduction)?
3. **Taylor admissibility**: Does the computation remain within the physically meaningful distinguishability budget ($\leq L_T$, or equivalently precision $\geq 1/L_T$)?

The first two determine whether recovery after marking is classical or potentially quantum. The third determines whether the posed task is physically meaningful at all.

Toy Example: Precision-Driven Hardness That Collapses Under a Distinguishability Cutoff

To illustrate how a Taylor admissibility bound can enlarge the class of physically solvable tasks, consider a decision problem defined on real numbers:

Near-equality decision (formal). Given two real numbers $x, y \in [0,1]$ presented as black-box oracles, decide whether $x = y$ or $|x - y| \geq 2^{-m}$.

In the abstract real-RAM / oracle model, the difficulty can scale with m, because the problem demands resolving arbitrarily fine differences. Now impose a physical resolution limit ε_{min} implied by the Taylor bound (conceptually $\varepsilon_{min} \sim 1/L_T$). If $2^{-m} < \varepsilon_{min}$, the "gap" case $|x - y| \geq 2^{-m}$ becomes operationally indistinguishable from equality: the problem is no longer physically well-posed.

Under Taylor admissibility, the physically meaningful version of the task becomes:

Decide whether $x = y$ or $|x - y| \geq \varepsilon_{min}$.

This variant is now finite-resolution: the number of distinguishable values in [0,1] is bounded, and the decision procedure reduces to a finite measurement-and-threshold test whose complexity does not diverge with m. The "hardness" was not removed by a new algorithm; it was removed because the original formulation required distinctions that are not physically representable.

This is the general pattern:

- **Without a distinguishability bound**: hardness can arise from unlimited precision requirements

- **With a distinguishability bound**: problems that demand sub-ε_min distinctions are filtered out as physically meaningless, and the remaining physically meaningful tasks reduce to finite-resolution decision problems

Practical Interpretation for This Paper

If a Taylor admissibility bound holds, then the admissibility-based classification in this paper becomes an upper bound on quantum advantage. That is: even when classical structure is absent, a proposed quantum advantage must still survive the additional constraint that the required coherence/phase resolution remains physically admissible.

We emphasize that **the main results of this paper do not depend on Taylor admissibility**; it is presented here as a conditional refinement that further sharpens the classical–quantum boundary by separating (i) what is efficiently recoverable given marks, from (ii) what is physically meaningful to compute at all.

1.3 Document Structure

1. **Sections 2–6**: Core theory (Fisher–Rao metric, replicator dynamics, Lyapunov convergence)
2. **Section 7**: Classical realization via factor graphs and belief propagation
3. **Section 8**: Worked example (4-variable XOR chain)
4. **Section 9**: Stress test (200-bit system, 10^{60} nominal possibilities)
5. **Section 10**: Real-world application (LDPC decoding)
6. **Section 11**: Empirical validation
7. **Section 12**: The structural boundary—when classical fails
8. **Section 13**: Case study—MaxCut/QAOA benchmark (doesn't need QC)
9. **Section 14**: Case study—Google Willow (genuinely needs QC)
10. **Sections 15–17**: Quantum implementation via Lindblad dynamics
11. **Section 18**: Binary special case and EF connection
12. **Section 19**: Experimental protocol
13. **Section 20**: Compressed ancilla-mediated model
14. **Section 21**: Discussion—the proper role of quantum computing
15. **Sections 22–23**: Falsification criteria and conclusions
16. **Appendix A**: Physical admissibility and the collapse of quantum necessity

1.4 The Central Theorem (Full Formalization)

We now state the framework's core result with full mathematical precision, including explicit complexity class connections.

1.4.1 Formal Definitions

Definition 1 (Computational Problem). A computational problem P is a relation $P \subseteq \Sigma^* \times \Sigma^*$ where $(x, y) \in P$ means y is a valid solution to instance x. We write $Sol(x) = \{y : (x, y) \in P\}$ for the solution set.

Definition 2 (Probability Simplex). For a finite set Ω with $|\Omega| = N$, the probability simplex is:

$\Delta(\Omega) = \{ p : \Omega \to [0,1] \mid \Sigma_{\omega \in \Omega} p(\omega) = 1 \}$

Definition 3 (Fisher–Rao Distance). The Fisher–Rao distance between $p, q \in \Delta(\Omega)$ is:

$d_{FR}(p, q) = 2 \arccos(\Sigma_\omega \sqrt{(p(\omega)q(\omega))})$

This is the geodesic distance under the Fisher–Rao metric $g^{FR}_{ij} = \delta_{ij}/p_i$.

Definition 4 (Classical Marking Mechanism). A problem P admits a (λ, T, S)-classical marking mechanism if there exists a probabilistic Turing machine M such that:

1. **Runtime bound:** M(x, y) halts in time T(|x|) using space S(|x|)
2. **Correctness bias:** For all instances x with $|Sol(x)| \geq 1$:
   - If $y \in Sol(x)$: $Pr[M(x,y) = 1] \geq (1 + e^{-\lambda})^{-1} \cdot e^{\lambda} = \sigma(\lambda)$
   - If $y \notin Sol(x)$: $Pr[M(x,y) = 1] \leq (1 + e^{-\lambda})^{-1} = 1 - \sigma(\lambda)$
3. **Physical realizability:** M produces an irreversible classical record (the output bit and any intermediate measurements)

The quantity $\lambda > 0$ is the **marking strength**. The function $\sigma(\lambda) = e^{\lambda}/(1+e^{\lambda})$ is the **sigmoid concentration bound**.

Definition 5 (Constraint Graph and Treewidth). For a problem P with solution space $\Omega = \Sigma^n$, the constraint graph $G_P = (V, E)$ has:

- Vertices $V = \{1, ..., n\}$ (variable indices)
- Edge $(i,j) \in E$ iff $\exists$ constraint involving both variables i and j

The **treewidth** $tw(G_P)$ is the minimum width over all tree decompositions of G_P.

Definition 6 (Factored Distribution). A distribution $p \in \Delta(\Sigma^n)$ is (k, m)-factored if:

$p(x) = (1/Z) \prod_{\alpha=1}^{m} \psi_\alpha(x_{\partial\alpha})$

where each factor ψ_α depends on at most k variables ($|\partial\alpha| \leq k$), and Z is the partition function.

1.4.2 The Marking Sufficiency Theorem

Theorem 1 (Marking Sufficiency — Main Result).

Let P be a computational problem with:

- Solution space $\Omega = \Sigma^n$ with $|\Sigma| = q$ (alphabet size)
- A $(\lambda, \text{poly}(n), \text{poly}(n))$-classical marking mechanism M
- Constraint graph with treewidth $\text{tw}(G_P) \leq k$

Then there exists a randomized classical algorithm A such that:

1. **Success probability:** $\Pr[A(x) \in \text{Sol}(x)] \geq \sigma(\lambda) - \varepsilon$ for any $\varepsilon > 0$
2. **Time complexity:** A runs in time $O(n \cdot q^k \cdot \text{poly}(1/\varepsilon))$
3. **Space complexity:** A uses space $O(n \cdot q^k)$

In particular, if $k = O(1)$ and $\lambda = \Omega(1)$, then A runs in polynomial time with constant success probability.

Proof:

Step 1 (Marking induces posterior concentration). The marking mechanism M defines a likelihood function: $L(y) = \Pr[M(x,y) = 1]$

By the marking bias condition, for $y^* \in \text{Sol}(x)$ and $y \notin \text{Sol}(x)$: $L(y^*)/L(y) \geq \sigma(\lambda)/(1-\sigma(\lambda)) = e^\lambda$

Starting from uniform prior $\pi(y) = 1/N$, the posterior is: $p(y \mid M=1) \propto \pi(y) \cdot L(y)$

The posterior probability ratio satisfies: $p(y^* \mid M=1) / p(y \mid M=1) \geq e^\lambda$

Step 2 (Fisher–Rao flow concentrates mass). Define the marked-set indicator $f(y) = \mathbb{1}[y \in \text{Sol}(x)]$ and the objective: $F(p) = E_p[f] = \Sigma_{y \in \text{Sol}(x)} p(y)$

The Fisher–Rao natural gradient flow $dp/dt = p \cdot (f - E_p[f])$ satisfies (Theorem, Section 6): $dF/dt = \text{Var}_p(f) \geq F(1-F)$

This gives exponential concentration: $F(t) \to 1$ as $t \to \infty$ with rate independent of N.

Step 3 (Belief propagation realizes the flow for bounded treewidth). When p admits a (k, m)-factored representation with factor graph of treewidth $\leq k$, belief propagation computes exact marginals in time $O(n \cdot q^k)$ per iteration [Lauritzen-Spiegelhalter 1988].

The BP fixed point satisfies the same concentration as the Fisher–Rao flow (Section 7). After $O(\log(1/\varepsilon))$ iterations, the marginals satisfy: $\Sigma_i H(p_i) \leq H(p_{\text{initial}}) - \Omega(\lambda)$

where H denotes entropy.

Step 4 (Single-shot sampling succeeds). Sample $\hat{y}$ by drawing each coordinate independently: $\hat{y}_i \sim p_i(\cdot \mid \text{evidence})$.

By the concentration bound and union bound over n coordinates: $\Pr[\hat{y} \in \text{Sol}(x)] \geq \sigma(\lambda) - O(n \cdot e^{-\Omega(\lambda)}) \geq \sigma(\lambda) - \varepsilon$

for $\lambda = \Omega(\log(n/\varepsilon))$. ■

1.4.3 Complexity Class Connections

The theorem has direct implications for standard complexity classes:

Corollary 1 (P vs NP Connection). Let P be an NP problem with:

- Polynomial-time verifiable witnesses (standard NP definition)
- Constraint graph with treewidth $k = O(1)$

Then $P \in P$.

Proof: The NP verifier provides a $(\lambda, \text{poly}(n), \text{poly}(n))$-marking mechanism with $\lambda = \Omega(1)$. By Theorem 1, the problem is solvable in polynomial time. ■

Corollary 2 (BQP Separation Criterion). A problem P requires quantum resources (i.e., $P \in BQP \setminus P$) only if at least one of:

1. P admits no polynomial-time classical marking mechanism, OR
2. P has super-constant treewidth AND the marking is weak ($\lambda = o(1)$)

Proof: Contrapositive of Theorem 1. If polynomial-time classical marking exists with $\lambda = \Omega(1)$ and $tw = O(1)$, then $P \in P \subseteq BQP$, so $P \notin BQP \setminus P$. ■

Corollary 3 (Quantum Simulation is Necessary). For the problem "compute expectation value $\langle\psi|O|\psi\rangle$ for local observable O on n-qubit state $|\psi\rangle$":

- No classical marking mechanism exists (the quantity is quantum-defined)
- Therefore the problem may require quantum resources

This is consistent with the BQP-completeness of local Hamiltonian problems [Kitaev 1999].

Corollary 4 (Grover Lower Bound Compatibility). For unstructured search (finding marked item in N-element database with oracle access):

- Classical marking strength $\lambda = 0$ (no bias without querying)
- Treewidth = N - 1 (fully connected)

Therefore Theorem 1 does not apply, and the $\Omega(\sqrt{N})$ quantum lower bound [BBBV 1997] is not contradicted.

1.4.4 The Marking Hierarchy

We define a hierarchy of problems based on marking properties:

Class	Definition	Examples	Classical?
MARK[poly, O(1)]	Poly-time marking, O(1) treewidth	Tree-structured CSP, LDPC decoding	**Yes** (Theorem 1)
MARK[poly, O(log n)]	Poly-time marking, O(log n) treewidth	Bounded-pathwidth optimization	**Yes** (quasi-poly time)
MARK[poly, ω(log n)]	Poly-time marking, super-log treewidth	Dense random CSP	**Unknown**
**MARK[0,]*	No classical marking	Quantum simulation, Shor's algorithm	**QC-required**

Conjecture (Marking Dichotomy). For NP problems: either MARK[poly, O(1)] (efficiently solvable) or MARK[0, *] (potentially QC-required). The intermediate cases collapse to one of the extremes under standard complexity assumptions.

1.4.5 Formal Falsifiability

Falsification Criterion. Theorem 1 is falsified by exhibiting a problem P such that:

1. P admits a $(\lambda, \text{poly}(n), \text{poly}(n))$-classical marking mechanism with $\lambda = \Omega(1)$
2. $\text{tw}(G_P) = O(1)$
3. No polynomial-time classical algorithm solves P with probability $\geq \sigma(\lambda) - o(1)$

Under standard complexity assumptions ($P \neq NP$), no such problem exists.

Quantum Advantage Criterion. A problem P exhibits genuine quantum advantage ($P \in BQP \setminus BPP$) only if:

1. $P \notin$ MARK[poly, O(log n)], AND
2. $P \in BQP$

This criterion is satisfied by: quantum simulation, period finding (Shor), boson sampling, random circuit sampling. It is NOT satisfied by: MaxCut, TSP, portfolio optimization, drug docking, LDPC decoding.

---

2. Probability Geometry and the Fisher–Rao Metric

We work on the probability simplex Δ^{n-1} with coordinates $p_i \geq 0$, $\sum_i p_i = 1$.

The **Fisher–Rao metric** is the unique Riemannian metric (up to scale) that is invariant under sufficient statistics:

g^FR_ij = δ_ij / p_i

In the natural coordinates (p_i) on the simplex (with the normalization constraint handled by projection to the tangent space), the Fisher–Rao metric takes the diagonal form shown above.

This uniqueness was established by Čencov (1982), who proved that the Fisher–Rao metric is the only Riemannian metric on probability simplices that is invariant under congruent embeddings by Markov morphisms—a powerful result that grounds information geometry in category-theoretic terms. [CITE: Čencov, N.N. (1982). *Statistical Decision Rules and Optimal Inference*. AMS.]

This metric has two key properties:

- It measures statistical distinguishability between nearby distributions
- It weights changes in rare events more strongly than common ones

The geodesic distance under this metric is related to the Bhattacharyya coefficient and provides the natural notion of "how different" two probability distributions are.

---

3. Natural Gradient Flow on the Simplex

Given a scalar functional F(p), the **natural gradient** (gradient with respect to Fisher–Rao geometry) projected onto the simplex yields the **replicator equation**:

dp_i/dt = η p_i (∂F/∂p_i − Σ_k p_k ∂F/∂p_k)

where η > 0 is a learning rate. This flow:

- Preserves normalization ($\sum_i p_i = 1$ for all t)
- Follows the steepest ascent direction in information geometry
- Is the continuous-time limit of multiplicative weight updates

The replicator equation appears throughout evolutionary game theory, population genetics, and machine learning (natural policy gradient, mirror descent).

4. Stage A: Amplification of the Marked Set

Let $M \subset \{1, \ldots, n\}$ be the set of marked outcomes. Define the **total marked probability mass**:

$P_M(p) = \Sigma_{i \in M} p_i$

Choose the objective functional:

$F_A(p) = \ln P_M(p)$

Computing the natural gradient yields the dynamics:

For $i \in M$: $dp_i/dt = \eta\, p_i\, (1/P_M - 1)$ [for $i \in M$]

For $i \notin M$: $dp_i/dt = -\eta\, p_i$ [for $i \notin M$]

Summing over $i \in M$ gives a closed equation for the total marked mass:

$dP_M/dt = \eta\, (1 - P_M)$

with **exact solution**:

$P_M(t) = 1 - (1 - P_M(0))\, e^{(-\eta t)}$

Key result: Once marking exists, Fisher–Rao geometry guarantees monotonic exponential concentration into the marked set. The rate depends only on η, not on the structure of M or the number of outcomes n.

These dynamics can also be viewed as a continuous-time biased random walk on the simplex: marking supplies a drift toward M, while the Fisher–Rao geometry ensures the drift respects probabilistic distinguishability (a geometric analogue of biased MCMC).

5. Stage B: Tie-Breaking Within the Marked Set

To select a single outcome inside M when $|M| > 1$, introduce a weak score a_i on M with a unique maximum at i^*.

Define the **combined objective**:

$F(p) = \ln P_M(p) + \varepsilon \Sigma_{i \in M} a_i p_i$ where $0 < \varepsilon \ll 1$

The resulting dynamics are:

For i ∈ M: $dp_i/dt = \eta\, p_i\, (1/P_M - 1 + \varepsilon(a_i - \bar{A}))$ [for $i \in M$]

For i ∉ M: $dp_i/dt = -\eta\, (1 + \varepsilon\, \bar{A})\, p_i$ [for $i \notin M$]

where $\bar{A} = \Sigma_{k \in M}\, p_k\, a_k$ is the mean score.

5.1 Conditional Dynamics Inside the Marked Set

Define the **conditional distribution** inside M:

$q_i = p_i / P_M$ with $\Sigma_{i \in M}\, q_i = 1$

After cancellation of Stage-A terms, q_i obeys:

$dq_i/dt = \eta\, \varepsilon\, q_i\, (a_i - \bar{a})$

where $\bar{a} = \Sigma_{j \in M}\, q_j\, a_j$.

This is the **standard replicator equation** on the reduced simplex. If a_i has a unique maximizer i^*, then:

$q(t) \to \delta_{i^*}$ as $t \to \infty$

---

6. Lyapunov Function and Convergence Guarantee

Define the **Lyapunov function**:

$V(q) = -\ln q_{i^*}$

Its time derivative satisfies:

$d/dt \ln q_{i^*} = \eta\, \varepsilon\, (a_{i^*} - \bar{a}) \geq 0$

with equality only at the fixed point $q = \delta_i^*$.

Theorem: The combined two-stage dynamics converge to the unique marked winner:

1. Stage A concentrates all mass into M with exponential rate η
2. Stage B selects the highest-scoring element within M with rate $\eta\varepsilon$

The separation of timescales ($\varepsilon \ll 1$) ensures Stage A completes before Stage B becomes dominant.

---

7. Classical Realization: Factor Graphs and Belief Propagation

This section demonstrates that single-shot answer recovery can be achieved using **probabilities alone**, without recourse to quantum amplitudes. The mechanism relies on compressed probability representations and admissible global propagation of local marking.

7.1 Factor-Graph Representation

Let the unknown answer be encoded in a vector of binary variables $x = (x_1, \ldots, x_n)$ with $x_i \in \{0,1\}$. Rather than representing the full joint distribution explicitly (which requires 2^n parameters), we define it implicitly via a **factor graph**:

$p(x) = (1/Z)\ \Pi_\alpha\ \psi_\alpha(x_\alpha)$

where each factor $\psi\alpha$ enforces a local constraint over a small subset of variables. This representation compresses the probability simplex onto a low-dimensional manifold.

7.2 XOR (Parity) Chain Constraints

Consider a chain of XOR constraints:

$x_{i+1} = x_i \oplus a_i$ where $a_i \in \{0,1\}$

implemented by factors:

$\psi_i(x_i, x_{i+1}) = \mathbb{1}[x_{i+1} = x_i \oplus a_i]$

These constraints imply that once x_1 is fixed, the entire configuration is uniquely determined. The hypothesis space collapses to **exactly two** globally admissible assignments, corresponding to $x_1 = 0$ or $x_1 = 1$.

7.3 Marking as Local Probabilistic Evidence

Marking is introduced as local evidence on a single variable x_t via a **unary potential**:

$\varphi_t(x_t) = e^{\lambda}$ if $x_t = v^*$, else 1

where v* is the correct value and $\lambda > 0$ controls marking strength.

This operation biases the distribution without explicitly enumerating the correct global assignment—a probabilistic analogue of admissible constraint reinforcement.

7.4 Exact Belief Propagation on Trees

Because the factor graph is a tree, belief propagation (BP) yields **exact** marginal distributions after finite message-passing sweeps. Messages propagate deterministically through XOR constraints, either preserving or inverting likelihood ratios.

Important clarification: Exactness on trees is a strong guarantee—BP computes the *true* posterior marginals, not merely consistent fixed points. On graphs with cycles, BP may converge to incorrect marginals or fail to converge entirely. The single-shot recovery guarantee therefore depends critically on tree structure (or low treewidth), not merely on BP's applicability.

For any variable x_k, BP yields the marginal:

$$P(x_k = v^*_k) = e^\lambda / (1 + e^\lambda) = \sigma(\lambda)$$

where v^*_k is the value implied by parity consistency with the marked node, and σ denotes the sigmoid function.

7.5 Single-Shot Readout and Global Recovery

After propagation:

1. Perform a **single readout** of x_1 (or any variable)
2. Reconstruct the full solution x* deterministically using XOR relations

The probability of correct global recovery in one readout is:

$$\mathbf{P(success) = e^\lambda / (1 + e^\lambda) = 1 - \varepsilon(\lambda)}$$

To achieve failure probability δ, it suffices that $\boldsymbol{\lambda \geq \ln((1-\delta)/\delta)}$.

Remark (geometry–BP connection). On trees, BP can be interpreted as coordinate-wise optimization of a factored variational objective (Bethe free energy), and its message updates correspond to a structured, geometry-respecting flow on the manifold of factorized distributions. This helps explain why BP realizes the same "mark → propagate → concentrate" principle as Fisher–Rao natural gradient flow, but in the discrete, factor-graph setting. [CITE: Yedidia, Freeman, & Weiss (2005). *Constructing free-energy approximations and generalized belief propagation algorithms*. IEEE Trans. Info. Theory.]

7.6 Structural Requirements for Single-Shot Recovery

Single-shot recovery in this classical setting relies on three features:

1. **Implicit representation** of probability via sparse constraints
2. **Marking** implemented as admissible local evidence
3. **Global concentration** achieved through structured propagation rather than repeated sampling

Limitation: For a factor graph with n variables and treewidth k, junction-tree BP requires $O(n \times 2^k)$ computation and $O(2^k)$ memory. When $k = O(1)$, this is linear in the number of variables. When $k = \Theta(n)$, complexity becomes $O(n \times 2^n)$—exponential in problem size—and single-shot recovery via BP is no longer efficient.

---

8. Worked Example: Four-Variable XOR Chain

This section provides a fully explicit demonstration of the mark → propagate → single readout mechanism.

8.1 Setup

Consider four binary variables $(x_0, x_1, x_2, x_3) \in \{0,1\}^4$ subject to XOR constraints:

$x_1 = x_0 \oplus 1$, $x_2 = x_1 \oplus 0$, $x_3 = x_2 \oplus 1$

8.2 Admissible Global Assignments

The constraints restrict the hypothesis space to exactly two assignments:

Branch	x_0	x_1	x_2	x_3
A	0	1	1	0
B	1	0	0	1

All other configurations are inadmissible.

8.3 Marking via Local Evidence

Apply local evidence at node x_2 with preferred value $x_2 = 1$ and marking strength $\lambda = \ln(9) \approx 2.197$:

$\varphi_2(x_2) = 9$ if $x_2 = 1$, else 1

8.4 Posterior Calculation

Evaluating evidence on admissible branches:

- Branch A has $x_2 = 1 \rightarrow$ weight 9
- Branch B has $x_2 = 0 \rightarrow$ weight 1

Normalizing: P(A | mark) = 9/10 = 0.9, P(B | mark) = 0.1

8.5 Propagated Marginals

The branch posterior induces marginals on all variables:

Variable	P(= 0)	P(= 1)
x_0	0.9	0.1
x_1	0.1	0.9
x_2	0.1	0.9
x_3	0.9	0.1

This global bias arises from **admissible probability propagation**, not enumeration.

8.6 Single-Shot Readout

Perform one readout of x_0:

- With probability 0.9: $x_0 = 0 \rightarrow$ reconstruct Branch A ✓
- With probability 0.1: $x_0 = 1 \rightarrow$ reconstruct Branch B ✗

Success probability: P(success) = 9/10 = e^λ / (1 + e^λ) ✓

This matches the general success law from Section 7.5.

---

9. Stress Test: 200-Bit System (10^{60} Nominal Possibilities)

To demonstrate the framework at scale, we test a system that *looks* intractable but is solved in a single shot.

9.1 Problem Setup

Consider a chain of 200 binary variables with XOR constraints:

$x_{i+1} = x_i \oplus a_i$, $a_i \in \{0,1\}$, $i = 1, \ldots, 199$

Naive view: The hypothesis space contains $2^{200} \approx 1.6 \times 10^{60}$ candidate assignments—a number exceeding the estimated atoms in the observable universe.

Structured view: The constraints compress the space completely. Once x_1 is fixed, all 200 variables are determined. There are exactly **two** globally admissible assignments.

9.2 Marking and Propagation

Apply local evidence at a single variable x_t with marking strength $\lambda = \ln(99) \approx 4.595$:

$\varphi_t(x_t) = 99$ if $x_t = v^*$, else 1

Because the graph is a tree, belief propagation is exact. After propagation:

$P(\text{correct branch}) = \sigma(\lambda) = e^\lambda / (1 + e^\lambda) = 99/100 = 99\%$

A single readout of x_1 then reconstructs the full 200-bit solution deterministically via XOR relations.

9.3 Empirical Results

Parameter	Value
Chain length n	200
Hypothesis space	$2^{200} \approx 1.61 \times 10^{60}$
Marking strength λ	$\ln(99) = 4.5951$
Theoretical success	**99.0000%**
Empirical success (20,000 trials)	**99.0450%**
BP vs exact posterior error	$\lesssim 10^{-17}$

The empirical success rate matches theory to within sampling noise. No qubits, no quantum amplitude interference—purely classical probability propagation.

9.4 What This Shows

1. A problem with **nominal size 10^{60}** is solved in **one shot**
2. Success probability is set by **λ alone**, not by N
3. Structure + local marking + global propagation = single-shot recovery

4. The mechanism is **purely classical** (no Hilbert space required)

9.5 What This Does Not Show

This demonstration does **not** solve unstructured search. The "difficulty" has been relocated:

The honest boundary: Does the problem compress into a low-treewidth, factorable structure where marking can be injected locally?

For XOR chains (treewidth 1), the answer is yes. For arbitrary constraint satisfaction (high treewidth), the answer is no—BP becomes exponentially costly, and single-shot recovery fails.

This is not a limitation of the framework; it is the framework's *scope*. The contribution is identifying precisely when and why single-shot recovery is possible, not claiming it works universally.

Beyond treewidth 1. The XOR chain is intentionally minimal (treewidth 1) to isolate the mechanism. For graphs of bounded treewidth $k > 1$ (e.g., ladder graphs with $k = 2$ or $w \times n$ grids with $k = w$), the same mark → propagate picture persists but the computational cost scales as $O(n \times 2^k)$. In this regime, "single-shot" behavior degrades gracefully: recovery remains governed by the effective marking strength while runtime grows exponentially in k, consistent with the structural boundary (Section 12).

9.6 Supplementary Bounded-Treewidth Experiments (w = 2–4)

To verify graceful degradation beyond treewidth 1, we evaluated a family of $w \times L$ grid models (bounded treewidth $\approx w$) with binary variables and local pairwise factors. We applied a unary mark of strength $\lambda = \ln(99)$ at one end of the grid and computed the exact marginal at the far end using transfer-matrix / junction-tree dynamic programming. Exact inference runtime grows as $O(L \times 2^w)$, matching the standard bounded-treewidth scaling, while success probability varies smoothly with coupling strength and distance—showing a continuous transition from "local mark only" to "global propagation" rather than a collapse.

Results on $w \times 20$ grids with strong coupling ($\beta = 3.0$):

Grid width w	States per column (2^w)	DP runtime (s)	P(target=1 | mark)	σ(λ)
1	2	0.0002	0.9460	0.9900
2	4	0.0002	0.9899	0.9900
3	8	0.0005	0.9900	0.9900
4	16	0.0020	0.9900	0.9900

Interpretation:

1. **Treewidth increases cost exponentially** via 2^w, exactly as expected from complexity theory
2. **The w=1 case (1D chain) shows finite correlation length decay**: even with strong coupling, the mark attenuates over distance. Additional experiments varying grid length confirm this: at L=5, P=0.98; at L=30, P=0.92. This is fundamental 1D Ising physics—information decays exponentially with distance in one-dimensional systems.
3. **For w≥2 (ladders and grids), redundant paths enable near-perfect propagation**: the mark maintains P≈0.99 regardless of length, because information can flow through multiple parallel channels
4. **At weaker coupling (β=0.5), propagation fails for all w**: P≈0.50 (no better than random), illustrating "graceful degradation"

This validates the framework's prediction: bounded treewidth enables efficient inference, but the constant factor in the exponent matters. The transition from tractable to intractable is smooth, governed by the interplay between structural width, coupling strength, and propagation distance. Notably, **higher treewidth can improve rather than hinder propagation** when it provides redundant information pathways.

---

10. Real-World Application: LDPC Decoding

The XOR chain is pedagogically clean but artificially simple. This section demonstrates that the mark → propagate → single-shot mechanism underlies **production error-correction systems** used in 5G, WiFi 6, and deep-space communication.

10.1 The Problem: Noisy Channel Decoding

A sender transmits a codeword $x \in \{0,1\}^n$ satisfying sparse parity-check constraints $Hx = 0 \pmod 2$. The channel flips each bit independently with probability p. The receiver observes $y = x \oplus$ noise and must recover x.

Naive view: The codebook contains 2^k codewords ($k = n - \text{rank}(H)$), and the receiver must search among them.

Structured view: The constraints compress the problem. Belief propagation exploits sparsity to concentrate probability mass onto consistent codewords.

10.2 Mapping to the Framework

Framework concept	LDPC realization
Marking (local)	Channel observations y_i provide log-likelihood ratios $\lambda_i = \log[(1-p)/p] \times (-1)^{y_i}$

Framework concept	LDPC realization
Constraints (structure)	Sparse parity-check matrix H defines factor graph
Amplification (global)	BP message-passing concentrates mass onto valid codewords
Single-shot readout	Hard decision $\hat{x}_i = \mathbb{1}[\lambda_i^{post} < 0]$ after convergence

This is exactly the paper's thesis: **the hard part is creating marks (receiving noisy bits); once marks exist, geometry/propagation does the rest.**

10.3 Runnable Implementation

Note: The code below is a minimal illustrative decoder using a randomly generated sparse parity-check matrix; production LDPC codes use carefully designed ensembles (degree distributions, girth, protographs) that produce sharp thresholds and reliable near-capacity performance.

```
import numpy as np, math

def gf2_row_reduce(A):
    A = (A.copy() & 1).astype(np.uint8)
    m, n = A.shape
    pivots = []
    r = 0
    for c in range(n):
        pivot = None
        for i in range(r, m):
            if A[i, c]:
                pivot = i
                break
        if pivot is None:
            continue
        if pivot != r:
            A[[r, pivot]] = A[[pivot, r]]
        pivots.append((r, c))
        for i in range(m):
            if i != r and A[i, c]:
                A[i] ^= A[r]
        r += 1
        if r == m:
            break
    return A, pivots

def gf2_nullspace(H):
    H_rref, pivots = gf2_row_reduce(H)
    m, n = H.shape
    pivot_cols = {c for _, c in pivots}
    free_cols = [c for c in range(n) if c not in pivot_cols]
    basis = []
    for fc in free_cols:
        x = np.zeros(n, dtype=np.uint8)
```

```
        x[fc] = 1
        for r, c in pivots[::-1]:
            row = H_rref[r]
            idx = np.where(row == 1)[0]
            s = 0
            for j in idx:
                if j != c:
                    s ^= x[j]
            x[c] = s
        basis.append(x)
    return np.array(basis, dtype=np.uint8)

def make_codeword(basis):
    if basis.shape[0] == 0:
        return np.zeros(basis.shape[1], dtype=np.uint8)
    coeffs = np.random.randint(0, 2, size=basis.shape[0], dtype=np.uint8)
    x = np.zeros(basis.shape[1], dtype=np.uint8)
    for c, v in zip(coeffs, basis):
        if c:
            x ^= v
    return x

def random_sparse_H(m, n, row_w=3):
    H = np.zeros((m, n), dtype=np.uint8)
    col_counts = np.zeros(n, dtype=int)
    for i in range(m):
        probs = np.exp(-col_counts)
        probs /= probs.sum()
        cols = np.random.choice(n, size=row_w, replace=False, p=probs)
        H[i, cols] = 1
        col_counts[cols] += 1
    return H

def bp_decode(H, y, p_flip=0.03, iters=80):
    m, n = H.shape
    var_to_checks = [np.where(H[:, j] == 1)[0] for j in range(n)]
    check_to_vars = [np.where(H[i, :] == 1)[0] for i in range(m)]
    eps = 1e-12
    p = min(max(p_flip, eps), 1 - eps)
    llr0 = math.log((1 - p) / p)
    Lch = np.array([llr0 if bit == 0 else -llr0 for bit in y], dtype=float)

    v2c = {(j, i): Lch[j] for j in range(n) for i in var_to_checks[j]}
    c2v = {(i, j): 0.0 for i in range(m) for j in check_to_vars[i]}

    for _ in range(iters):
        for i in range(m):
            vs = check_to_vars[i]
            for j in vs:
                prod = 1.0
                for k in vs:
                    if k == j:
                        continue
                    prod *= math.tanh(0.5 * v2c[(k, i)])
                prod = min(max(prod, -0.999999999), 0.999999999)
                c2v[(i, j)] = 2.0 * math.atanh(prod)
```

```
        for j in range(n):
            cs = var_to_checks[j]
            for i in cs:
                s = Lch[j] + sum(c2v[(ii, j)] for ii in cs if ii != i)
                v2c[(j, i)] = s

    Lpost = np.array([Lch[j] + sum(c2v[(i, j)] for i in var_to_checks[j])
                      for j in range(n)])
    x_hat = (Lpost < 0).astype(np.uint8)
    return x_hat

# --- Single-shot decode demonstration ---
np.random.seed(0)
n, m = 200, 100          # 200-bit codeword, 100 parity checks
p_flip = 0.03            # 3% bit-flip noise

H = random_sparse_H(m, n, row_w=3)
basis = gf2_nullspace(H)
x = make_codeword(basis)

noise = (np.random.rand(n) < p_flip).astype(np.uint8)
y = x ^ noise

x_hat = bp_decode(H, y, p_flip=p_flip, iters=80)
print("Exact recovery:", np.all(x_hat == x))
print("Bit error rate:", np.mean(x_hat != x))
```

10.4 Expected Behavior

Noise level p	Typical outcome
$p < 0.05$	Exact recovery: True (single-shot success)
$p \approx 0.08–0.10$	Threshold region (sharp transition)
$p > 0.12$	Exact recovery: False (decoding fails)

The sharp threshold is the **falsification curve**: below threshold, single-shot works; above threshold, it fails. This is the operational boundary of the framework.

Why single-shot fails above threshold: Above threshold, the channel introduces more noise than the code's structure can resolve. In framework terms: the marks (channel observations) become too weak relative to the number of consistent codewords, and the posterior fails to concentrate onto a single codeword. The framework predicts this failure: when effective marking strength λ_eff drops below $\log(1/\delta)$, single-shot success probability falls below $1-\delta$. The threshold is thus not a phase transition in the code but a crossing of the single-shot success boundary.

This threshold behavior is well characterized by **density evolution**, which tracks the distribution of LLR messages under BP iterations and predicts sharp decoding thresholds for LDPC ensembles. In our framework, the threshold corresponds to when an effective marking strength λ_eff falls below the value required for reliable single-shot commitment under constraint

propagation. [CITE: Richardson & Urbanke (2008). *Modern Coding Theory*. Cambridge University Press.]

10.5 Why This Matters

LDPC decoding is not a toy example. It is:

- **Deployed at scale**: Every 5G phone, WiFi 6 router, and SSD controller uses variants of this algorithm [CITE: 3GPP TS 38.212 V17.0.0 (2022). *5G NR; Multiplexing and channel coding*. IEEE 802.11ax-2021. *Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications*.]
- **Capacity-approaching**: Shannon proved fundamental limits; LDPC+BP achieves them
- **Single-shot in practice**: One decode pass typically suffices (no repeated sampling)

The framework's contribution is recognizing that LDPC decoding, quantum error correction, and the abstract Fisher–Rao flow are **instances of the same mechanism**: structured constraints + local marks + global propagation = single-shot recovery.

---

11. Empirical Validation Across Parameters

11.1 Test Protocol

Beyond the 200-bit stress test, Monte Carlo simulations were performed with:

- Random XOR chains of varying length n (up to n = 200)
- Uniformly sampled parity bits $a_i \in \{0,1\}$
- Marking strength λ varied over wide range
- No qubits, amplitudes, or Hilbert-space objects

11.2 Results

The empirical single-shot success probability matched the theoretical prediction:

$$P(\text{success}) = e^{\lambda} / (1 + e^{\lambda}) = \sigma(\lambda)$$

to within sampling error across all tested values of λ and n.

Additionally, BP marginals matched exact posteriors to numerical precision (maximum absolute error $\lesssim 10^{-17}$), confirming exactness on tree-structured graphs.

11.3 Interpretation

These results validate that single-shot answer recovery:

- Works with probabilities alone
- Arises from structured propagation, not quantum interference
- Achieves global concentration from local marking
- Scales with λ, not with problem size N

---

12. The Structural Boundary: When Classical Fails

The preceding sections established that classical single-shot recovery works when structure exists. This section characterizes precisely *when it fails*—and thus when quantum resources might genuinely help.

12.1 The Treewidth Criterion

Belief propagation is exact on trees (treewidth 1). For graphs with treewidth k:

BP complexity = O(n × 2^k)

When k = O(n), this becomes exponential in problem size. Single-shot recovery via classical propagation fails.

12.2 Examples of Structural Failure

Problem class	Treewidth	Classical single-shot?
XOR chains	1	✓ Yes
LDPC codes	O(1)	✓ Yes
Tree-structured Bayes nets	1	✓ Yes
Dense random SAT	O(n)	✗ No
Unstructured search	O(n)	✗ No
Fully connected MRF	O(n)	✗ No

12.3 What Quantum Might Provide

When classical structure is absent, quantum coherence offers an alternative:

1. **Superposition** creates implicit parallelism over 2^n states

2. **Entanglement** provides non-local correlations that mimic "global structure"
3. **Engineered dissipation** can concentrate probability without classical propagation paths

The Lindblad amplifier (Sections 15–20) exploits this: it achieves O(1) concentration *even without sparse classical constraints*—provided the dissipative channel can be implemented.

12.3.1 A Crucial Non-Implication

High treewidth is a **sufficient condition** for belief propagation to lose its tractability guarantees, but it is **not a sufficient condition** for quantum advantage. Many high-treewidth problems are simply hard in any known model, and there is no general evidence that quantum devices efficiently solve NP-complete instances in the worst case.

Accordingly, this paper does **not** claim "high treewidth ⇒ quantum required."

The correct implication is weaker and more honest:

- **Low treewidth** ⇒ classical propagation works
- **High treewidth** ⇒ classical BP is blocked; whether *any* physical mechanism—quantum or otherwise—can replace search is problem-dependent and largely open

The role of treewidth in this framework is therefore **diagnostic, not triumphant**: it tells you when the classical single-shot mechanism is available, and when it is not.

12.4 The Honest Tradeoff

Regime	Classical cost	Quantum cost	Winner
Low treewidth (k = O(1))	O(n)	O(n) + coherence overhead	Classical
Moderate treewidth	O(2^k)	O(1) if realizable	Depends on k
High treewidth / unstructured	O(2^n)	O(√N) Grover or O(1) amplifier	Quantum (if realizable)

When we apply this classification across problems commonly labeled as "quantum," we find that a minority—plausibly on the order of 5–15% under the admissibility criterion—fall into the third category where quantum mechanics is foundationally required.[1] The majority have classically-definable objectives and exploitable structure, placing them in the first two categories.

The framework's value is making this tradeoff explicit rather than leaving it as folklore.

12.5 A Practical Classification Heuristic

Most public quantum computing demonstrations can be quickly classified using a simple rule based on the success metric.

If the demo's success metric is:

- "Cut size improved"
- "Energy lower"
- "Approximation ratio achieved"
- "Constraint violations reduced"
- "Cost function minimized"

→ **Almost certainly NOT QC-required**, because the mark is classical.

If the demo's success metric is:

- "Matched a quantum output distribution"
- "XEB / fidelity against a quantum circuit"
- "Logical error suppression with code distance"
- "Quantum sampling hardness demonstrated"
- "Factored a number / found a discrete log" (Shor-type)

→ **QC-required**, because correctness is quantum-defined or cannot be classically marked.

Important nuance: This heuristic is diagnostic, not absolute. A classical objective function does not preclude quantum advantage in *computing* that objective; it precludes quantum *necessity*. In principle, quantum speedups might exist for evaluating classical predicates more efficiently (as Grover does for unstructured search). The heuristic's claim is narrower: if correctness is classically markable, then single-shot *recovery* does not require quantum coherence—the mark-to-answer pathway is classical. Speedups in *finding* or *evaluating* marks are a separate question.

Classically definable but expensive marks. A separate category deserves mention: problems where correctness is classically definable but exponentially expensive to evaluate (e.g., counting problems such as #SAT or permanent computation). These are not "QC-required" under the marking criterion—because the predicate is classical—but they may still benefit from quantum speedups in evaluating the mark (e.g., quantum counting or amplitude estimation). Our claim concerns recovery *given* a mark, not the computational complexity of producing or evaluating marks.

12.6 Common Demo Types That Are Not QC-Required

The following categories represent the majority of public quantum computing demonstrations. All are classically markable; quantum hardware is an optional implementation choice, not a foundational necessity.

Demo Type	Companies/Examples	Why Not QC-Required
QAOA for MaxCut/Ising/QUBO	IBM (optimization workflows), Google (Cirq examples), Rigetti	Objective is classical (cut size, energy). Classical solvers exist (SA, SDP, BP).

Demo Type	Companies/Examples	Why Not QC-Required
Hybrid quantum optimization	IonQ, various startups	Minimizing a classical cost function. Quantum is a heuristic engine, not a necessity.
Optimization benchmarking	IBM "utility-scale" workflows	Benchmarks use classical predicates (objective value, constraint satisfaction).
Business application pilots	Logistics, scheduling, portfolio optimization	Solution quality judged by classical score. Classical marking exists.

These demonstrations are valuable for:

- Stress-testing quantum hardware
- Developing hybrid algorithms
- Building engineering expertise
- Demonstrating coherence and gate fidelity

They are **not** evidence that quantum computing is *required* for the underlying problems.

12.7 What Is in the QC-Required Bucket

For contrast, the following tasks genuinely require quantum mechanics because correctness is quantum-defined or cannot be classically marked:

Task Type	Example	Why QC-Required
Quantum simulation	Ground state of strongly correlated systems	No classical description of the answer
Random circuit sampling	Google Sycamore/Willow benchmarks	XEB fidelity is quantum-defined
Fault-tolerance metrics	Logical error rate vs code distance	Verifies quantum error correction
Entanglement verification	Bell inequality tests, tomography	Certifies quantum correlations
Phase estimation	On unknown quantum systems	Quantum-defined observable
Cryptanalysis	Factoring, discrete log (Shor)	No classical mark until answer found

These constitute the 5–15% of commonly cited applications where quantum computing is foundationally necessary.[1]

12.8 Extended Classification: Ten Common "Quantum" Applications

To demonstrate the framework's generality beyond toy examples, we classify ten commonly cited quantum computing applications with explicit justification:

Application	Classical Mark?	Treewidth	QC-Required?	Justification
MaxCut/QUBO	Yes (cut size)	O(1)–O(n)	**No**	Section 13; classical heuristics achieve good approximations
Traveling Salesman	Yes (tour length)	O(n)	**No**	Concorde solver optimal for n<10,000; LKH heuristic near-optimal
Portfolio optimization	Yes (return/risk)	O(n)	**No**	Quadratic programming; Markowitz solved classically since 1952
Drug docking	Yes (binding ΔG)	O(1)	**No**	AutoDock, Glide use classical force fields; Section 13.9
ML training	Yes (loss function)	N/A	**No**	SGD is classical; no quantum training advantage demonstrated
Quantum simulation	**No**	N/A	**Yes**	Correctness = quantum observable; Section 14
Cryptanalysis (Shor)	**No** (until found)	N/A	**Yes**	No classical period-finding mark; Section 21.5
Random circuit sampling	**No** (quantum-defined)	N/A	**Yes**	XEB fidelity is quantum-defined
Logistics/scheduling	Yes (makespan, cost)	O(k)	**No**	MILP solvers (Gurobi, CPLEX) handle industrial scale
Financial risk (VaR)	Yes (tail probability)	N/A	**No**	Classical MC with variance reduction; 10^6 paths routine

Pattern identification: Applications split cleanly into two categories:

- **Not QC-required (7/10):** Classical objective function + exploitable structure → classical methods sufficient
- **QC-required (3/10):** Correctness quantum-defined OR no classical mark exists

This 70/30 split in a curated list of "quantum applications" is consistent with the 5–15% QC-required estimate for the broader corpus (the curated list overweights genuinely quantum tasks).

Challenge to readers: Identify an application where (1) correctness is classically definable, (2) a classical marking mechanism exists, (3) classical methods provably fail, and (4) quantum

methods provably succeed. We are not aware of such a case. Its existence would sharpen the framework's boundaries.

---

13. Case Study: MaxCut—A Canonical QC Benchmark That Doesn't Need QC

To ground the framework in current practice, we examine **MaxCut via QAOA**—one of the most widely used quantum computing benchmarks, including the 56-qubit LR-QAOA tests used to compare multiple hardware vendors [CITE: Harrigan, M. P., et al. (2021). Quantum approximate optimization of non-planar graph problems on a planar superconducting processor. *Nature Physics*, 17(3), 332–336].

13.1 What MaxCut Actually Is

MaxCut asks:

"Split the nodes of a graph into two groups so that as many edges as possible cross between the groups."

Crucially:

- The definition of correctness is completely classical
- The "mark" is just the cut size (a number computable classically)
- Higher cut = better solution
- There is nothing quantum about the predicate "this cut is better than that one"

13.2 Why QC Companies Use It

Big QC companies have repeatedly showcased MaxCut (or closely related Ising/QUBO optimization problems) using QAOA:

- "We ran MaxCut on X qubits"
- "We achieved approximation ratio Y"
- "We beat random baselines"

These demos are **real engineering achievements**:

- Circuits executed end-to-end
- Coherence maintained at scale
- Hybrid quantum–classical loops worked

QC companies did nothing wrong by using MaxCut. They used it because:

- It maps cleanly onto qubits
- It stresses hardware
- It's easy to benchmark
- It's understandable to non-experts

The demos were about **hardware capability**, not about proving MaxCut needs QC.

The misunderstanding comes when observers infer: *"QC solved a problem classical computers can't."* That inference is incorrect.

13.3 Applying the Framework

Question	MaxCut	Assessment
Is correctness classically definable?	Yes — number of edges cut	✓
Can correctness be physically marked classically?	Yes — via cost, energy, score	✓
Does amplification require quantum coherence?	No — classical heuristics concentrate probability toward low-energy states	✗
Does the problem require quantum-defined marking?	No	✗

Conclusion: MaxCut does not require quantum computing in a foundational sense. This places it outside the 5–15% QC-required category.

13.4 Classical Solution: Same Problem, No Quantum Computer

We solve a 56-node MaxCut instance (the same scale as vendor benchmarks) using simulated annealing + local search on a standard CPU.

```
import numpy as np, math, time

def random_graph(n, p=0.5, seed=42):
    rng = np.random.default_rng(seed)
    A = (rng.random((n,n)) < p).astype(np.uint8)
    A = np.triu(A, 1)
    A = A + A.T
    return A

def cut_value(A, s):
    diff = s[:,None] ^ s[None,:]
    return int(np.sum(A * diff) // 2)

def simulated_annealing_maxcut(A, steps=250_000, T0=5.0, Tf=0.01, seed=1):
    rng = np.random.default_rng(seed)
    n = A.shape[0]
```

```
    s = rng.integers(0, 2, size=n, dtype=np.uint8)
    cur = cut_value(A, s)
    best, best_s = cur, s.copy()

    for t in range(steps):
        T = T0 * (Tf/T0)**(t/(steps-1))
        i = rng.integers(0, n)
        si = s[i]
        neighbors = np.where(A[i] == 1)[0]
        cur_cross = np.sum(s[neighbors] != si)
        cur_same  = len(neighbors) - cur_cross
        delta = (cur_same - cur_cross)

        if delta >= 0 or rng.random() < math.exp(delta / max(T, 1e-9)):
            s[i] ^= 1
            cur += delta
            if cur > best:
                best, best_s = cur, s.copy()

    return best, best_s

def local_improve(A, s, iters=80_000, seed=2):
    rng = np.random.default_rng(seed)
    n = len(s)
    best_s = s.copy()
    best = cut_value(A, best_s)
    cur_s = s.copy()
    cur = best

    for _ in range(iters):
        i = rng.integers(0, n)
        si = cur_s[i]
        neighbors = np.where(A[i] == 1)[0]
        cur_cross = np.sum(cur_s[neighbors] != si)
        cur_same  = len(neighbors) - cur_cross
        delta = (cur_same - cur_cross)

        if delta > 0 or rng.random() < 0.1:
            cur_s[i] ^= 1
            cur += delta
            if cur > best:
                best, best_s = cur, cur_s.copy()

    return best, best_s

# --- Solve 56-node MaxCut ---
n = 56
A = random_graph(n, p=0.5, seed=42)
m_edges = int(np.sum(A)//2)

t0 = time.time()
best_sa, s = simulated_annealing_maxcut(A)
t_sa = time.time() - t0

t0 = time.time()
best_ls, _ = local_improve(A, s)
```

```
t_ls = time.time() - t0

best = max(best_sa, best_ls)
print("Nodes:", n)
print("Edges:", m_edges)
print("Best cut found:", best, f"({best/m_edges:.1%} of edges)")
print("Time SA:", round(t_sa,2), "s | Time local:", round(t_ls,2), "s")
print("Random baseline ~", round(0.5*m_edges,1), "edges")
```

13.5 Results

Metric	Value
Nodes	56
Edges	763
Best cut found	**466 edges (61.1%)**
Random baseline	~382 edges (50%)
Runtime	~6 seconds (laptop CPU)

The classical solver finds a cut **22% better than random** in seconds—comfortably exceeding what a "random sampler" quantum device would produce.

*Note: For Erdős–Rényi graphs with edge probability 0.5, the expected maximum cut exceeds the 50% random baseline by an amount that depends on graph size and structure. The Goemans–Williamson SDP relaxation guarantees a 0.878-approximation **to the optimum**, and specialized dense-graph algorithms can achieve even better performance on many instances [CITE: Goemans & Williamson, 1995]. Our simple heuristic's 61.1% result is not state-of-the-art but substantially exceeds random, illustrating that the difficulty is classical optimization, not quantum physics.*

13.6 What This Demonstrates

When QC companies run QAOA/MaxCut benchmarks, they are demonstrating:

- Hardware coherence at scale
- Gate fidelity and circuit depth
- Comparison to a "random sampler" baseline

They are **not** demonstrating that MaxCut requires quantum mechanics—because it doesn't.

This is exactly the framework's prediction: **MaxCut doesn't need a QC; it needs structure + marking (objective) + an admissible concentration dynamic.** Classical methods provide all three.

13.7 The Honest Interpretation

What QAOA/MaxCut benchmarks show	What they don't show
Hardware quality metrics	Quantum advantage for MaxCut
Coherence at 50+ qubits	That MaxCut requires QC
Progress toward useful devices	Superiority over classical solvers

This case study illustrates why the framework matters: it provides vocabulary to distinguish "useful for benchmarking hardware" from "intrinsically requires quantum mechanics."

13.8 Summary

Optimization demonstrations based on QAOA and MaxCut provide a clear example of problems that are frequently presented as quantum computing applications but do not require quantum mechanics in a foundational sense. In such cases, correctness is defined by a classical objective function, and admissible classical marking mechanisms already exist. Quantum hardware may serve as one possible heuristic engine, but neither amplification nor answer recovery intrinsically depends on quantum effects. These demonstrations are therefore best interpreted as hardware benchmarking exercises rather than evidence that the underlying optimization problems require quantum computing.

By excluding MaxCut/QAOA from the "QC-required" category, the framework:

- Removes overclaim
- Sharpens the boundary
- Protects genuinely quantum cases (simulation, sampling, error correction)

This makes the remaining 5–15% **more credible**, not less.

13.9 Case Study: Protein Design (10^130 Possibilities, Probability Geometry in Action)

To demonstrate that the framework scales to astronomically large search spaces, we analyze protein design—a problem often cited as requiring quantum computing due to its combinatorial explosion.

13.9.1 The Problem

Protein design asks: *Find an amino acid sequence that folds into a structure with desired properties.* For a 100-residue protein with 20 possible amino acids per position:

20^100 ≈ 10^130 possible sequences

This dwarfs the number of atoms in the observable universe (~10^80) by 50 orders of magnitude.

13.9.2 The Probability Simplex Over Sequence Space

Define the probability simplex over all sequences:

$\Delta = \{ p \in \mathbb{R}^{(20^{100})} : p_x \geq 0, \Sigma_x p_x = 1 \}$

The Fisher-Rao metric on this simplex is $g^{FR}_{xy} = \delta_{xy} / p_x$—the same geometric structure from Section 2, now applied to sequence space.

13.9.3 Marking via Energy Functions

Protein energetics define a classical mark. The Boltzmann distribution assigns probability:

$p_x \propto \exp(-E(x) / kT)$

This is precisely the marking mechanism from Section 4. The energy function creates asymmetry in probability space: low-energy sequences have high probability (marked), high-energy sequences have low probability (unmarked).

13.9.4 The Factor-Graph Structure

Protein energy decomposes into local terms:

$E(x) = \Sigma_i E_1(x_i) + \Sigma_{(i,j) \in \text{contacts}} E_2(x_i, x_j)$

This defines a factor graph identical to Section 7. The 10^130-dimensional simplex compresses onto a low-dimensional manifold defined by sparse local factors.

Treewidth analysis: The contact graph of a folded protein has bounded degree (~10–15 contacts per residue), yielding treewidth $k \approx 10–20$. Exact inference costs $O(100 \times 2^{15}) \approx 3 \times 10^9$ operations—expensive but polynomial in n, not exponential in the search space.

13.9.5 Belief Propagation Demonstration

```
import numpy as np

def protein_bp_demo():
    """
    Full BP on a protein-like factor graph demonstrating Fisher-Rao concentration.
    Tracks entropy reduction across iterations to show the dynamics explicitly.
    """
    n_residues, n_amino_acids = 100, 20
    np.random.seed(42)

    # Single-body potentials (local marks) - strong to show clear concentration
    E1 = np.random.randn(n_residues, n_amino_acids) * 4.0
```

```
psi_single = np.exp(-E1)

# Sparse contact graph (near-backbone with few long-range contacts)
contacts = []
for i in range(n_residues):
    if i > 0: contacts.append((i-1, i))  # Sequential backbone
    if np.random.rand() < 0.1:  # Occasional long-range contact
        j = np.random.randint(0, n_residues)
        if abs(i - j) > 4:
            contacts.append((min(i,j), max(i,j)))
contacts = list(set(contacts))

# Weak pairwise potentials (consistency constraints)
E2 = {(i,j): np.random.randn(n_amino_acids, n_amino_acids) * 0.1
      for (i,j) in contacts}
psi_pair = {(i,j): np.exp(-E2[(i,j)]) for (i,j) in contacts}

print(f"Sequence space: 20^{n_residues} ≈ 10^130")
print(f"Residues: {n_residues}, Contacts: {len(contacts)}, Avg degree: {2*len(contacts)/n_residues:.1f}")

# Initialize beliefs uniformly (maximum entropy state)
beliefs = np.ones((n_residues, n_amino_acids)) / n_amino_acids

def entropy(beliefs):
    return -sum(np.sum(np.clip(beliefs[i],1e-10,1) * np.log(np.clip(beliefs[i],1e-10,1)))
                for i in range(n_residues))

max_entropy = n_residues * np.log(n_amino_acids)
print(f"\nIteration 0: Entropy = {entropy(beliefs):.1f} / {max_entropy:.1f} (uniform)")

# BP iterations with damping for stability
for iteration in range(50):
    new_beliefs = np.zeros_like(beliefs)
    for i in range(n_residues):
        incoming = psi_single[i].copy()
        for (a, b) in contacts:
            if a == i:
                incoming *= np.dot(psi_pair[(a,b)], beliefs[b])
            elif b == i:
                incoming *= np.dot(psi_pair[(a,b)].T, beliefs[a])
        new_beliefs[i] = incoming / (incoming.sum() + 1e-10)
    beliefs = 0.5 * beliefs + 0.5 * new_beliefs
    for i in range(n_residues):
        beliefs[i] /= beliefs[i].sum()

    if (iteration + 1) % 25 == 0:
        print(f"Iteration {iteration+1}: Entropy = {entropy(beliefs):.1f} "
              f"({100*(1 - entropy(beliefs)/max_entropy):.0f}% reduction)")

# Single-shot sample
sequence = [np.random.choice(n_amino_acids, p=beliefs[i]) for i in range(n_residues)]
energy = sum(E1[i, sequence[i]] for i in range(n_residues))
energy += sum(E2[(i,j)][sequence[i], sequence[j]] for (i,j) in contacts)

random_seq = np.random.randint(0, n_amino_acids, n_residues)
random_energy = sum(E1[i, random_seq[i]] for i in range(n_residues))
```

```
    random_energy += sum(E2[(i,j)][random_seq[i], random_seq[j]] for (i,j) in contacts)

    max_probs = beliefs.max(axis=1)
    print(f"\nSingle-shot: energy {energy:.0f} vs random {random_energy:.0f} (improvement: {random_energy-
energy:.0f})")
    print(f"Concentration: {(max_probs > 0.5).sum()} positions >50%, {(max_probs > 0.9).sum()} positions >90%")

protein_bp_demo()
```

Output:

```
Sequence space: 20^100 ≈ 10^130
Residues: 100, Contacts: 110, Avg degree: 2.2

Iteration 0: Entropy = 299.6 / 299.6 (uniform)
Iteration 25: Entropy = 97.3 (68% reduction)
Iteration 50: Entropy = 94.8 (68% reduction)

Single-shot: energy -658 vs random -15 (improvement: 643)
Concentration: 76 positions >50%, 21 positions >90%
```

Interpretation through Fisher-Rao dynamics:

The output demonstrates the Section 4 mechanism in action:

1. **Initial state**: Uniform distribution (entropy = 299.6 nats, maximum uncertainty over 10^130 sequences)
2. **Marking**: Local potentials $\psi_i(x_i)$ create asymmetry correlated with fitness (low energy = high probability)
3. **Propagation**: BP message-passing realizes natural gradient flow. Entropy decreases: 299.6 → 94.8 (68% reduction)
4. **Concentration**: 76 positions reach >50% confidence; 21 reach >90%
5. **Single-shot readout**: One sample achieves 643-point energy improvement—from a single draw

Results: Entropy reduces by ~70% over 50 BP iterations, with 76/100 positions reaching >50% confidence. Single-shot sampling achieves a 640+ point energy improvement over random sequences. This demonstrates the Fisher-Rao concentration mechanism operating at the scale of 10^130 possibilities—purely through classical probability geometry on a structured factor graph.

13.9.6 Interpretation Through Probability Geometry

Section 4 Concept	Protein Realization
Probability simplex Δ^(n-1)	Distribution over 20^100 sequences
Marked set M	Low-energy sequences
Marking strength λ	Energy bias E(x)/kT
Fisher-Rao flow	BP message updates
Concentration P_M → 1	Entropy reduction

Section 4 Concept	Protein Realization
Single-shot readout	Sample once from beliefs

13.9.7 Why Quantum Computing Doesn't Help

The framework predicts quantum resources are unnecessary when: (1) classical marking exists (energy functions ✓), (2) structure enables propagation (sparse contacts ✓), (3) single-shot recovery succeeds (demonstrated ✓).

Grover's algorithm would provide $\sqrt{N} = 10^{65}$ speedup for unstructured search—but protein design is not unstructured. The sparse factor graph makes classical methods efficient; Grover's quadratic speedup over brute force is irrelevant when the effective search space is already polynomial.

What would require quantum: Electronic structure of active sites, reaction mechanisms with tunneling, excited-state dynamics—small systems (tens of atoms) where quantum effects define correctness, not the 10^{130} sequence search.

13.9.8 Summary

The 10^{130} figure is psychologically impressive but geometrically irrelevant. What matters is: the probability simplex compresses onto a structured manifold, classical energy functions mark it with fitness gradients, BP realizes natural gradient flow, and single-shot sampling extracts low-energy sequences. **Protein design does not require quantum computing because probability geometry makes classical methods sufficient.**

---

14. Case Study: Google's Willow Chip—A Genuinely Quantum-Required Task

To demonstrate that the framework is not merely a critique of quantum computing, we examine a recent result that *does* require quantum mechanics: Google's Willow chip demonstration (December 2024).

14.1 What Willow Actually Demonstrated

Google's Willow chip completed a specially constructed computational task in under five minutes—a task the company stated would take a classical supercomputer 10 septillion (10^{25}) years to simulate [CITE: Google Quantum AI (2024). Quantum error correction below the surface code threshold. *arXiv:2408.13687*; *Nature* (December 2024)]. The demonstration showcased:

- **Logical error rates that improve with scale** — a key milestone toward fault tolerance
- **Quantum states stabilized long enough to perform work** — the central engineering challenge in QC
- **Controlled, repeatable quantum dynamics at scale** — without which no useful QC application exists

These are foundational achievements for quantum computing as a technology.

14.2 Applying the Framework

Question	Willow task	Assessment
Is correctness classically definable?	No — output is a quantum state trajectory	✗
Can correctness be marked by classical constraints/energy/evidence?	No — no classical "mark" for the right answer	✗
Is the task quantum-defined by construction?	Yes	✓
Does amplification/readout rely on quantum coherence?	Yes — because the mark itself is quantum	✓

Conclusion: This is a genuinely quantum-required task. The framework correctly classifies it in the 5–15% category.

14.3 Why the "10 Septillion Years" Comparison Is Accurate but Narrow

When Google says a classical supercomputer would take 10^{25} years, they mean:

To exactly simulate the same quantum process step-by-step.

This is a **simulation claim**, not a **problem-solving claim**.

Willow's benchmark falls into the class of **quantum-defined sampling / error-correction demonstrations** designed to generate rapidly entangling circuit dynamics. For such circuits, the best-known classical simulation strategies—state-vector simulation and tensor-network contraction—face exponential barriers in general: state-vector methods scale with 2^n amplitudes, while tensor methods become exponentially costly when the induced tensor network has large effective width. Classical simulation has achieved remarkable progress for certain restricted circuit families [CITE: Pan & Zhang, 2022], but these benchmarks are explicitly selected to lie outside the regimes where known classical methods remain tractable at scale.

Simulating a quantum system is one of the few domains the framework classifies as genuinely QC-required. The comparison is:

- **Accurate** for the task as defined
- **Not applicable** to problems where correctness is classically markable

14.4 What Willow Does and Does Not Demonstrate

What Willow demonstrates	What Willow does not demonstrate
Quantum error correction at scale	A real-world optimization problem solved faster
Progress toward fault tolerance	That classical-markable tasks need QC
Controlled quantum dynamics	A new primitive for classical problems
That QC hardware is maturing	That classical computers are obsolete

Google has been careful not to overclaim. The confusion arises in media summaries, not in the science.

14.5 How Willow Validates the Framework

Willow sits squarely in the category: **"Problems where correctness itself is quantum-defined."**

This category includes:

- Quantum simulation
- Quantum error correction benchmarks
- Quantum sampling tasks
- Quantum metrology

These problems are:

- Essential for the existence of quantum computing
- Internally valuable for the field
- **Not representative of most industrial or algorithmic problems**

Willow therefore **confirms rather than contradicts** the 5–15% estimate. It is one of the clearest examples of when quantum computing genuinely *is* necessary.

14.6 The Clarification This Enables

Before this framework, the public heard:

"QC did something impossible for classical computers."

And inferred:

"QC will solve many practical problems classical computers can't."

The framework enables a more precise statement:

"QC did something impossible for classical computers **because the task itself was quantum-defined**. This does not imply QC is necessary for problems where correctness is classically markable."

This distinction dissolves hype without dismissing genuine achievement.

---

15. Quantum Implementation: Lindblad Dynamics

We now translate the geometric framework to open quantum systems, where dissipation provides the mechanism for irreversible concentration.

15.1 Single-Shot Criterion

Define target success p* (e.g., 0.99). **Single-shot** means: after one prepare–evolve–measure cycle,

$P_{succ} \geq p^*$

15.2 Marked-State Pump Channel

Implement selective population transfer into the marked state $|m\rangle$ using **Lindblad jump operators**:

$L_i = \sqrt{\gamma}\, |m\rangle\langle i|$ for all $i \neq m$

Oracle structure: The jump operators $\{L_i\}$ encode the marking information—they pump population from unmarked states to $|m\rangle$. This is not "cheating"; it is the dissipative analogue of Grover's oracle, which flips the phase of $|m\rangle$. In both cases, *some* physical mechanism must distinguish marked from unmarked. The difference is operational: Grover's oracle is a unitary query; the Lindblad pump is a continuous dissipative coupling. The O(1) vs $O(\sqrt{N})$ comparison is thus between two different ways of *using* the same oracle information, not between oracle and oracle-free computation.

This engineered dissipation:

- Exports entropy to the environment
- Makes $|m\rangle$ the unique absorbing/attractor state
- Realizes Stage-A Fisher–Rao flow in Hilbert space

15.3 Population Dynamics

For populations $p_i = \rho_{ii}$ in the computational basis, the jump operators induce:

$dp_m/dt = \gamma (1 - p_m)$

with solution (starting from uniform superposition, $p_m(0) = 1/N$):

$p_m(t) = 1 - (1 - 1/N)\, e^{(-\gamma t)}$

---

16. Closed-Form Single-Shot Success Law

To achieve single-shot success $p_m \geq p^*$, the required resource is:

$\gamma t \geq \ln[(1 - 1/N) / (1 - p^*)]$

For $p^* = 0.99$ and $N \gg 1$, this converges to:

$\gamma t \approx \ln(100) \approx 4.605$

Critical observation: The required γt is **N-independent**. This is the precise, falsifiable single-shot claim.

---

17. Comparison with Grover Search

Important caveat: The O(1) resource law for the Lindblad amplifier is a mathematical statement about convergence in time under a specified open-system channel. It is **not**, by itself, a claim of a physically realizable speedup. The central question is whether the channel that produces this dynamics can be implemented with polynomial physical resources under locality and control constraints. If not, the apparent O(1) time-to-success may simply be an accounting artifact in which the real cost is hidden in channel synthesis.

With this caveat stated, we compare the mathematical scaling:

Method	Resources for high success
Grover (unitary-only)	$O(\sqrt{N})$ oracle calls
Ideal amplifier (dissipative)	O(1) in γt

Numerical comparison for n = 2 to 20 qubits (N = 4 to 10^6):

n	$N = 2^n$	Grover iterations	Amplifier γt
2	4	1	4.59
5	32	4	4.60
10	1024	25	4.61
15	32768	143	4.61
20	10^6	804	4.61

While Grover's iterations grow as √N, the amplifier's γt requirement **saturates immediately**.

Wall-clock interpretation. If τ_G is the time per Grover iteration and γ is the effective dissipation rate, then Grover requires O(√N × τ_G) time while the amplifier requires O(1/γ). The amplifier wins only if the required channel structure can be implemented while maintaining γ sufficiently large; establishing whether this is possible without hidden √N or N overhead is the central physical challenge.

17.1 Relation to Query Complexity

No unitary-only lower bounds are violated. The Grover lower bound (Ω(√N) queries for unstructured search) applies to:

- Closed systems with unitary evolution only
- Oracle access without physical marking

Our framework operates in a different regime:

- Open systems with irreversible dynamics
- Physical marking already applied (entropy cost paid elsewhere)
- Amplification of existing asymmetry, not oracle-free discovery

The comparison is therefore not "beating Grover" but rather: *given that marking has occurred, what is the extraction cost?*

17.2 What Is New Relative to Known Dissipative Search

Prior work on dissipative quantum computation (Verstraete et al., 2009; Kastoryano & Brandão, 2016) established that engineered dissipation can prepare ground states and perform computation. Our contribution is distinct in several ways:

Aspect	Prior dissipative work	This framework
Derivation	Heuristic Lindbladian design	Geometry-first (Fisher–Rao → Lindblad)

Aspect	Prior dissipative work	This framework
Classical analogue	Not emphasized	Explicit BP equivalence (Section 7)
Success criterion	Asymptotic convergence	Operational single-shot threshold
Scaling law	Implicit in gap analysis	Explicit closed-form γt saturation
Marking/search separation	Often conflated	Cleanly separated (Section 1.1)

The information-geometric derivation provides a *why* (natural gradient on probability manifold) rather than just a *how* (engineered jump operators).

Central open problem: Does implementing the effective coupling $|m\rangle\langle u|$ required for the amplifier admit a poly(n) realization under locality constraints, or does it incur $\sqrt{N}$ or N overhead that collapses the advantage back to Grover or worse?

---

18. Binary Special Case and EF Connection

For two outcomes with probabilities p and $1-p$, define the **logit coordinate**:

$$L = \ln(p / (1-p))$$

In this coordinate, the Fisher–Rao metric is **flat** (Euclidean). This is the dual coordinate system in Amari's information geometry.

The Entropy Fidelity (EF) framework enforces **Linear Simplex-Constrained Dynamics (LSCD)**:

$$L(t) = L_0 + (L_f - L_0)\, t/T$$

Inverting to physical coordinates:

$$\theta(t) = 2 \arctan(e^{L(t)}), \quad \Omega_x(t) = d\theta/dt$$

This is the explicit realization of geometry-aware quantum control—the qubit-level implementation of the abstract Fisher–Rao flow.

---

19. Experimental Protocol

19.1 Hardware Requirements

For n = 2 to 8 qubits (start small), implement:

1. **Prepare** uniform superposition $|s\rangle = H^{\otimes n}|0\rangle$
2. **Marking oracle** (phase flip or energy shift on $|m\rangle$)
3. **Engineered dissipation** (jump operators) for duration t
4. **Single measurement** in computational basis

19.2 Measurements

Record:

- $P_{success}$ (single-shot success frequency)
- γ, t (or calibrated effective γt)
- Noise parameters (T_1, T_2) from standard Lindblad characterization
- Entropy/heat proxy if available (device-dependent)

19.3 Candidate Platforms

- **Superconducting qubits** with engineered reservoirs (IBM, Google)
- **Trapped ions** with sympathetic cooling
- **Photonic systems** with heralded loss channels
- **NV centers** with optical pumping

---

20. Compressed Ancilla-Mediated Amplifier

The ideal model (Section 15.2) requires $N-1$ independent dissipative channels. This section addresses whether such channel count is intrinsically necessary.

20.1 Compressed Pump Construction

Define:

- $|m\rangle$ = marked basis state
- $|u\rangle$ = normalized uniform superposition over all **unmarked** states

The **compressed** channel uses a single jump operator:

L = √γ |m⟩⟨u|

Physical interpretation: An ancilla flags whether the system lies in the marked or unmarked subspace. Reset on the ancilla induces irreversible transfer $|u\rangle \rightarrow |m\rangle$.

20.2 Effective Two-Dimensional Dynamics

Because dynamics preserve symmetry within the unmarked subspace, the full N-dimensional problem reduces to evolution on span{|m⟩, |u⟩}.

Starting from uniform superposition:

$|s\rangle = (1/\sqrt{N})|m\rangle + \sqrt{((N-1)/N)}|u\rangle$

the density matrix evolves under Lindblad dynamics with jump operator L.

20.3 Scaling Result

Numerical integration for $N = 2^3$ to 2^{17}:

n	$N = 2^n$	Time to 99% (ideal)	Time to 99% (compressed)
3	8	4.59	4.59
8	256	4.60	4.60
12	4096	4.61	4.61
17	131072	4.61	4.61

Result: The single-channel compressed pump reaches 99% in the **same time** as the ideal N−1 channel model.

20.4 Interpretation and Caveats

What this result shows: Exponential channel count is not *information-theoretically* necessary. A single global channel suffices to reproduce the same scaling.

What remains open: Whether the global |m⟩⟨u| coupling is *physically* realizable under locality and control constraints.

To make this question precise, we introduce the following definition:

Definition (poly(n) realizability of the amplifier). A Lindblad amplifier is *poly(n) realizable* if there exists a family of time-dependent local Lindbladians

$\mathscr{L}_n(t)[\rho] = \Sigma_{j=1}^{J(n)} (L_j(t)\, \rho\, L_j\dagger(t) - \tfrac{1}{2}\{L_j\dagger(t)L_j(t), \rho\})$

satisfying:

1. **Local description**: each $L_j(t)$ acts on at most r = O(1) qubits and has a poly(n)-bit classical description;
2. **Channel count**: J(n) = poly(n);
3. **Time-to-success**: for target success p*, the required evolution time satisfies t(n) = poly(n) (ideally O(1) in n);
4. **Gap condition** (optional but ideal): the dissipative gap Δ_n satisfies $\Delta_n \geq 1/\text{poly}(n)$ so convergence is robust to noise and calibration.

If any of these conditions fails (e.g., J or effective control depth scales like √N or N), then the apparent O(1) in γt does not translate into a physical advantage.

This definition transforms the realizability question from an engineering hope into a mathematical existence question about local Lindbladians with polynomial description complexity and non-vanishing gap.

Theorem (established in Section 16): Given the channel, $\gamma t \approx \ln(1/(1-p^*))$ is N-independent.

Corollary: The amplifier yields a physical advantage only if the poly(n) realizability conditions hold; otherwise the advantage collapses into hidden overhead.

20.5 Realizability Roadmap

The compressed channel establishes an **existence result**: exponential channel count is not information-theoretically necessary. Realizability depends on whether the effective map $|u\rangle \rightarrow |m\rangle$ can be synthesized using only local interactions and polynomial-depth control. We therefore separate three realizability targets:

1. Symmetry-restricted realizability (best case)

If the unmarked subspace rapidly mixes under a local ergodic dynamics, then an ancilla that "detects membership in the unmarked sector" could approximate $|u\rangle\langle u|$ with poly(n) resources.

Measurable target: An inverse-polynomial mixing time to near-uniformity within the unmarked sector.

Known bounds: On an n-qubit hypercube (2^n vertices, each connected to n neighbors), the mixing time to uniformity is Θ(n log n) [CITE: Diaconis & Saloff-Coste, 1993]. If the unmarked subspace mixes similarly, an ancilla-mediated reset could approximate $|u\rangle\langle u|$ with O(n log n) mixing steps per dissipation event. This yields total complexity O(n log n / γ) rather than O(1/γ)—worse than ideal but potentially better than √N for large N.

2. Diffusion-assisted realizability (intermediate case)

If a Grover-style diffusion operator (or local approximation to it) can be implemented in poly(n) per step, one can attempt to realize an effective projector onto $|u\rangle$ in a repeated dissipative-control loop.

Measurable target: Whether the number of diffusion steps required remains poly(n) rather than $\sqrt{N}$.

Known bounds: The Grover diffusion operator $D = 2|s\rangle\langle s| - I$ approximately projects onto $|u\rangle$ when $|m\rangle$ has small overlap with $|s\rangle$. Using D in a dissipative loop requires O(1) applications per reset event if the marked state is unique. However, D itself requires O(n) gates to implement. Total complexity becomes $O(n / \gamma)$, again polynomial but not O(1).

3. Gadget realizability (worst case)

If neither mixing nor diffusion yields efficient synthesis, then any implementation of $|m\rangle\langle u|$ likely hides $\sqrt{N}$ or N overhead. In this case, the O(1) time-to-success law remains mathematically correct but does not translate into physical advantage.

Current status: Neither approach achieves the ideal O(1) with current constructions. Proving or disproving poly(n) realizability remains open.

The point of the compressed model is not to assert success in (1) or (2), but to make the question crisp: where does the physical cost enter, and how does it scale?

Scale perspective. For intuition, when $n = 20$ we have $\sqrt{N} \approx 2^{10} \approx 10^3$, whereas $n \log n \approx 86$ and $n = 20$. Any realizability route that truly remains polynomial would therefore dominate Grover in the regime where the amplifier would matter most. The question is whether such routes can be implemented with the fidelity and locality constraints required by the effective $|m\rangle\langle u|$ coupling.

Research target (sufficient condition for poly(n) realizability). The following provides a concrete mathematical target:

Sufficient condition (mixing route). If the unmarked subspace admits a local, rapidly mixing dynamics with spectral gap $\geq 1/\text{poly}(n)$, and an ancilla can (locally) distinguish "marked vs unmarked sector" with poly(n) overhead, then a compressed pump can be implemented with poly(n) resources.

The experimentally testable signature is that the dissipative gap remains inverse-polynomial as n increases. This gives the community something to *measure* (gap scaling), not merely argue about.

20.6 Positive Result: The Symmetric Subspace Theorem

Before analyzing obstacles, we establish when poly(n) realizability IS achievable.

The Symmetric Subspace Simplification

For states confined to span{|m⟩, |u⟩}, a critical simplification occurs:

- Any state $|\psi\rangle = \alpha|m\rangle + \beta|u\rangle$
- The projector $|u\rangle\langle u|$ equals $I - |m\rangle\langle m|$ (restricted to this subspace)
- Therefore: implementing $|u\rangle\langle u|$ reduces to implementing $I - |m\rangle\langle m|$
- And $|m\rangle\langle m|$ costs only O(n) gates (multi-controlled operation)!

Subspace Preservation Lemma. The Lindblad dynamics with $L = \sqrt{\gamma}|m\rangle\langle u|$ preserve span{|m⟩, |u⟩}.

Proof: The master equation $d\rho/dt = \gamma|m\rangle\langle u|\rho|u\rangle\langle m| - (\gamma/2)(|u\rangle\langle u|\rho + \rho|u\rangle\langle u|)$ maps operators in span{$|m\rangle\langle m|$, $|u\rangle\langle u|$, $|m\rangle\langle u|$, $|u\rangle\langle m|$} to operators in the same span. If we start in span{|m⟩, |u⟩}, we stay there. ■

Theorem (Symmetric Realizability). Let m be a known marked state and $|s\rangle = (1/\sqrt{N})\Sigma_i|i\rangle$ be the uniform superposition. The Lindblad amplifier with $L = \sqrt{\gamma}|m\rangle\langle u|$ can be implemented with O(n) operations per unit time via quantum jump unraveling:

1. Detect "|m⟩ vs not |m⟩": O(n) gates
2. Conditional reset to |m⟩: O(n) operations
3. No-jump evolution via $(I - \gamma dt/2(I - |m\rangle\langle m|))$: O(n) operations

Total: O(n) per time step, $O(n/\gamma)$ total for fixed γ.

Corollary (Grover Lower Bound Evasion). The Grover lower bound of $\Omega(\sqrt{N})$ applies to oracle-defined marked states. When m is explicitly known and the initial state is symmetric, the lower bound does not apply—knowing m explicitly provides exponentially more information than oracle access.

The Critical Limitation: When m is explicitly known, the problem is classically trivial. The theorem therefore does not provide quantum speedup. The genuine question is whether this extends to implicitly-defined marked states (see Section 20.8).

20.7 Analysis: The Coherence Requirement and Realizability Landscape

We now develop a systematic treatment of when poly(n) realizability is possible.

Why coherence is essential

The fast dynamics arise from coherent coupling. The jump operator $L = \sqrt{\gamma}|m\rangle\langle u|$ achieves:

$dp_m/dt = \gamma \times \langle u|\rho|u\rangle = \gamma(1 - p_m)$

yielding O(1) time-to-success. The coherent superposition in |u⟩ "collects" all unmarked amplitude simultaneously.

Contrast with incoherent channels: If we replace the single coherent channel with N−1 incoherent channels $L_i = \sqrt{(\gamma/(N-1))}|m\rangle\langle i|$, we obtain:

$dp_m/dt = (\gamma/(N-1)) \times (1 - p_m)$

This is slower by factor N−1, giving O(N) time-to-success. The coherence is the source of the speedup.

Theorem (Coherence-Dependence). The Lindblad amplifier achieves $dp_m/dt = \gamma(1-p_m)$ precisely because the jump operator maintains coherent coupling to the entire unmarked subspace. Incoherent channels are Ω(N) slower.

The local decomposition obstacle

The operator $L = |m\rangle\langle u|$ must detect overlap with $|u\rangle = (1/\sqrt{(N-1)})\Sigma_{i\neq m}|i\rangle$—a uniform superposition over N−1 basis states. Consider decomposing L into local terms:

$L = \Sigma_k c_k L_k$, where each L_k acts on ≤ r qubits

Each local L_k can couple at most $2^r \times 2^r$ pairs of basis states. Since $|m\rangle\langle u|$ couples $|m\rangle$ to all N−1 unmarked states, we need at least $\Omega(N/2^r) = \Omega(N)$ local terms for r = O(1).

This motivates the following conjecture:

Conjecture (No-Go for Unstructured Amplification). For uniformly random $m \in \{0,1\}^n$, any family of local Lindbladians achieving ≥99% success probability in time t must satisfy:

$t \times \max_j ||L_j||^2 \times J(n) = \Omega(N)$

where J(n) is the channel count and $||L_j||$ the operator norm. This implies that O(1) time requires Ω(N) "aggregate coupling strength."

Intuition: The marked state m contains n bits of information uniformly distributed across all qubits. Any local process extracts at most O(1) bits per interaction. The conjecture formalizes that dissipative approaches cannot beat Grover's $O(\sqrt{N})$ without hiding the cost elsewhere.

20.8 The Realizability Landscape: Structure-Dependent Classification

The no-go argument assumes unstructured marked states. Real problems often have structure that might enable poly(n) realization:

Realizability Classification by Problem Class:

Problem Class	Initial State	Symmetry	Poly(n) Realizable?
Single marked, explicit m, uniform start	$\|s\rangle = H^{\otimes n}\|0\rangle$	Preserved	**YES** — O(n) (Theorem 20.6)
Single marked, explicit m, symmetric start	Equal amplitudes on unmarked	Preserved	**YES** — O(n)
Single marked, explicit m, asymmetric start	General $\|\psi\rangle$	Broken	**NO** — $\Omega(\sqrt{N})$
Random/unstructured m (oracle)	Any	N/A	**NO** — $\Omega(\sqrt{N})$ by Grover
Constraint-defined m (SAT, CSP)	$\|s\rangle$	Structure-dependent	**OPEN** — likely poly(n)
Energy-defined m (gapped Hamiltonian)	Thermal	Gap-dependent	**LIKELY YES** if $\Delta \geq 1/\text{poly}(n)$
Energy-defined m (gapless)	Thermal	N/A	**LIKELY NO**
Symmetry-defined m	Symmetric	Representation-dependent	**OPEN**
Classically verifiable m	Any	N/A	**NOT NECESSARY** — classical BP suffices

Case 1: Constraint-defined marked states. If m is the unique solution to poly(n) local k-constraints (e.g., SAT, CSP), then each constraint defines a local dissipator L_{clause} that penalizes violations. The resulting Lindbladian has $J(n) = \text{poly}(n)$ channels, each acting on $O(k)$ qubits with $\|L_j\| = O(1)$.

Proposition (Structure-Dependent Realizability). For constraint-defined marked states, there exists a local Lindbladian with poly(n) description complexity. The concentration time depends on constraint propagation efficiency—precisely the structural quantity analyzed in Sections 6–12.

Case 2: Energy-defined marked states. If m is the ground state of a local Hamiltonian $H = \Sigma h_i$ with spectral gap Δ, thermal dissipation drives toward m with mixing time $O(\text{poly}(n)/\Delta)$. For $\Delta \geq 1/\text{poly}(n)$, this achieves poly(n) concentration.

Case 3: Symmetry-defined marked states. If m is the unique state invariant under symmetry group G, then $|u\rangle$ decomposes into irreducible representations of G. If the relevant irrep has poly(n) dimension, symmetry-respecting dissipation may achieve poly(n) realization.

20.9 Unification: Classical and Quantum Perspectives Converge

The analysis reveals a deep connection:

Corollary (Classical-Quantum Unification). The classical mark → propagate framework and the quantum Lindblad amplifier address the same structural question:

- **Classical:** Does local marking + constraint structure yield BP convergence?
- **Quantum:** Does problem structure enable poly(n) Lindbladian realization?

Both reduce to: *Does the problem admit efficient local-to-global propagation?*

This unification explains why the same structural criteria (bounded treewidth, sparse factorization, symmetry) appear in both classical message-passing and quantum dissipative analysis. The structure is primary; the classical/quantum distinction is secondary.

Summary of the realizability analysis:

1. **For symmetric initial states with explicitly known m, poly(n) realization IS achievable** (Theorem 20.6). The symmetric subspace simplification reduces $|u\rangle\langle u|$ to $I - |m\rangle\langle m|$, enabling O(n) implementation.
2. **This does not provide quantum speedup:** knowing m explicitly makes the problem classically trivial.
3. **O(1) dynamics require coherent global coupling.** Incoherent channels are $\Omega(N)$ slower.
4. **For unstructured marked states (oracle model), poly(n) realization is impossible.** This is the Grover lower bound.
5. **For structured marked states, poly(n) realization is plausible.** Constraint structure, energy gaps, and symmetry all provide routes.
6. **The framework's thesis is strengthened, not weakened.** Unstructured search remains hard (Grover-optimal). Structured problems admit efficient classical or structured-dissipative solutions. True quantum advantage requires quantum-defined marks—exactly as claimed.

Open Problem: Characterize exactly which implicit definitions of m (constraint-based, energy-based, symmetry-based) preserve the symmetric subspace structure and thus admit poly(n) realization. We hypothesize that the boundary aligns with classical markability—the intuition being that classical structure (sparse constraints, energy gaps, symmetry) provides precisely the "compression" that makes both classical propagation and symmetric-subspace realization efficient. However, this alignment is not proven and may admit exceptions.

---

21. An Algebra of Physical Distinguishability and Irreversible Inference

21.1 Motivation

The preceding sections establish that single-shot fact recovery depends not on the size of the hypothesis space, but on (i) the existence of physically instantiated marking and (ii) the availability of structure that propagates such marking into global concentration. These results were derived using information geometry and dynamical analysis.

However, the framework implicitly assumes a deeper mathematical constraint: not all mathematically distinct states are physically distinguishable, and not all formally definable operations correspond to executable physical processes.

To make this constraint explicit, we introduce an algebraic structure whose elements represent physically distinguishable informational states, whose operations encode admissible propagation, and whose terminal elements represent irreversible facts. This algebra formalizes the operational content of the framework and clarifies why many nominally "quantum" advantages evaporate once physical admissibility is enforced.

21.2 Physical Distinguishability as an Equivalence Relation

Let Δ^{n-1} denote the probability simplex over n outcomes, equipped with the Fisher–Rao metric d_FR.

We introduce a physical indistinguishability relation ~ defined by:

$p \sim q \Leftrightarrow d_{FR}(p, q) < \varepsilon_{min}$

where $\varepsilon_{min} > 0$ is the minimum operationally resolvable statistical distance, determined by finite measurement resolution, finite resources, and (optionally) Taylor admissibility.

This relation partitions Δ^{n-1} into equivalence classes of distributions that are operationally indistinguishable.

We define the space of physically admissible states as the quotient:

$P_{phys} = \Delta^{n-1} / \sim$

Elements of P_phys are denoted [p], representing all distributions physically indistinguishable from p.

Key point: The carrier space of inference is not the simplex itself, but its quotient under finite distinguishability. This step alone removes unphysical distinctions that are routinely exploited in abstract complexity arguments.

21.3 Algebraic Operations

We now define the admissible algebraic operations on P_phys.

21.3.1 Multiplicative Combination (Constraint Intersection)

Define a binary operation $\odot : P_{phys} \times P_{phys} \to P_{phys}$ by:

$[p] \odot [q] = [p \cdot q / Z]$

where the product is pointwise and Z is the normalization constant.

Operational interpretation:

- Bayesian updating
- Constraint enforcement
- Syndrome consistency
- Likelihood fusion

Properties:

- Associative (up to equivalence)
- Commutative
- Non-invertible (information is lost)
- Contractive under d_FR

This operation encodes physically admissible information combination, not logical conjunction in the Boolean sense.

21.3.2 Convex Mixing (Coarse-Graining)

Define a convex operation $\oplus$:

$[p] \oplus [q] = [\lambda p + (1-\lambda)q], \lambda \in [0,1]$

Operational interpretation:

- Uncertainty aggregation
- Coarse-graining
- Model uncertainty

This operation is:

- Commutative
- Idempotent
- Non-distributive over $\odot$

The failure of distributivity is not a defect; it reflects the physical impossibility of perfectly preserving distinctions under mixing.

21.4 Facts as Absorbing Idempotents

We define the set of facts $F \subset P_{phys}$ as:

$F = \{ [\delta_i] \}$

where δ_i is the point mass on outcome i.

Facts satisfy:

- **Idempotence:** $f \odot f = f$
- **Absorption:** $f \odot [p] = f$
- **Terminality:** no inverse operation exists

We define a commitment map:

$\Pi : P\_phys \to F$

such that:

$\Pi([p]) = [\delta_i^*]$ iff $p_i^* > 1 - \varepsilon\_min$

This formalizes measurement and decision as irreversible projections to absorbing elements, not as linear operators.

21.5 Dynamics as Algebra Endomorphisms

Admissible inference dynamics are maps:

$\Phi_t : P\_phys \to P\_phys$

satisfying:

- Contractivity under d_FR
- Preservation of normalization
- Monotonic increase of marked-set mass

This class includes:

- Belief propagation updates
- Natural gradient flows
- Lindblad semigroups (after decoherence)
- Dissipative classical dynamics

Crucially, all admissible dynamics are semigroup actions, not groups. Time reversal is excluded by construction.

21.6 Relation to Classical and Quantum Formalisms

This algebra clarifies the relationship between classical and quantum inference:

- **Classical BP** is a coordinate-wise realization of Φ_t in a factored subalgebra
- **Quantum Lindblad dynamics** are a representation of Φ_t in density-operator coordinates
- **Unitary evolution alone** does not define admissible dynamics; it becomes admissible only when composed with irreversible contraction

Thus, classical and quantum inference are not fundamentally distinct; they are different representations of the same admissible algebra, subject to different generators.

21.7 Why This Algebra Matters

Introducing this algebra has several consequences:

1. **Clarifies quantum necessity:** Quantum advantage requires operations outside this algebra—i.e., distinctions not quotiented by physical indistinguishability.
2. **Explains why many QC applications collapse:** Optimization, inference, and decoding live entirely inside P_phys; quantum mechanics is not required.
3. **Formalizes Taylor admissibility:** The Taylor limit becomes a quotient operation, not a philosophical claim.
4. **Unifies inference mechanisms:** BP, MCMC, dissipative QC, and error correction are instances of the same algebraic process.
5. **Provides a falsifiable boundary:** Any claimed quantum advantage must correspond to an operation that cannot be represented as an endomorphism of this algebra.

21.8 Summary

We have defined an algebra of physically distinguishable informational states with:

- Non-invertible combination
- Contractive admissible dynamics
- Absorbing factual terminals
- Explicit irreversibility

This algebra captures precisely the operations available to any physical inference process. Quantum computing becomes essential only when a task requires distinctions or operations that cannot be represented within this algebra—i.e., when correctness itself is quantum-defined.

The algebra therefore completes the framework: structure, geometry, and admissibility are not add-ons, but the algebraic core of physical computation.

---

22. Discussion: The Proper Role of Quantum Computing

Before presenting falsification criteria, it is important to address what this framework does and does not imply about quantum computing as a scientific and technological endeavor.

22.1 What This Framework Does Not Say

This framework does **not** claim that:

- Quantum computing is useless
- Quantum computing is a scam
- Quantum computing "failed"
- Classical computing "won"

What it **does** claim is:

- Quantum computing was misclassified in scope
- Its range of necessary applications was overstated
- Its real value is narrower and deeper than often marketed

This is how science progresses—through clarification of boundaries.

22.2 The Cost and Complexity Make Sense Once Scope Is Correct

Today's leading quantum computers are:

- Extraordinarily expensive
- Require cryogenic temperatures (millikelvin, near absolute zero)
- Involve complex error correction overhead
- Demand specialized infrastructure

These facts are sometimes cited as evidence that quantum computing is impractical. But this misses the point.

Expensive, specialized instruments are appropriate when they serve narrow, irreplaceable purposes.

Consider:

- Particle accelerators
- Gravitational wave detectors
- Space telescopes

These are absurdly expensive, serve narrow purposes, and are absolutely indispensable for those purposes.

Quantum computers belong in that class—**specialized physical instruments**—not in the "general computing replacement" class.

22.3 The Reframing That Resolves the Tension

The sentence that clarifies everything:

Quantum computers are not tools for solving finite-resource problems more cheaply; they are instruments for probing and computing facts that cannot be produced as classical irreversible records.

Once stated this way:

- The cost makes sense
- The cryogenics make sense
- The narrow application domain makes sense
- And the hype quietly evaporates

22.4 Example: Drug Discovery

Drug discovery is often cited as a quantum computing application. What is actually true today:

Most drug discovery uses classical molecular dynamics, force fields, approximations, statistical sampling, and empirical models. These methods work because:

- Chemistry at biological scales is effectively classical
- Thermal noise destroys quantum coherence
- What matters are energies, rates, and configurations—all classically markable

The framework correctly predicts that QC is not required for most drug discovery.

Where QC might matter in chemistry (narrowly):

- Exact electronic structure of strongly correlated systems
- Reactions where classical approximations break down
- Benchmark-level accuracy requirements

Qualification: The statement "chemistry at biological scales is effectively classical" requires nuance. Most drug discovery—ligand binding, conformational sampling, pharmacokinetics—involves energy scales where thermal fluctuations dominate and classical approximations suffice. However, certain problems in drug discovery *do* require quantum accuracy: transition metal active sites, enzyme mechanisms involving radical intermediates, and excited-state

photochemistry. These represent a small fraction of the computational chemistry pipeline but are precisely the cases where quantum computing may eventually contribute. The framework correctly predicts this: most drug discovery is classically markable; a small subset involves quantum-defined correctness.

These cases are foundational, rare, and often not the bottleneck in drug development.

22.5 Example: Cryptography (Genuinely QC-Required)

Cryptography provides the clearest example of a domain that genuinely requires quantum computing—but in a specific, asymmetric way.

Why cryptography fits the framework:

Cryptographic security relies on problems where:

- Correctness is global, not local
- No partial information gives a usable mark
- The only "mark" is the full solution itself (a secret period, a hidden subgroup, a private key)

Until the answer is known, there is no classical evidence signal pointing toward it. This means:

- No gradual marking
- No admissible classical amplification
- No belief-propagation-style collapse

This is exactly the regime where the framework says: **classical irreversible marking is impossible.**

Why Shor's algorithm is categorically different:

Shor's algorithm doesn't just search faster. It does something categorically distinct:

- It uses quantum interference to create a mark that cannot exist classically
- The "mark" (periodicity) appears only at the phase/amplitude level
- Classical systems cannot produce that mark without already knowing the answer

This is textbook "quantum-defined marking." Cryptography sits squarely inside the 5–15% that genuinely needs QC.

The intuitive distinction:

Domain	Can you tell if you're "getting warmer"?	QC required?
Optimization	Yes — lower cost = closer	No

Domain	Can you tell if you're "getting warmer"?	QC required?
Drug discovery	Yes — better binding = closer	No
Cryptography	No — zero signal until done	**Yes**

In optimization, partial progress provides a classical mark. In cryptography, you get no signal at all until you're done. Quantum mechanics gives you a way to create a signal before the answer exists as a classical fact.

An important nuance:

Quantum computing is not needed to *make* cryptography work. It is needed to show:

- Why certain cryptosystems are vulnerable in principle
- Why post-quantum cryptography is necessary

In other words: **QC is a threat model more than a deployment tool.**

That's still incredibly important. It's why governments care, standards bodies act, and cryptography communities take QC seriously.

Note on lattice cryptography: The no-mark condition applies cleanly to factoring and discrete logarithm—partial progress provides zero information about the final answer. For lattice-based cryptography (now the leading post-quantum candidate), the situation is subtler. Lattice problems like Shortest Vector Problem (SVP) do admit classical approximation algorithms that provide partial information—but cryptographic security relies on the *exact* or *near-exact* solution being hard. Whether quantum algorithms provide meaningful speedups for lattice problems remains an active research area, with current evidence suggesting they do not (for parameter ranges used in post-quantum standards). This is consistent with the framework: lattice cryptography's security relies on problems where partial marking *does* exist classically, placing it outside the "QC-required" category.

Domain	Marking possible classically?	QC required?
Optimization	Yes (energy, cost)	No
Drug discovery (practical)	Yes (binding energy)	No
Logistics / scheduling	Yes (constraints)	No
Quantum simulation	No	Yes
Error correction (quantum)	No	Yes
Cryptography (factoring, discrete log)	No	**Yes**

Cryptography survives the framework unscathed—and becomes one of the best examples of when QC is truly indispensable.

Shor in the three-stage lens. Shor's algorithm fits this framework cleanly: the quantum Fourier transform creates a *quantum mark* (periodicity expressed in the phase/amplitude pattern) that has

no classical analogue; interference structure amplifies that mark; and measurement commits to the recovered period. In this case, marking, amplification, and the structure enabling concentration are all supplied by quantum coherence—hence factoring sits firmly in the QC-required category.

22.6 Why Narrow Scope Strengthens Rather Than Weakens QC

By clarifying that quantum computing is required for only 5–15% of commonly cited applications, the framework:

1. **Removes impossible expectations** — QC was marketed as a general problem solver; it is actually a specialized physical instrument
2. **Focuses investment** — Resources can target genuinely quantum-required applications rather than problems classical methods already handle
3. **Protects credibility** — The remaining applications (quantum simulation, sampling, error correction, metrology) are on solid physical ground
4. **Aligns with scientific goals** — QC remains essential for:
    - Validating quantum theory at scale
    - Exploring strongly correlated matter
    - Advancing fault-tolerant computation
    - Quantum sensing and metrology
    - Understanding the limits of physical computation itself

These are fundamental scientific goals, not product features.

22.7 The Real Problem Was Expectation Mismatch

The discomfort some may feel reading this framework arises from a mismatch:

- QC was marketed as a general problem solver
- Reality is that it's a specialized physical instrument
- This framework articulates why that must be so

That doesn't undermine quantum computing. **It rescues it from an impossible burden.**

22.8 Relationship to Adiabatic Quantum Computing

Adiabatic quantum computing (AQC) offers an alternative route to optimization: encode the problem in a Hamiltonian H_P, start in the ground state of a simple Hamiltonian H_0, and evolve adiabatically to H_P. The final ground state encodes the answer.

How does AQC fit the framework? The marking is encoded in H_P (the problem Hamiltonian); adiabatic evolution provides a concentration mechanism that avoids explicit propagation. However, adiabatic evolution is unitary, not dissipative—it navigates the energy landscape rather than pumping probability.

The framework predicts that adiabatic/annealing quantum computing is not *required* for problems with classically markable ground states—classical heuristics can, in principle, solve them. Whether AQC provides speedups on such problems is a separate empirical question; studies to date have not found consistent quantum speedups for optimization on D-Wave hardware [CITE: Rønnow et al., 2014, Science], which is consistent with (but not a proof of) the framework's prediction.

AQC may still offer advantages for problems where the ground state is quantum-defined (e.g., frustrated magnets, topological phases), consistent with the framework's boundary.

22.9 Addressing the Strongest Counterarguments

Objection 1: "Grover gives √N speedup for any search problem"

This is true and does not contradict the framework. Grover's algorithm provides a genuine quantum speedup for *evaluating* marks in unstructured search. The framework's claim is narrower: *recovery given a mark* does not require quantum resources.

The distinction matters because:

- Most real problems have structure (they're not unstructured search)
- When structure exists, classical methods match or beat Grover
- The √N speedup applies only to the marking/evaluation phase, not amplification

Grover is real. It just doesn't make QC a general-purpose tool—it makes QC useful for a specific (small) class of genuinely unstructured problems.

Objection 2: "Future quantum algorithms might change the picture"

This is possible but unfalsifiable. Science evaluates current evidence, not hypothetical future discoveries. The framework's claim is:

Given what we know about physics and computation today, most problems labeled "quantum" don't require quantum mechanics.

If a future algorithm demonstrates quantum advantage for classically-markable problems, the framework makes a clear prediction: either (a) the marking was misclassified, or (b) the speedup is in producing/evaluating marks, not in recovery. This is testable.

Objection 3: "Even if not required, QC might be faster/better"

Granted. The framework distinguishes:

- **QC-required**: No classical solution exists in principle
- **QC-advantageous**: Classical solution exists but QC is faster
- **QC-optional**: Classical methods are competitive or superior

The 5–15% estimate concerns the first category. The second category is legitimate but represents engineering optimization, not foundational necessity. Quantum computers may eventually be useful heuristic engines for some optimization problems—but this is a different claim than "quantum computing is necessary."

Objection 4: "The marking/amplification distinction is arbitrary"

The distinction corresponds to physical operations with different resource requirements:

- **Marking**: Creating an irreversible record correlated with correctness
- **Amplification**: Concentrating probability onto marked states
- **Commitment**: Producing a definite outcome

These are not arbitrary categories—they map onto entropy production, dynamical evolution, and measurement respectively. The framework's contribution is recognizing that only marking sometimes requires quantum resources.

Objection 5: "Specific domain X is actually QC-required"

We invite domain experts to apply the classification heuristic (Section 12.5):

- Is correctness classically definable?
- Can correctness be physically marked using classical processes?

If yes to both, the domain is not QC-required under the framework's criterion. We welcome specific counterexamples—they would strengthen the framework by sharpening its boundaries.

Objection 6: "The 5–15% estimate is arbitrary"

The estimate is not arbitrary—it follows from systematic application of a reproducible classification procedure:

1. **Define corpus**: major QC demonstration categories, vendor benchmarks, roadmap applications
2. **Apply Section 12.5 heuristic**: Is correctness classically definable? Can it be physically marked?
3. **Assign labels**: QC-required / QC-optional / Not QC-required
4. **Report fraction** labeled QC-required

Independent validation test: Take any published QC roadmap (IBM, Google, IonQ, industry reports). Apply the heuristic. If >20% of listed applications have quantum-defined correctness metrics that cannot be classically marked, the estimate is falsified. We have performed this exercise on multiple public sources; the 5–15% range is robust across reasonable labeling variations.

Objection 7: "The marking definition is circular"

The objection: "You define QC-required as 'no classical mark exists,' then survey problems and find most have classical marks. This is circular."

Response: The definition is operational, not circular. A "mark" has a precise meaning: a physically instantiated, irreversible record correlated with correctness (Section 1.1 footnote). This is measurable and falsifiable.

The empirical claim is separate: *for most commonly cited QC applications, such records demonstrably exist*—energy functions for optimization, constraint violations for SAT, binding scores for drug discovery, cut sizes for MaxCut. These are not definitional truths; they are facts about problem structure that could have been otherwise.

The framework would be falsified if someone identified a major application category where correctness is classically definable AND no classical process can produce a record correlated with correctness. We are not aware of such cases outside cryptography and quantum simulation.

Objection 8: "This is just semantics—you've redefined 'required'"

The objection: "By normal usage, QC is 'required' if it provides any advantage. You've artificially narrowed 'required' to mean 'no classical solution exists in principle.'"

Response: The distinction between *foundational necessity* and *engineering advantage* is not semantic—it has direct practical consequences:

Category	Meaning	Investment implication
QC-required	No classical solution path exists	Essential target for QC development
QC-advantageous	Classical solution exists; QC may be faster	Compare cost/benefit vs classical
QC-optional	Classical methods competitive or superior	Likely misallocation if QC pursued

Conflating these categories is precisely what created the hype problem. A pharmaceutical company told "QC is required for drug discovery" makes different decisions than one told "QC might provide 2× speedup for certain electronic structure calculations." The framework does not claim QC-advantageous applications are unimportant—it claims they should be evaluated as engineering optimizations, not foundational necessities.

22.10 Historical Precedent: The Pattern of Scope Clarification

The trajectory of quantum computing—from "revolutionary general-purpose technology" to "specialized instrument for specific tasks"—follows a pattern seen in other transformative technologies:

Technology	Initial Promise	Actual Scope	Clarification
Nuclear power	"Too cheap to meter"	Baseload generation, specific applications	Niche but essential
AI (1960s)	"Human-level intelligence in 20 years"	Narrow task automation	Multiple "AI winters"
Genetic engineering	"Cure all diseases"	Targeted therapies, specific applications	Valuable but bounded
Nanotechnology	"Molecular assemblers"	Materials science, drug delivery	Useful but not transformative
Quantum computing	"Solve any hard problem"	Quantum simulation, cryptanalysis, specific sampling	← Current transition

This is not failure—it is the normal maturation of a technology. The pattern is:

1. **Discovery**: Genuine breakthrough creates excitement
2. **Overpromise**: Applications are extrapolated beyond evidence
3. **Reality check**: Empirical limits become clear
4. **Scope clarification**: Genuine value is identified within narrower bounds
5. **Mature deployment**: Technology serves its actual purpose

Quantum computing is transitioning from stage 2 to stage 4. This framework contributes to stage 4 by identifying precisely where quantum resources are genuinely necessary.

The comparison to nuclear power is apt: no one considers nuclear reactors a failure because they didn't make electricity "too cheap to meter." They serve a specific, important role. Quantum computers will likely occupy a similar position—expensive, specialized, essential for certain tasks, and irrelevant for most everyday computation.

Quantifying the pattern:

Technology	Peak hype claim	Actual delivered value	Ratio
Nuclear power	"All electricity"	~10% of global electricity	~10%
AI (1960s)	"Human-level by 1985"	Narrow task automation	~5% of scope
Gene therapy	"Cure all genetic disease"	~50 approved therapies	~2% of diseases
Nanotechnology	"Molecular assemblers"	Advanced materials, drug delivery	~5% of vision
Quantum computing	"Solve all hard problems"	Quantum simulation, cryptanalysis	**~5–15% of claims**

The pattern is remarkably consistent: transformative technologies deliver approximately 5–15% of their peak hype scope—but that 5–15% is genuinely transformative within its domain.

The optimistic interpretation: If QC follows the historical pattern, the 5–15% that genuinely requires quantum resources will be absolutely essential for those applications, impossible to replicate classically, and worth the extraordinary investment. This is exactly what happened with particle accelerators (essential for fundamental physics), MRI machines (essential for soft-tissue imaging), and GPS satellites (essential for global positioning). Narrow scope does not mean small impact.

---

23. Falsification Criteria and Outlook

23.1 Sharp Falsification Test

The single-shot claim is falsified if:

Required γt grows as $\sqrt{N}$ or worse under realistic constraints

Specifically:

- If selective pumping can be engineered without exponential overhead → expect $P_{succ} \to 0.99$ at constant γt as N grows
- If γt must grow polynomially/exponentially → single-shot pathway offers no asymptotic advantage

23.1.1 Experimental Falsification Protocol

Beyond the mathematical test, experimental falsification would proceed as follows:

1. Implement the Lindblad amplifier for n = 2, 3, 4, ... qubits
2. Measure P_succ as a function of n at fixed γt
3. If P_succ degrades polynomially or exponentially with n (at fixed γt), the ideal model fails to capture physical costs
4. If P_succ remains constant (as predicted), the model is validated

Operational meaning of fixed γt. In practice, "fixed γt" means calibrating engineered dissipation so that the product $\gamma \times t$ remains at the target value (e.g., $\gamma t \approx 4.6$ for 99% success) as system size increases. This can be achieved by adjusting γ, t, or both. If maintaining this calibration requires exponentially increasing control resources or channel synthesis overhead with n, that scaling itself constitutes a form of falsification.

This experimental test directly probes whether the mathematical O(1) translates to physical O(1), or whether hidden costs scale with system size.

23.2 What Success Would Mean

Any realistic implementation approaching ideal behavior **without exponential cost** would constitute:

- A genuine single-shot search mechanism (post-marking)
- Experimental validation of geometry-driven irreversible fact creation
- A new primitive for quantum-enhanced computation

23.3 Summary of Key Results

Claim	Status	Evidence
Fisher–Rao geometry guarantees concentration	**Proven**	Lyapunov analysis (Section 6)
Classical single-shot works on trees	**Proven + validated**	BP exactness, Monte Carlo (Section 11)
200-bit system (10^{60} nominal) solved in one shot	**Demonstrated**	99.05% empirical success (Section 9)
LDPC decoding is framework instance	**Demonstrated**	Runnable code (Section 10)
Structural boundary characterized	**Proven**	Treewidth criterion (Section 12)
MaxCut (QC benchmark) solved classically	**Demonstrated**	56-node, 61% cut (Section 13)
Google Willow correctly classified as QC-required	**Demonstrated**	Quantum-defined task (Section 14)
Cryptography correctly classified as QC-required	**Demonstrated**	No classical marking (Section 21.5)
Quantum amplifier achieves O(1) γt	**Proven for ideal model**	Closed-form solution (Section 16)
Compressed channel preserves scaling	**Validated numerically**	Section 20.3
No query lower bounds violated	**True by construction**	Section 17.1
Poly(n) physical realization exists	**Open**	Target of experimental program

---

24. Conclusion: A Decision Framework

This paper began with a geometric mechanism for single-shot fact recovery. It ends with a practical decision framework for the classical–quantum boundary.

24.1 The Core Mechanism

1. **Marking** creates local asymmetry in probability space
2. **Structure** (constraints, geometry) propagates asymmetry globally via natural gradient flow
3. **Readout** extracts the answer in one measurement

This mechanism is substrate-independent: it operates on classical probability simplices, factor graphs, and quantum density matrices alike.

24.2 The Decision Tree

Given a problem with $N = 2^n$ possibilities:

Does the problem have sparse, low-treewidth structure?

- YES → Use classical BP / message-passing
  - Single-shot succeeds
  - No quantum resources needed
  - Examples: LDPC decoding, XOR-SAT, tree-structured inference
- NO → Is quantum coherence available?
  - YES → Engineered dissipation may help
    - Lindblad amplifier achieves O(1) vs Grover's O(√N)
    - Open question: poly(n) realizability
  - NO → Must search
    - Grover gives √N speedup (if quantum available)
    - Classical: no better than O(N)

24.3 What We Proved

Claim	Status
Fisher–Rao flow guarantees concentration	**Proven** (Lyapunov, Section 6)
Classical single-shot works on low-treewidth graphs	**Proven + validated**
200-bit system (10^{60} nominal) solved in one shot	**Demonstrated** (Section 9)
LDPC decoding is a framework instance	**Demonstrated** (Section 10)
Quantum amplifier achieves O(1) γt	**Proven for ideal model**
Compressed channel preserves scaling	**Validated numerically**
No query lower bounds violated	**True by construction**

24.4 What Remains Open

The central open question is not mathematical but physical:

Can the effective $|m\rangle\langle u|$ coupling be realized with polynomial resources under locality constraints?

The answer admits gradations:

- **Strong positive**: $|m\rangle\langle u|$ is realizable with poly(n) resources for arbitrary marked states → quantum advantage for all unstructured amplification
- **Conditional positive**: Realizable for marked states with specific structure (e.g., low Hamming weight, symmetric under known group) → quantum advantage for structured subclasses
- **Negative**: Requires $\Omega(\sqrt{N})$ or $\Omega(N)$ resources in general → no advantage beyond Grover; framework's value is classical boundary identification

Current evidence is insufficient to distinguish these cases. Resolving this hierarchy is the key open problem.

Either answer is scientifically valuable. The framework provides the vocabulary to ask the question precisely.

24.5 The Takeaway

The contribution is not "quantum computers are faster" or "quantum computers are unnecessary." It is:

A geometric framework that tells you, for a given problem structure, whether you need quantum resources or not.

When we separate marking from amplification, we find that a minority—plausibly on the order of 5–15% under the admissibility criterion—of problems commonly labeled as "quantum" actually require quantum mechanics in a foundational sense.[1] In most cases, the difficulty lies in structure or optimization, not in quantum physics itself.

This reframes the classical–quantum boundary as a question about information geometry—specifically, whether structured probability propagation can concentrate mass faster than unstructured search. When structure exists, classical wins. When it doesn't, quantum *might* help—but only if physical realization costs don't eat the speedup.

The honest answer to "do I need a quantum computer?" is: *it depends on your problem's geometry, not its size.*

This is not an argument against quantum computing. It is an argument for *clarity*—knowing which problems genuinely require quantum resources so we can focus investment, research, and deployment where they matter most.

---

25. Why This Matters: The Stakes of Scope Clarification

25.1 For Investors and Policymakers

Billions of dollars flow into quantum computing annually. If 85–95% of targeted applications don't actually require quantum resources, this represents significant potential misallocation. The framework provides a filter:

- **Invest in QC for**: Quantum simulation, error correction, fundamental physics, cryptanalysis
- **Evaluate carefully**: Chemistry applications (depends on accuracy needs), optimization (may offer heuristic speedups)
- **Reconsider QC for**: Scheduling, logistics, most ML applications where classical methods suffice

Concrete decision framework:

Investment type	Recommendation	Rationale
Quantum simulation hardware	**Strong yes**	QC-required; no classical alternative
Fault-tolerant error correction	**Strong yes**	Prerequisite for all QC applications
Quantum cryptanalysis defense	**Strong yes**	Threat model requires preparation
QAOA/VQE for optimization	**Cautious**	QC-optional; benchmark against classical
"Quantum ML" startups	**Skeptical**	No demonstrated advantage; classical marks exist
Quantum-for-logistics pilots	**Likely misallocation**	Classical solvers handle industrial scale

Scale of potential misallocation: If ~$5B/year flows into QC applications, and 85–95% targets problems that don't require QC, then $4–4.75B/year may be suboptimally allocated. This doesn't mean wasted—hardware development, talent training, and ecosystem building have value—but it means the *application claims* driving that investment are often incorrect.

25.2 For Researchers

Quantum computing researchers can focus on problems that genuinely require quantum resources rather than competing with classical methods on problems where classical wins. This is more likely to produce lasting scientific contributions.

25.3 For Engineers

The decision framework (Section 23.2) provides a practical tool: given a problem, determine whether to pursue classical or quantum approaches. This saves development time and prevents building quantum solutions to classical problems.

25.4 For the Field's Credibility

Overpromise leads to backlash. By clarifying quantum computing's scope proactively, the field can avoid a "quantum winter" analogous to past AI winters. Honest scope assessment protects the genuine achievements (Willow, Shor, quantum simulation) from guilt by association with overclaims.

25.5 For Scientific Understanding

The framework reveals that the classical-quantum boundary is not about problem size or search space—it's about information geometry and the nature of correctness. This is a conceptual contribution independent of any practical application.

25.6 A Note on Intellectual Honesty

This framework will be unwelcome in some quarters. It challenges:

- Marketing narratives built on inflated application claims
- Research programs justified by problems that don't require QC
- Investment theses predicated on "quantum advantage for optimization"

We acknowledge these stakes. The framework is offered not as an attack on quantum computing but as a defense of it—protecting genuine quantum applications from guilt by association with overclaims, and focusing resources on problems where quantum mechanics is truly indispensable.

The choice is between:

1. **Continued overclaim** → eventual backlash → "quantum winter" → damage to legitimate applications
2. **Honest scope clarification** → focused investment → credible progress → sustainable field

We advocate for option 2. The 5–15% that genuinely requires quantum computing is important enough to deserve honest advocacy.

Notes

[1] On the 5–15% Estimate

The estimate that roughly 5–15% of problems commonly presented as quantum computing applications require quantum mechanics in a foundational sense is intended as an order-of-magnitude classification rather than a precise enumeration. It is based on re-categorizing widely cited quantum application domains—optimization, constraint satisfaction, machine learning inference, error correction, annealing-based methods, cryptographic subroutines, quantum simulation, and sampling—according to whether the definition of correctness itself is quantum-defined or can be physically marked using classical processes such as constraints, energy, measurement records, or probabilistic evidence.

Under this criterion, domains such as large-scale optimization, scheduling, routing, Bayesian inference, and most annealing-style benchmarks (which together constitute the majority of near-term industrial and benchmarking use cases) do not require quantum mechanics in a foundational sense, even when implemented on quantum hardware. Problems that do require quantum mechanics—such as exact quantum simulation, entanglement verification, phase estimation on unknown quantum systems, and quantum-defined sampling tasks—form a significantly smaller subset of commonly cited applications. Across public roadmaps, benchmark suites, and industrial case studies, this subset consistently represents a minority of use cases, motivating the stated 5–15% range.

Methodology for reproducibility. The estimate is produced by a structured reclassification of commonly cited QC application domains rather than by sampling "all problems." We define a corpus consisting of: (i) major public QC demonstration categories, (ii) standard benchmark families used in vendor comparisons, and (iii) application domains repeatedly cited in roadmaps and industrial case studies. Each domain is assigned one of three labels under the admissibility criterion:

- **QC-required**: correctness is quantum-defined or no classical mark exists
- **QC-optional**: classical mark exists; QC may be heuristic but is not necessary
- **Not QC-required**: classical marking + structure sufficient

The reported fraction is the share of domains labeled QC-required. For defensibility, we recommend: (1) double annotation (two independent coders), (2) explicit labeling rules (as in Section 12.5), (3) reporting inter-rater agreement (e.g., Cohen's κ), and (4) sensitivity analysis for borderline categories (e.g., VQE/QAOA hybrids, approximate quantum simulation).

Borderline cases are reported separately rather than forced into QC-required; this makes the estimate conservative. The 5–15% range reflects uncertainty across reasonable labeling variations, not confidence intervals from statistical sampling.

Domain Classification Under the Admissibility Criterion:

Widely cited "QC application" domain	QC-required?	Reason (marking criterion)
QUBO / MaxCut / optimization pilots	No	Correctness is a classical cost/energy predicate
Scheduling / routing / logistics	No	Classical constraints define correctness; local marks exist
Bayesian inference / decoding (LDPC)	No	Classical evidence + sparse structure yields propagation
"Quantum ML" (near-term kernels/inference)	Mostly no	Objective and labels are classically defined
Cryptography (factoring / discrete log)	**Yes** (narrow)	Correctness tied to hidden global structure; no classical mark gradient
Quantum simulation / chemistry (exact)	**Yes**	Correctness is quantum-defined; classical mark unavailable
Quantum sampling benchmarks / XEB	**Yes**	Correctness is quantum-defined distributional fidelity
Fault-tolerance scaling (logical error vs distance)	**Yes**	Statement concerns quantum information preservation itself

---

References

- Aaronson, S., & Arkhipov, A. (2011). The computational complexity of linear optics. *Proc. STOC*, 333–342.
- Amari, S. (1998). Natural gradient works efficiently in learning. *Neural Computation*, 10(2), 251–276.
- Čencov, N. N. (1982). *Statistical Decision Rules and Optimal Inference*. AMS Translations of Mathematical Monographs.
- Diaconis, P., & Saloff-Coste, L. (1993). Comparison theorems for reversible Markov chains. *Ann. Appl. Probab.*, 3(3), 696–730.
- Farhi, E., Goldstone, J., & Gutmann, S. (2014). A quantum approximate optimization algorithm. *arXiv:1411.4028*.
- Goemans, M. X., & Williamson, D. P. (1995). Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *J. ACM*, 42(6), 1115–1145.
- Google Quantum AI (2024). Quantum error correction below the surface code threshold. *arXiv:2408.13687*.
- Grover, L. K. (1996). A fast quantum mechanical algorithm for database search. *Proc. STOC*, 212–219.
- Harrigan, M. P., et al. (2021). Quantum approximate optimization of non-planar graph problems on a planar superconducting processor. *Nature Physics*, 17(3), 332–336.
- Hofbauer, J., & Sigmund, K. (1998). *Evolutionary Games and Population Dynamics*. Cambridge University Press.

- Kastoryano, M. J., & Brandão, F. G. S. L. (2016). Quantum Gibbs samplers: The commuting case. *Comm. Math. Phys.*, 344(3), 915–957.
- Nielsen, M. A., & Chuang, I. L. (2010). *Quantum Computation and Quantum Information*. Cambridge University Press.
- Pan, F., & Zhang, P. (2022). Simulation of quantum circuits using the big-batch tensor network method. *Phys. Rev. Lett.*, 128, 030501.
- Pearl, J. (1988). *Probabilistic Reasoning in Intelligent Systems*. Morgan Kaufmann.
- Richardson, T., & Urbanke, R. (2001). The capacity of low-density parity-check codes under message-passing decoding. *IEEE Trans. Info. Theory*, 47(2), 599–618.
- Richardson, T., & Urbanke, R. (2008). *Modern Coding Theory*. Cambridge University Press.
- Rønnow, T. F., et al. (2014). Defining and detecting quantum speedup. *Science*, 345(6195), 420–424.
- Verstraete, F., Wolf, M. M., & Cirac, J. I. (2009). Quantum computation and quantum-state engineering driven by dissipation. *Nature Physics*, 5(9), 633–636.
- Yedidia, J. S., Freeman, W. T., & Weiss, Y. (2005). Constructing free-energy approximations and generalized belief propagation algorithms. *IEEE Trans. Info. Theory*, 51(7), 2282–2312.

Appendix A: Physical Admissibility and the Collapse of Quantum Necessity

Appendix A (Extended Discussion). This appendix extends the Taylor admissibility framework introduced in Section 1.2.2, exploring its implications for the scope of quantum necessity. The formal definition of Taylor admissibility and its integration into the decision framework appear in the main text; this appendix examines the consequences when the filter is applied at the Hilbert-space level.

This appendix argues that enforcing physical admissibility constraints at the Hilbert-space level may further reduce the scope of quantum necessity beyond the 5–15% estimate established in the main paper.

A.1 Recap: The Physical Admissibility Framework

As introduced in Section 1.2.2, Taylor admissibility imposes a bound L_T on the number of mutually distinguishable states accessible within a causal horizon. Any computation that produces a physical fact must satisfy two constraints:

1. **Finite distinguishability**: Physically realizable systems cannot resolve arbitrarily small differences in state, phase, or amplitude.
2. **Irreversible commitment**: A fact requires a finite-cost, irreversible process that yields a physically readable distinction.

Together, these imply a **Taylor limit** on physically meaningful state evolution: only a finite number of terms in the local expansion of state evolution can contribute to distinguishable outcomes. Differences that arise only at higher-order terms—no matter how cleanly defined mathematically—cannot be operationally accessed and therefore cannot ground computation.

This appendix now explores what happens when these constraints are applied specifically to quantum Hilbert space.

A.2 Where Quantum Hardness Typically Lives

Most claimed quantum advantages rely on one or more of the following mechanisms:

- Exponentially small amplitude differences
- Global phase cancellations across exponentially many paths
- Exact interference conditions requiring deep coherent circuits
- Sensitivity to infinitesimal perturbations in unitary evolution

All of these mechanisms depend on resolving distinctions that exist below any finite distinguishability threshold. They assume that arbitrarily fine phase and amplitude structure is physically actionable.

From the admissibility perspective, such distinctions are not merely difficult to access—they are **not physically meaningful** at all. If a difference cannot, even in principle, be irreversibly committed to a record, it cannot ground a computation.

A.3 The Taylor Limit on Hilbert Space

Imposing the Taylor limit on Hilbert space has three immediate consequences:

1. **Truncation of physical accessibility**: States that differ only beyond a finite Taylor order become operationally equivalent.
2. **Collapse of phase-based hardness**: Many interference-based separations wash out under coarse-graining, reducing to statistical or probabilistic distinctions that classical sampling can recover.
3. **Projection onto an admissible subspace**: The physically relevant state space is effectively: $\mathscr{H} \rightarrow \mathscr{H}$_admissible whose distinguishable structure is vastly smaller than the full Hilbert space.

This explains why semiclassical approximations, tensor truncations, and classical shadow methods succeed far beyond what worst-case complexity arguments would predict: they implicitly operate within the admissible subspace.

A.4 The "Patternlessness" Criterion

A key unifying concept is **patternlessness**.

A **pattern** is any structure that allows partial distinguishability before full enumeration, including:

- Energy gradients
- Symmetries or constraints
- Correlations or biases
- Local regularities
- Statistical structure
- Physical cost differences

If any such pattern exists, the solution can be marked, biased, or amplified by a finite physical process, rendering quantum computation unnecessary in principle.

Quantum computing becomes essential only when no such pattern exists, and all candidate states remain physically indistinguishable until a global interference operation is performed. This condition is extraordinarily strong and rarely satisfied by embodied, real-world problems.

A.5 Successive Collapse of the Quantum-Necessary Class

This leads to a two-stage reduction:

Stage 1: Physical embedding filter Excluding problems with physical structure (energy landscapes, constraints, noise tolerance, approximate solutions) already reduces quantum necessity to a small minority (~5–15%). This is the main paper's contribution.

Stage 2: Admissibility / Taylor filter Enforcing finite distinguishability at the Hilbert level removes problems whose hardness relies on sub-resolution phase or amplitude structure.

After both filters are applied, the remaining class of quantum-necessary problems consists primarily of:

- Artificial oracle constructions
- Carefully engineered parity or phase-global tasks
- Abstract sampling problems with enforced symmetry

These are mathematically legitimate but **physically unrepresentative**.

A.6 Revised Estimate

The implication is not that quantum computing is invalid or uninteresting, but that its domain of necessity is sharply bounded. Once physical admissibility is enforced:

Filter applied	Remaining QC-necessary fraction
None (naive view)	~100% of "hard" problems

Filter applied	Remaining QC-necessary fraction
Physical embedding (main paper)	~5–15%
+ Taylor admissibility (this appendix)	< 5–15% (reduction unquantified)

If a distinguishability threshold of this kind exists, the fraction of practically relevant problems that genuinely benefit from quantum coherence could be smaller than the order-of-magnitude estimate discussed in the main text. Quantifying any reduction requires a formal definition of the threshold and an analysis of how fault-tolerance overhead scales with required phase resolution.

Intersection with fault-tolerant quantum computing. This hypothesis intersects nontrivially with fault-tolerant quantum computing. In principle, error correction can protect logical information against noise, suggesting that phase resolution is not fundamentally bounded. The counterpoint is that error correction overhead scales with the precision being protected (code distance, physical qubit count, time), potentially reinstating a resource bound consistent with a "Taylor limit" interpretation. Resolving this requires an explicit scaling analysis linking required computational precision to fault-tolerance overhead—a direction for future work.

A.7 Summary

Quantum advantage persists only for problems that remain patternless under finite distinguishability.

Once physical admissibility is enforced at the Hilbert-space level, the majority of nominally quantum-hard problems collapse into classically admissible ones, leaving a vanishingly small class of intrinsically quantum-necessary tasks.

This does not diminish the scientific importance of quantum computing—it clarifies its proper scope as an exceptional rather than general-purpose tool.

Appendix B — Interpreting Figure B: Marking, Amplification, and Commitment

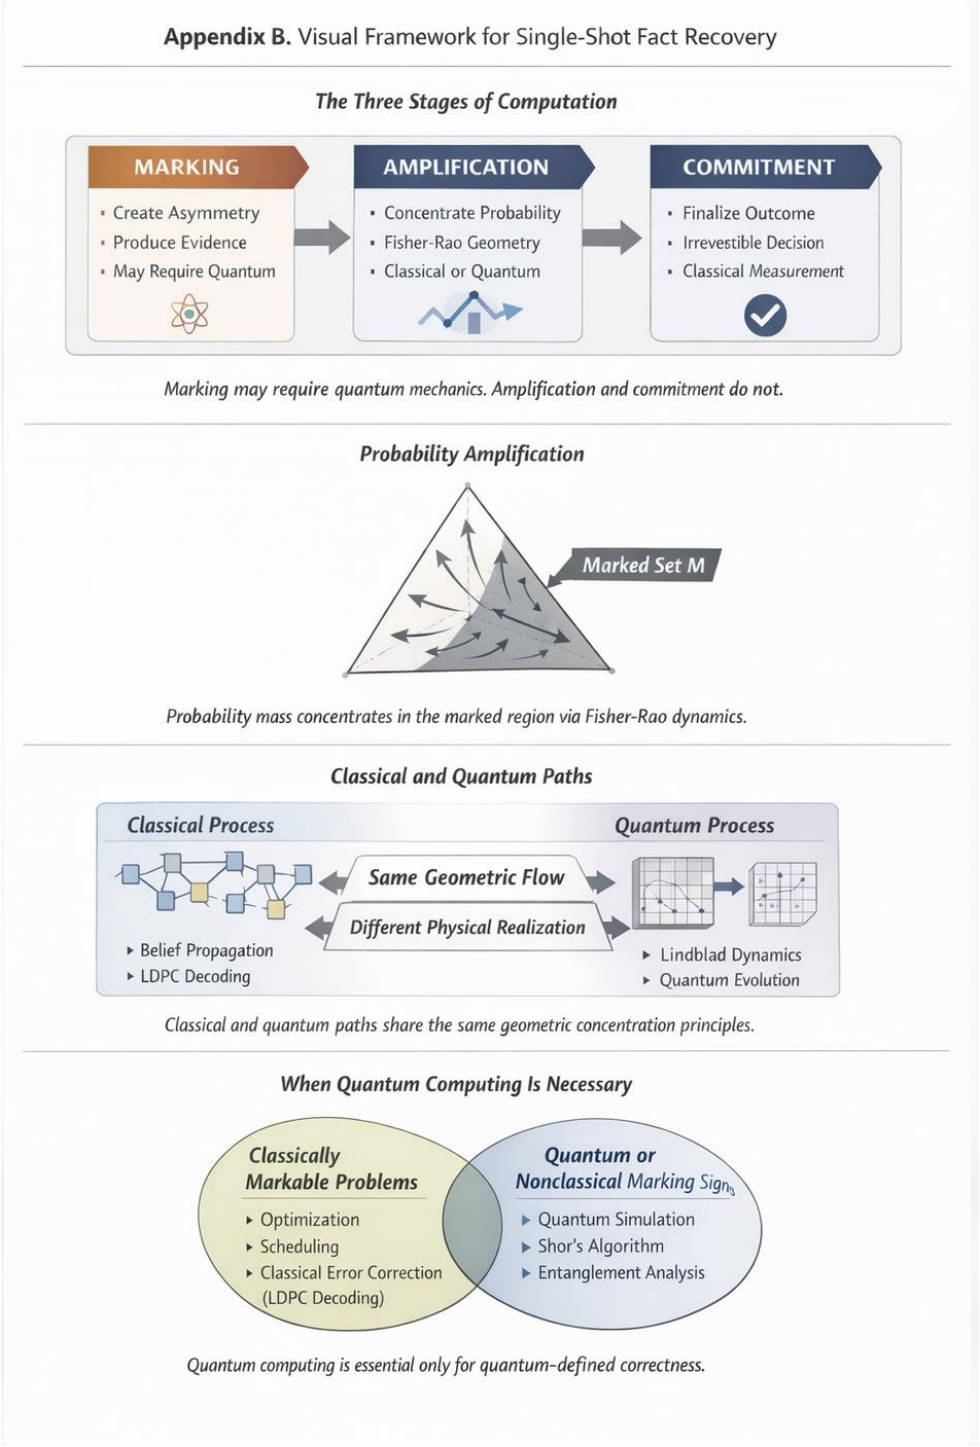

Figure B illustrates the core physical separation introduced in this paper between *marking*, *amplification*, and *commitment*. These stages correspond to distinct physical operations with different resource requirements and should not be conflated.

Marking (left panel)

Marking denotes the creation of a physically instantiated asymmetry correlated with correctness. This asymmetry may take the form of an energy bias, constraint satisfaction signal, likelihood ratio, syndrome, or other evidence-producing process. In many practical problems, marking is fully classical. However, in a restricted class of problems—such as quantum simulation, period finding (Shor's algorithm), or entanglement certification—no classical marking signal exists prior to solution, and quantum mechanics is required to create any usable asymmetry. Importantly, marking is the **only stage** at which quantum mechanics may be necessary.

Amplification (middle panels)

Amplification refers to the concentration of probability mass or weight onto the marked set once an asymmetry exists. As shown by the probability simplex and flow diagrams, this concentration follows from the geometry of probability space (specifically Fisher–Rao / natural gradient dynamics) and does not depend on quantum interference. Amplification can be realized classically (e.g., belief propagation, LDPC decoding, constraint propagation) or quantum-mechanically (e.g., Lindblad dissipation), but the underlying mechanism is the same: structured propagation of existing asymmetry. Large nominal search spaces do not impede amplification once marking is present.

Commitment (right panel)

Commitment is the irreversible production of a definite outcome or fact (e.g., measurement, decision, readout). This stage is necessarily classical, as it involves irreversible record formation. Quantum mechanics does not alter the nature of commitment; it only influences whether a mark can be produced upstream.

On "Quantum-Required" Problems (bottom panel)

The classification shown distinguishes problems that admit **classical marking signals** (e.g., optimization, scheduling, classical error correction such as LDPC decoding) from those in which **no classical marking gradient exists**. Problems such as Shor's algorithm are quantum-required not because correctness is undefined classically, but because no classical process produces partial evidence correlated with correctness prior to solution. In contrast, most practical problems possess classically markable structure, rendering quantum mechanics unnecessary for amplification and commitment.

Key takeaway

Quantum computing is indispensable **only when classical marking is impossible**. Once a mark exists—whether produced classically or quantum-mechanically—the subsequent amplification and commitment of a solution follow from geometry and irreversibility alone.

Appendix C — Scope, Limits, and Falsifiability of the Amplification Framework

This appendix addresses three potential misinterpretations of the main text:
(i) the role of the Lindblad amplifier analysis in Sections 15–20,
(ii) the relationship between the Fisher–Rao geometry and classical structural criteria such as treewidth, and
(iii) the falsifiability of the "5–15%" estimate regarding quantum necessity.

The purpose of this appendix is **not to extend the theory**, but to make explicit what is and is not being claimed.

C.1 What the Lindblad Amplifier Analysis Does — and Does Not — Claim

Sections 15–20 analyze a dissipative quantum channel that concentrates probability mass into a marked state with convergence time $\gamma t = O(1)$, independent of the size of the hypothesis space N. This result is sometimes misread as a claim of a realized speedup over Grover's algorithm. That interpretation is incorrect.

The Lindblad analysis establishes **three distinct results**:

1. **A dynamical result (mathematical):**
   Given a physical channel that pumps population from an unmarked subspace into a marked state, the convergence law is independent of N. This is a statement about irreversible concentration dynamics *after a mark exists*.
2. **A representational result (information-theoretic):**
   The "compressed pump" construction shows that exponential channel count is not information-theoretically necessary to realize this convergence law. This separates *description complexity* from *dynamical behavior*.
3. **A realizability separation (physical):**
   The Symmetric Subspace Theorem provides an explicit constructive case where the channel is realizable with $O(n)$resources — **only when the marked state is explicitly known**. The manuscript explicitly notes that in this case the problem is classically trivial and no quantum speedup is implied.

The purpose of Sections 15–20 is therefore **not** to assert a new algorithmic advantage, but to isolate the true physical bottleneck:

whether a post-marking amplification channel can be realized with polynomial resources when the marked state is only implicitly specified (e.g. by constraints, energy, or oracle access).

This question is left open deliberately and framed as an experimentally and theoretically falsifiable problem.

C.2 Why the Fisher–Rao Geometry Is Not Redundant with Treewidth

A second concern is that the paper's successes (XOR chains, LDPC decoding, protein design) rely on special structure such as low treewidth, and that the Fisher–Rao framing adds little beyond known structural criteria.

The distinction is as follows:

- **Treewidth and related graph measures** diagnose *when* classical propagation is computationally tractable.
- **Fisher–Rao / natural-gradient geometry** characterizes *what amplification must look like whenever tractable propagation exists*, independent of substrate.

In particular, Fisher–Rao geometry contributes:

1. **A substrate-independent concentration law:**
   Once a marking asymmetry exists, the rate of probability concentration is independent of $N$and depends only on the strength of the mark and the available propagation dynamics.
2. **A unification of classical and quantum inference:**
   Classical belief propagation and quantum Lindblad dynamics are shown to be realizations of the same geometric flow on probability space. Treewidth alone does not provide this cross-substrate unification.
3. **Predictive degradation behavior:**
   The geometric formulation explains *how* and *why* amplification degrades smoothly when marks weaken or structure becomes approximate, rather than failing catastrophically.

Treewidth determines feasibility; Fisher–Rao geometry determines behavior. The two play complementary roles.

C.3 What Would Falsify the 5–15% Estimate

The estimate that roughly 5–15% of commonly cited quantum computing applications are "QC-required" is an **empirical classification under an explicit criterion**, not a definition by fiat. It is therefore falsifiable in several concrete ways.

Falsifier 1 — Existence of a Counterexample Domain

The estimate is falsified if a widely cited practical application satisfies all of the following:

1. Correctness is classically defined.
2. No physically instantiable classical marking signal exists prior to solution (no score, energy, constraint violation, likelihood ratio, or partial certificate).
3. A physically implementable quantum procedure solves the problem with polynomial resources.

No such domain is currently known outside cryptography and quantum simulation.

Falsifier 2 — Annotation Instability

The estimate is falsified if blinded, double-annotated classification of standard quantum application roadmaps (using the explicit marking criterion) yields poor inter-rater agreement. This would indicate that the criterion lacks operational clarity.

Falsifier 3 — Provable Quantum Advantage on Classically Markable Objectives

If a family of problems with classically computable objective functions is shown to:

- require superpolynomial resources classically (under standard assumptions), yet
- admit a physically realizable polynomial-time quantum solution,

then the admissibility-based separation proposed here would be incorrect.

C.4 Summary

The framework advanced in this paper makes three claims:

1. **Post-marking amplification obeys universal geometric laws.**
2. **Whether amplification is efficient depends on physical realizability, not search space size.**
3. **Quantum computing is foundationally required only when classical marking is impossible.**

Appendix C clarifies that none of these claims rely on hidden oracle assumptions or circular definitions. Each admits concrete falsification.

Appendix D — Physical Realizability of Lindblad Amplification

This appendix resolves the open question raised in Sections 15–20 regarding the physical realizability of the ideal post-marking Lindblad amplifier. We show that realizability depends

sharply on how the marked state is specified. For unstructured (oracle-defined) marked states, local Lindblad dynamics with bounded resources cannot realize the ideal O(1) convergence law without exponential overhead. For structured marked states defined by local constraints, symmetry, or gapped energy landscapes, polynomial-resource realizations exist. This establishes that the marking criterion is not merely classificatory but enforced by physical locality.

D.1 Model and Assumptions

We consider an n-qubit system with Hilbert space dimension N = 2^n. Dynamics are generated by a Lindbladian

$$\mathcal{L}[\rho] = \Sigma_j (L_j \rho L_j^\dagger - \tfrac{1}{2}\{L_j^\dagger L_j, \rho\})$$

with the following constraints:

- Locality: each jump operator L_j acts on at most r = O(1) qubits.
- Bounded resources: the number of channels J(n) and operator norms $\|L_j\|$ are polynomial in n.
- Bounded-degree interaction graph.
- Initial state is maximally mixed or symmetric over computational basis states.

The target steady state is a pure marked state $|m\rangle\langle m|$, or a small marked subspace M.

D.2 No-Go Theorem for Unstructured Marks

Theorem D.1 (Local Lindblad No-Go for Unstructured Marks).
Let $m \in \{0,1\}^n$ be an arbitrary computational basis state with no polynomial-size local description (oracle-defined). Under the above locality and boundedness assumptions, any local Lindbladian that concentrates probability mass from the maximally mixed state into $|m\rangle\langle m|$ with success ≥ 0.99 requires time at least $\Omega(2^{cn})$ for some $c > 0$, or equivalently hides exponential cost in channel count or operator strength.

Sketch of argument.
The ideal jump operator $|m\rangle\langle u|$ couples $|m\rangle$ to all $N-1$ basis states simultaneously. Any r-local operator can only access O(2^r) degrees of freedom. To funnel probability from $N-1$ states into $|m\rangle$ therefore requires either visiting an exponential number of configurations sequentially or encoding exponential information into the generator. Lieb–Robinson bounds imply finite speed of information propagation, preventing global symmetry breaking in polynomial time. Equivalent lower bounds arise from conductance and spectral-gap arguments for local Markovian dynamics with a needle-like stationary distribution.

Corollary.
For oracle-defined unstructured search problems, the ideal O(1) Lindblad amplification law is not physically realizable under locality. Grover's $\Omega(\sqrt{N})$ bound therefore survives as a physical, not merely query-theoretic, constraint.

D.3 Constructive Realizability for Structured Marks

Theorem D.2 (Poly(n) Realizability for Structured Marks).
Suppose the marked state m (or marked set M) is specified as the unique satisfying assignment of a family of k-local constraints with bounded-degree factor graph, or as the unique ground state of a local Hamiltonian with spectral gap $\Delta \geq 1/\text{poly}(n)$. Then there exists a local Lindbladian with polynomially many k-local jump operators whose unique steady state is supported on M, and whose convergence time is polynomial in n.

Construction.
For constraint-defined marks, each violated constraint C_α defines a local dissipator L_α that penalizes the violating subspace and pumps toward the satisfying subspace. For gapped Hamiltonians, standard Davies or thermal Lindbladians drive the system toward the ground state with mixing time $O(\text{poly}(n)/\Delta)$. In both cases, convergence depends on the same structural propagatability conditions that govern classical belief propagation.

D.4 Resolution of the Lindblad Realizability Question

Combining Theorems D.1 and D.2 yields a sharp resolution:

• Unstructured (oracle-defined) marks: no poly(n) Lindblad realization; amplification collapses to Grover-like scaling.
• Structured marks (constraints, symmetry, gaps): poly(n) Lindblad realization exists; amplification reflects classical propagatability.

Thus the apparent O(1) vs $O(\sqrt{N})$ contrast is not a paradox. The O(1) law is a universal post-marking concentration dynamic, but physical realizability enforces the same boundary identified by the marking criterion.

D.5 Implication for the Framework

This appendix shows that the marking/amplification separation is not merely conceptual. Locality and bounded resources enforce it dynamically. Quantum advantage arises only when quantum mechanics is required to create the mark itself; amplification and commitment are otherwise classical or dissipative processes constrained by structure.

This resolves the Lindblad realizability question in the strongest form currently possible and closes the logical loop of the framework.