Perspective

Vast World of Quantum Advantage

Hsin-Yuan Huang,[1,2] Soonwon Choi,[3] Jarrod R. McClean,[2] and John Preskill[1,4]
[1]California Institute of Technology, Pasadena, California 91125, USA
[2]Google Quantum AI, Venice, California 90291, USA
[3]Massachusetts Institute of Technology, Cambridge, Massachusetts 02139, USA
[4]AWS Center for Quantum Computing, Pasadena, California 91125, USA

(Received 7 August 2025; revised 19 March 2026; published 9 July 2026)

The quest to identify quantum advantages lies at the heart of quantum technology. While quantum devices promise extraordinary capabilities, from exponential computational speedups to unprecedented measurement precision, distinguishing genuine advantages from mere illusions remains a formidable challenge. In this endeavor, quantum theorists are like prophets attempting to foretell the future, yet the boundary between visionary insight and unfounded fantasy is perilously thin. In this Perspective, we examine our mathematical tools for navigating the vast world of quantum advantages across computation, learning, sensing, and communication. We explore five keystone properties, predictability, typicality, robustness, verifiability, and usefulness that define an ideal quantum advantage, and envision what new quantum advantages could arise in a future with ubiquitous quantum technology. We prove that some quantum advantages are inherently unpredictable using classical resources alone, suggesting a landscape far richer than what we can currently foresee. While mathematical rigor remains our indispensable guide, the ultimate power of quantum technologies may emerge from advantages we cannot yet conceive.

DOI: 10.1103/tn89-g1xz Subject Areas: Quantum Information

I. QUEST FOR QUANTUM ADVANTAGE

Quantum entanglement is often portrayed as the epitome of quantum strangeness. Even when two entangled particles are separated by light years, measuring one particle causes the other to instantly "know" and collapse to a correlated state. This phenomenon appears to enable faster-than-light communication, violating the fundamental speed limit that governs our classical Universe. Could this mysterious quantum effect enable instant galactic communication? The answer, as we now understand, is no. Yet this question reveals a deeper challenge that continues to perplex physicists: When does quantum physics truly outperform classical physics?

This quest to identify genuine quantum advantages, tasks where quantum systems fundamentally surpass what is achievable by their classical counterparts, lies at the heart of quantum technology. While quantum devices are heralded for their revolutionary potential, from quantum computers promising exponential speedups to quantum sensors achieving unprecedented precision, distinguishing genuine quantum advantages from mere illusions has proven remarkably subtle. Entangled particles illustrate this perfectly: Despite exhibiting correlations that seem to defy classical physics, quantum entanglement does not necessarily enable capabilities beyond what classical systems can achieve.

A. Puzzle 1: Is this spooky?

Consider two observers, Alice and Bob, sharing an entangled quantum state:

$$\frac{1}{\sqrt{2}}(|\uparrow\rangle \otimes |\uparrow\rangle + |\downarrow\rangle \otimes |\downarrow\rangle). \tag{1.1}$$

Alice and Bob are separated to opposite ends of the galaxy.

(i) When Alice measures her particle's spin in the $\{|\uparrow\rangle, |\downarrow\rangle\}$ basis, Bob's particle instantly collapses to $|\uparrow\rangle$ or $|\downarrow\rangle$ depending on Alice's outcome.
(ii) When Bob measures his spin in the $\{|\uparrow\rangle, |\downarrow\rangle\}$ basis, he immediately knows Alice's result.
(iii) This happens instantly and is faster than light could travel between them.

Can we simulate this behavior using only classical objects, like a pair of socks? Does the answer change if Alice and Bob randomly choose to measure in either the $\{|\uparrow\rangle, |\downarrow\rangle\}$ basis or the rotated basis $\{|\nearrow\rangle, |\swarrow\rangle\}$? The solution can be found in Appendix A 1.

This puzzle illustrates a crucial principle about quantum advantages: They can be remarkably subtle. When Alice

and Bob measure only in the $\{|\uparrow\rangle, |\downarrow\rangle\}$ basis, their quantum correlations can be perfectly mimicked by classical systems, such as pairs of socks with predetermined properties. Yet a small change, allowing measurements in multiple bases, reveals a genuine quantum advantage that cannot be replicated by any classical system with local interactions. This fundamental impossibility was established in Bell's elegant theorem [1], which demonstrated the nonlocal nature of quantum correlations. Such delicate distinctions between quantum and classical capabilities pervade the study of quantum advantages.

As quantum technology advances, distinguishing genuine advantages from *pseudoadvantages* has become increasingly crucial. A pseudoadvantage arises when a quantum protocol appears to outperform all classical approaches, only to be matched by a clever classical strategy that was previously overlooked. Puzzle 1 demonstrates that while quantum entanglement seems to enable nonlocal instantaneous correlations, the same statistical outcomes can be achieved classically with pairs of socks, at least when measurements are restricted to a single basis.

The subtlety of quantum advantages extends far beyond this example. Even when quantum algorithms appear to achieve extraordinary feats impossible for classical computers, careful analysis sometimes reveals surprising classical alternatives. The following puzzle illustrates how quantum advantages can be particularly deceptive in the realm of data analysis and machine learning.

B. Puzzle 2: Exponential quantum advantage in analyzing big data?

Consider a massive dataset represented by M exponentially long vectors $\vec{x}_1, \ldots, \vec{x}_M \in \mathbb{R}^{2^n}$, each having unit norm. Alice and Bob want to analyze this massive dataset, but they have very different tools at their disposal. Alice can harness the power of quantum technology. Suppose she can create quantum states representing her data in just $\mathrm{poly}(n)$ time that eventually stores the state in $\mathcal{O}(n)$ qubits of memory:

$$|x_k\rangle = \sum_{i=1}^{2^n} x_{k,i}|i\rangle. \tag{1.2}$$

Using this ability and a basic quantum algorithm, Alice can efficiently compute inner products between any two vectors using the stored states. Even though the vectors have 2^n components, she can estimate

$$\langle x_k|x_\ell\rangle = \sum_{i=1}^{2^n} x_{k,i} \cdot x_{\ell,i} = \vec{x_k} \cdot \vec{x_\ell} \tag{1.3}$$

to precision $1/\mathrm{poly}(n)$ in $\mathrm{poly}(n)$ time. This seems almost magical as Alice can analyze relationships between vectors of exponential size in time that grows logarithmically in their dimensions. Meanwhile, Bob only has access to classical technology: a standard random-access memory (RAM) with size $\mathrm{poly}(n)$ able to retrieve chunks of the same data at a time. At first glance, his task seems hopeless, as reading all entries of a single vector would take exponential time in n.

Can Bob's classical computation somehow match Alice's quantum prowess? Or has Alice discovered a genuine quantum advantage? The solution can be found in Appendix A 2.

At first glance, Alice's quantum approach seems unbeatable. However, Bob has a clever solution using only classical importance sampling to exactly match Alice's quantum performance. The inner product can be estimated in classical polynomial time without any quantum technology. We also note that Alice's ability to store the information into the amplitudes of a state in $\mathrm{poly}(n)$ is sometimes referred to as quantum random access memory (QRAM). The best known scheme for creating fault-tolerant QRAM generally requires a quantum spacetime footprint scaling as $\mathcal{O}(2^n)$, so this advantage can disappear due to the existing challenge in physical implementations [2,3].

Such pseudoadvantages have appeared repeatedly throughout the development of quantum technology. Perhaps the most dramatic story began with the HHL algorithm, a quantum algorithm for solving linear systems of equations that exemplifies genuine quantum advantage. In 2009, Harrow, Hassidim, and Lloyd proved that quantum computers can solve linear systems superpolynomially faster than classical computers [4]. More precisely, they proved that this superpolynomial quantum advantage follows from the widely believed conjecture that quantum computer is strictly more powerful than classical computers in at least one computational tasks, or $\mathsf{BPP} \neq \mathsf{BQP}$. This breakthrough inspired hundreds of other *derived* quantum algorithms for machine learning, optimization, and numerical linear algebra, many claiming exponential speedups over classical approaches.

However, starting in 2018, this excitement collided with reality. Researchers discovered that some of these derived quantum algorithms could be matched by clever classical approaches [5,6], similar to Bob's solution in Puzzle 2. While the HHL algorithm still has genuine quantum advantage and may provide practical value in certain applications, many previously proposed applications of the HHL algorithm for practically useful tasks turn out to offer only modest computational speedups. This pattern, apparent quantum advantages dissolving under careful analysis, teaches us a profound lesson: Even when building upon genuine quantum advantages, we must rigorously verify that the advantage persists in our specific application of interest.

The story of quantum advantages often takes multiple unexpected turns. Consider one of the first enticing facts one learns about in the first chapter of a quantum

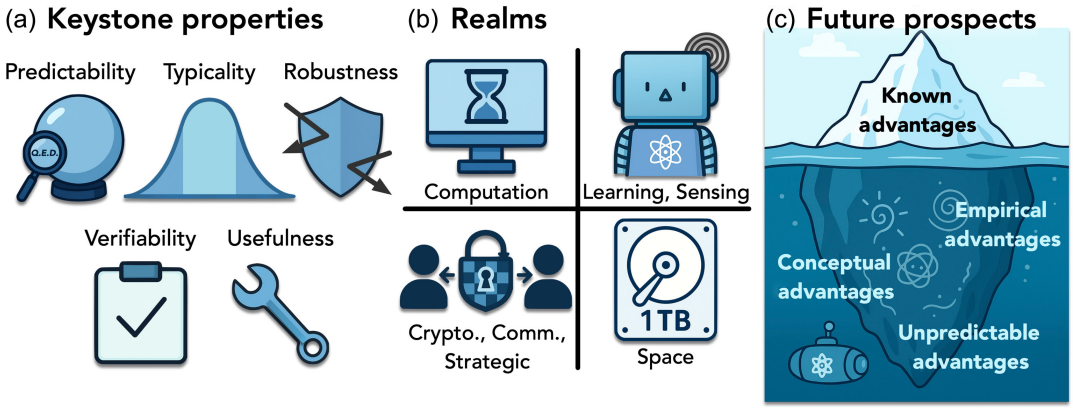

FIG. 1. Conceptual overview of the vast world of quantum advantage. (a) Keystones of quantum advantage. We propose five essential properties that collectively define an ideal quantum advantage: predictability (supported by rigorous evidence), typicality (applying to a substantial fraction of naturally occurring problem instances), robustness (maintaining advantage despite noise and imperfections), verifiability (enabling efficient validation of correctness), and usefulness (providing genuine practical value). (b) Realms of quantum advantage. Quantum advantages can be found across four primary domains: computation (algorithmic speedups), learning and sensing (enhanced methods for probing the world), cryptographic, communication, and strategic games (advantages in security and coordination), and space (improved memory efficiency). (c) Future of quantum advantage. The visible tip of the iceberg represents the currently known and predicted advantages, while the vast, submerged part represents the largely unexplored frontier of empirical, conceptual, and fundamentally unpredictable advantages that may only be discovered through future quantum technologies.

computing textbook: A quantum state of n qubits contains 2^n complex amplitudes, suggesting an exponential advantage in information storage compared to classical bits. However, the next chapter usually dashes these hopes as Holevo's bound [7] proves that only $\mathcal{O}(n)$ classical bits of information can be reliably retrieved from n qubits through measurement. Yet, remarkably, this is not the end of the story. The following puzzle illustrates how quantum advantages can reemerge in subtle ways.

C. Puzzle 3: Can quantum systems encode exponentially many bits?

Consider a large, normalized vector $\vec{x}$ of dimension $N = 2^n$ stored in Alice's classical memory. Alice wishes to compress this data into $\log N$ qubits to save space and bandwidth when sending it to Bob. She encodes her data into quantum amplitudes using $n = \log N$ qubits:

$$|x\rangle = \sum_{i=1}^{2^n} x_i |i\rangle. \tag{1.4}$$

Satisfied with her extremely compact encoding, Alice realizes she can now transmit or teleport just n qubits instead of the exponentially larger original data. She sends several copies to Bob and keeps a few herself before deleting the original data. Bob stores the quantum states for later processing but wonders, given Holevo's bound, could Alice have just sent $\mathcal{O}(n)$ classical bits and saved all the trouble of using quantum technology?

Did Alice and Bob gain anything by compressing the data into quantum states? Does compressing data into quantum states provide an advantage? Is this compression a genuine quantum advantage, or yet another illusion? The solution can be found in Appendix A 3.

At first glance one may conclude that quantum states offer a simple way to exponentially compress data for communication. However, Holevo's bound appears to quickly and completely shatter this hope as Bob can extract only $\mathcal{O}(n)$ classical bits from his quantum states, suggesting Alice could have achieved the same result by sending classical bits. Yet this is not the complete story. The answer crucially depends on what Bob wants to do with the data. For some tasks, like computing the overlap with another vector, the compression into a quantum state offers no advantage [5]. However, for other tasks, such as determining whether the vector belongs to a particular subspace or its orthogonal complement, compressing data into a quantum state provides a genuine exponential quantum advantage. Classical communication for solving this problem requires $\Theta(2^{n/2})$ bits, while quantum communication requires only $\mathcal{O}(n)$ qubits [8]. The maximum $\mathcal{O}(n)$ vs $\Omega(2^n)$ quantum-classical separations have since been established for related tasks [9,10], highlighting that quantum states can, in appropriate senses, encompass the power of 2^n classical bits. Moreover, these quantum advantages do not rely on any unproven computational conjecture. This type of advantage has found recent applications in quantum machine learning, such as for computing quadratic functions of data [11].

These examples underscore the importance of rigorous analysis of quantum advantages. As quantum technologies rapidly advance and require substantial investments of human effort and financial resources, establishing rigorous

theoretical foundations becomes essential to ensure these efforts are well directed. Quantum computing ventures now attract billions of dollars in investment, and governments worldwide launch major quantum initiatives, making the stakes for establishing genuine quantum advantages increasingly important. Without solid theoretical foundations, we risk pursuing technological directions that may ultimately prove less effective than anticipated, potentially misallocating significant resources and undermining confidence in quantum technologies. The quantum computing community continues to grapple with these challenges, debating evaluation criteria for quantum algorithms and addressing misconceptions about near-term capabilities, while recognizing the critical need for algorithmic breakthroughs to justify ongoing hardware investments [12–15].

At the same time, the quest for genuine quantum advantages offers rewards that extend far beyond responsible stewardship of resources. Seeking quantum advantages requires uncovering new insights into the fundamental distinctions between quantum and classical physics. These insights often transcend their original context, offering new perspectives in fields ranging from condensed matter physics to quantum gravity, from computational complexity to information theory. These profound intellectual rewards provide yet another compelling reason to pursue rigorous quantum advantages. The careful analysis of quantum advantages thus serves both the responsible stewardship of society's investment and our fundamental understanding of the quantum universe.

II. KEYSTONES OF QUANTUM ADVANTAGE

Having explored why identifying genuine quantum advantages is crucial for the future of quantum technology, we now turn to a systematic analysis of what constitutes compelling examples of quantum advantages. What properties must they possess to be both theoretically convincing and practically valuable? This question is particularly pressing in the current era, where quantum technologies are rapidly advancing yet remain imperfect, and where substantial resources are being allocated to quantum research and development. We propose five keystone properties that define an ideal quantum advantage. See Fig. 1(a) for an illustration of the five keystone properties.

A. Predictability

All theorists studying quantum technologies are, in essence, prophets trying to predict the future. In the current era, where large-scale quantum technology remains elusive, the exploration of quantum advantages requires careful theoretical guidance to distinguish genuine advantages from pseudoadvantages that may vanish upon closer examination. This motivates our first keystone property: Quantum advantages should be predictable, ideally through rigorous, substantive, and quantifiable evidence.

Definition 1 (Predictability). A quantum advantage is predictable if we have sufficient evidence supporting that, given the necessary quantum technology, we will achieve capabilities fundamentally beyond what is possible with classical technology.

The importance of predictability stems from the inherent difficulty of reasoning about quantum advantages. Establishing genuine quantum advantages is a subtle endeavor fraught with hidden pitfalls. Our classical intuitions often fail us when reasoning about quantum systems, and seemingly obvious quantum advantages can dissolve under careful analysis. Thus, we need robust theoretical frameworks not only to distinguish genuine quantum advantages from pseudoadvantages that merely reflect our incomplete understanding, but also to quantify the magnitude of these advantages with mathematical rigor.

Example 1 (Quantum recommendation systems). The evolution of quantum recommendation systems illustrates the subtle challenges in predicting quantum advantages. Early work [16] proposed quantum algorithms that appeared to offer exponential speedups for recommendation problems, representing a natural application area where quantum advantages seemed plausible. However, the wide array of classical algorithmic techniques made it difficult to definitively establish the advantage. Building on techniques similar to those in Puzzle 2, researchers later found that classical algorithms could achieve comparable performance [5], demonstrating how some apparent quantum advantages can dissolve when the full classical algorithmic landscape is thoroughly explored.

This and similar experiences have driven the quantum computing community to develop more rigorous methodologies for establishing predictable and quantifiable advantages. Several powerful paradigms have emerged, each with different strengths for providing reliable predictions about quantum capabilities. The gold standard is formal mathematical proof for separating the optimal performance achievable by classical technology from the performance using quantum technology. Such formal proofs allow us to predict quantum advantages with absolute certainty and precisely quantify their magnitude. Nonlocal games built upon Bell's inequality (explored in Puzzle 1) exemplify this approach.

A second powerful approach proceeds via reductions to widely believed conjectures. The key is identifying conjectures supported by diverse evidence across different fields. We then prove that the absence of quantum advantage would violate these conjectures. Common examples include $\mathsf{BPP} \neq \mathsf{BQP}$ (not all quantum polynomial-time problems are classically solvable in polynomial time), the noncollapse of the polynomial hierarchy [17], and the classical hardness of factoring [18]. Even quantum computational hardness assumptions, such as the hardness for quantum computers to solve the learning with errors problem [19], can sometimes be transformed into proofs of quantum advantage [20,21].

Example 2 (Shor's algorithm for factoring [18]). Shor's algorithm provides compelling evidence for quantum advantage through reduction to the hardness of factoring, a problem underpinning trillions of dollars of global digital commerce [22,23]. Rivest-Shamir-Adleman (RSA) encryption, standardized by NIST as a federal cryptographic standard [24], forms a critical component of this digital infrastructure. Despite enormous financial incentives and decades of intensive research by cryptographers and computer scientists, no efficient classical algorithm has emerged. The proven ability of Shor's algorithm to factor efficiently, combined with the sustained difficulty of classical approaches, provides strong evidence for quantum advantage.

Example 3 (Cooling physical systems [25]). A recent result shows how quantum advantage emerges in the seemingly simple task of cooling quantum systems. When a quantum system is placed in a cold environment, it naturally finds a local minimum of its energy. Local minima are not necessarily the absolute lowest (ground) state, but a state that cannot be improved by small changes. While finding ground states is believed hard even for quantum computers, cooling to local minima exhibits a clear quantum advantage: It is efficiently achievable by quantum computers but provably hard for classical computers under the widely believed conjecture that $\mathsf{BPP} \neq \mathsf{BQP}$. This result is particularly compelling because it connects to a ubiquitous physical process, the natural cooling of quantum systems, while providing rigorous evidence for quantum advantage through reduction to complexity assumptions.

Open question 1: New conjectures for building quantum advantages. Currently, provable quantum advantages rely on relatively few complexity-theoretic conjectures. Can we identify new mathematical conjectures, perhaps inspired by physical principles, that enable us to predict and develop new types of quantum advantages?

An emerging third paradigm elegantly combines empirical numerical calculations with mathematical proofs to identify *instance-specific quantum advantages*. For certain quantum algorithms, we can predict their performance on specific problem instances by computing particular quantities that are feasible using classical computers. While quantum technology remains essential for actually solving the original problem, this approach enables rigorous comparison between predicted quantum performance and the capabilities of the best known classical algorithms, allowing us to forecast and quantify quantum advantages before constructing the necessary quantum hardware. Unlike theoretical worst-case analysis, this paradigm enables evaluation on specific problem instances that arise in real applications. While this approach cannot rule out future classical algorithmic breakthroughs, it substantially expands our ability to analyze potential quantum advantages in practically relevant problem instances.

Example 4 (Quantum approximate optimization algorithm [26]). Recent advances in understanding the quantum approximate optimization algorithm (QAOA) [26] illustrate this paradigm. For specific problem classes, QAOA's performance can be predicted through classical numerical computation [27–29]. Remarkably, rigorous lower bounds can be established even when running the quantum algorithm is infeasible with current technology. For MAXCUT on high-girth three-regular graphs, this analysis has provided evidence for potential quantum advantages over known classical algorithms [29].

Example 5 (Decoded quantum interferometry [30]). Decoded quantum interferometry (DQI) [30] provides another compelling illustration. DQI reduces classical optimization problems to classical decoding problems through quantum interference. DQI's performance can be predicted by evaluating classical decoders, even for large-scale instances. This enables comparison with classical algorithms before building quantum hardware. For certain problems, such as optimal polynomial intersection, this approach has revealed that DQI achieves better approximation ratios than any known classical algorithm [30].

Example 6 (Quantum phase estimation [31–33]). This paradigm also manifests in quantum phase estimation (QPE) for finding low-energy states of a Hamiltonian H. After classical algorithms like DMRG [34] find their best approximation to the ground state, we can predict how quantum phase estimation would improve upon this result. If the energy variance $\mathrm{tr}(H^2\rho) - \mathrm{tr}(H\rho)^2$ of the classical state ρ is substantial, QPE can efficiently find lower-energy states [33]. This variance, often efficiently computable classically, provides a rigorous lower bound on the quantum advantage without requiring quantum hardware.

Open question 2: New theory for predicting quantum advantages. How can we expand our mathematical arsenal for predicting and quantifying quantum advantages before having full-fledged quantum technology? Can we develop a systematic theory for forecasting when and how quantum systems will surpass classical limitations, along with precise bounds on the magnitude of these advantages? The ability to rigorously predict and quantify quantum advantages may be as transformative as the advantages themselves.

B. Typicality

When researchers establish quantum advantages for computational, communication, learning, or sensing problems, they often prove that quantum technology can efficiently solve all instances within a problem class, while classical technology provably cannot. Though such results seem ideal, the reality is more nuanced. The distinction between worst-case and average-case complexity plays a crucial role in determining the practical relevance of quantum advantages.

Example 7 (Simulating quantum circuits). Consider the task of simulating a closed quantum system evolving under a time-dependent Hamiltonian described by a

polynomial-size quantum circuit C. The goal is to estimate local observables on the output state after evolving a simple product input state under circuit C. Under the widely believed conjecture that $\mathsf{BPP} \neq \mathsf{BQP}$, this task exhibits a superpolynomial quantum advantage in the worst case, that is, for the quantum circuits that are most difficult to simulate classically. However, this quantum advantage applies only for some quantum circuits. For most quantum circuits consisting of independent and generic circuit layers, the quantum chaos in the dynamics enables efficient classical simulation through low-weight Pauli propagation [35–43]. Moreover, it is widely believed that under natural Hamiltonian dynamics, local observables approach the predictions from generalized thermal equilibrium states characterized by a few conserved quantities such as particle numbers. This stark contrast between worst-case and average-case complexity exemplifies why worst-case advantages need not imply quantum advantages for typical instances.

The relationship between worst-case and average-case quantum advantages can be quite subtle. A problem may show no quantum advantage in the worst case, yet possess significant advantages for typical instances. Conversely, as illustrated in Example 7, a problem may demonstrate dramatic worst-case quantum advantage while admitting efficient classical solutions for most instances.

From a practical perspective, the ability to solve contrived worst-case instances is far less valuable than solving the typical instances encountered in real-world applications. This observation motivates our second criterion for meaningful quantum advantages:

Definition 2 (Typicality). A quantum advantage exhibits typicality if it applies to a substantial fraction of practically relevant problem instances, rather than being confined to particular instances that are especially hard classically.

This definition emphasizes that quantum advantages should apply broadly to the kinds of problems we actually want to solve, not just carefully crafted instances to prove quantum-classical separations. A typical quantum advantage promises practical impact across a broad spectrum of applications, making it far more valuable than a narrow advantage limited to pathological cases that rarely occur in practice.

To establish typical quantum advantages, we can build on average-case complexity theory, a subfield of theoretical computer science that provides tools for analyzing algorithmic performance on typical problem instances. Unlike worst-case complexity, which considers the maximum computational resources required across all possible inputs, average-case complexity examines the expected resource consumption under natural probability distributions of problem instances. While this theoretical framework is still developing, it has already yielded several compelling examples of typical quantum advantages.

Example 8 (Random circuit sampling [44–48]). Random circuit sampling (RCS) exemplifies how average-case hardness can establish typical quantum advantages. The task involves sampling from the output distribution of a random quantum circuit to within a specified total variation distance. For example, Google's Willow processor demonstrated quantum supremacy by sampling from a 103-qubit random circuit in just a few minutes, while a classical computer would require much longer than the age of the Universe to complete the same task [47]. The classical hardness of RCS stems from a reduction. For most circuit instances, approximate sampling to small total variation distance requires superpolynomial classical computation, under the widely believed assumption that the polynomial hierarchy does not collapse. These results establish that quantum advantages for RCS are typical. They hold for most random quantum circuits rather than carefully constructed worst-case instances.

Example 9 (Discrete logarithm problem). The discrete logarithm problem (DLP) provides another prototypical example of typical quantum advantage. Given a cyclic group G of prime order p with generator g, and an element $h \in G$, the task is to find x such that $g^x = h$. While Shor's algorithm solves the DLP in quantum polynomial time, the best-known classical algorithms require exponential time for most group elements when G is carefully chosen [49]. The generic nature of this quantum advantage stems from a powerful reduction: Finding discrete logarithms for a random element $h \in G$ is provably as hard as solving the problem in the worst case [50]. This reduction is known as random self-reducibility and holds for broad classes of groups, including those used in practical cryptographic systems, making the average-case hardness a foundation for widely deployed cryptographic protocols.

These examples highlight how average-case complexity can establish quantum advantages that are both theoretically rigorous and practically relevant. However, developing the theory of average-case complexity remains an ongoing challenge. The field must address several fundamental questions: identifying natural probability distributions over problem instances that capture real-world applications, finding more problems with properties like random self-reducibility that connect average-case to worst-case hardness, and developing new proof techniques for average-case lower bounds on classical computation. As of now, most superpolynomial quantum advantages are not known to be typical. The most well known examples are random circuit sampling, where the randomly generated bit strings are not known to be of much use, and Shor's algorithm for breaking cryptographic protocols, where one can reinstate security by using postquantum cryptography [19]. More recently, the problem of optimal polynomial interpolation (OPI) has been conjectured to be another such instance. For OPI, there are parameter regimes that are efficiently solvable by DQI [30] but which are not solved in polynomial time by any known classical algorithms, even for average-case instances, though this problem is not

known to have the property of random self-reducibility like DLP. Aside from such superpolynomial speedups, there are also a variety of average-case-hard problems, for example, circuit-SAT, that admit Grover-like quadratic speedups for typical instances. More recently, quantum algorithms have been discovered giving a superquadratic polynomial speedup for the problem of sparse learning parity with noise (LPN) [51]. As LPN is a canonical average-case-hard problem where even a quantum computer is not expected to offer an efficient solution in general, better-than-quadratic polynomial speedups for typical instances of this problem are conceptually significant. These observations raise the following open question.

Open question 3: Typical quantum advantage. Are there other problems not based on random circuit sampling, Shor's algorithm, or OPI that exhibit typical, superpolynomial quantum advantage, where the quantum advantage applies to most instances of the problem rather than just worst-case instances? Can we establish new tools enabling us to show that broad families of problems have typical quantum advantage?

C. Robustness

The journey from theoretical to empirical quantum advantage often faces a formidable enemy: noise. Quantum systems are notoriously fragile as interactions with the environment easily destroy quantumness. Moreover, quantum advantages may erode under deviations from idealized theoretical assumptions. This raises a crucial question: Can quantum advantages survive under noise and imperfection?

Definition 3 (Robustness). A quantum advantage is robust if it persists despite deviations from idealized theoretical assumptions, including imperfections in hardware, available resources, input data, etc.

For quantum computing tasks, we possess the remarkable theoretical framework of quantum fault tolerance that provides a path to universal hardware robustness. The threshold theorem proves that if we can reduce errors in our quantum hardware below a certain threshold, we can achieve arbitrarily long quantum computations at reasonable overhead using carefully designed error correction protocols. This means quantum computational advantages, like Shor's factoring algorithm or quantum simulation of many-body systems, are theoretically robust: They maintain their advantage even with noisy quantum hardware, provided the noise is below threshold.

Example 10 (Fault-tolerant quantum computation). The threshold theorem ensures that a quantum circuit with C perfect gates can be simulated by a noisy quantum circuit with error at most ϵ, provided the noise per gate is below a threshold p_{th}. Traditional approaches require $\mathcal{O}[C\log^{c}(C/\epsilon)]$ noisy gates for this simulation, where c is a constant. Recent breakthroughs have dramatically improved this overhead: Using constant-rate quantum locally testable codes, a depth-D quantum circuit on n qubits with C ideal gates can be simulated by a noisy quantum circuit that maintains a constant space overhead [using only $\mathcal{O}(n)$ qubits] and near-logarithmic time overhead [using only $D\log^{1+o(1)}(C/\epsilon)$ depth] [52,53].

However, the story changes dramatically when we consider other kinds of quantum advantages. Quantum sensing offers a striking example.

Example 11 (Entanglement-enhanced sensitivity). By exploiting quantum entanglement, quantum sensors can achieve precision scaling as $1/(NT)$ using an N-particle probe with total sensing time T; this so-called Heisenberg-limited quantum enhanced scaling should be compared to the standard $1/(\sqrt{N}T)$ scaling achieved by an unentangled probe. However, this advantage dissolves in the presence of generic noise. Recent no-go theorems [54] prove that, for many sensing tasks, even small amounts of noise can completely destroy the asymptotic quantum advantage, forcing the precision to the scaling $1/\sqrt{NT}$ for both entangled and unentangled sensors. The $\sqrt{N}$ scaling advantage achieved using entangled probes in the ideal case is fundamentally unachievable when generic noise is present, even if quantum error correction is strategically deployed. A simplified, self-contained proof is given in Appendix B.

This example presents a stark contrast to the case of computation, where the exponential advantage of a quantum computer over a classical computer survives even when the quantum computer is subject to generic noise provided the noise is sufficiently weak.

Beyond hardware noise, quantum advantages must also withstand deviations from other theoretical assumptions. The problem instances occurring in the real world may not satisfy properties that guarantee quantum advantages. In physical simulation, the initial condition or the specification of the dynamics may be slightly different from those considered theoretically. And in machine learning, the classical datasets may be inherently noisy or incomplete [55–57].

In recent years, theorists have envisioned quantum artificial intelligence (AI) agents that can learn about our Universe by retrieving information using quantum sensors, storing the information in quantum memory, and processing the information using quantum computers. Ideal noiseless quantum AI agents have been proven to achieve some learning tasks with exponentially fewer experiments than would be required by classical AI agents that store data in classical memory and process them with classical computers [58–60]. Because quantum AI agents integrate multiple quantum technologies, they face robustness challenges across multiple fronts simultaneously. The vulnerability of quantum sensing to noise may hamper data collection, while memory decoherence can corrupt stored information, and computational errors can compromise data processing. While some quantum learning advantages have been proven to persist under realistic

noise conditions [60–62] and experimentally demonstrated in intermediate-scale superconducting qubit platforms [60] and photonic systems [63], the robustness of many other learning advantages remains uncharacterized. In particular, it is not known whether universal fault-tolerant schemes exist that can preserve all exponential advantages enjoyed by ideal quantum AI agents.

Open question 1. Fault-tolerant quantum AI agents. Are there universal fault-tolerant schemes that enable a noisy quantum AI agent to efficiently simulate any noiseless ideal quantum AI agent? Such schemes must preserve the agent's learning advantages while maintaining robustness against multiple simultaneous sources of imperfection across sensing, memory, and computation.

D. Verifiability

Quantum technology operates at the frontier of physics, where highly complex entanglement and superposition dominate. As we push toward large-scale quantum systems, we venture into regimes of complexity and entanglement never before encountered in physical experiments. This raises a profound question: Can we trust quantum mechanics to remain valid in these unexplored regimes?

The history of physics suggests caution. Just as Newton's laws break down at high velocities or in strong gravitational fields, superseded by Einstein's special or general relativity, quantum mechanics might behave differently in regimes of high complexity and large-scale entanglement. Our confidence in quantum theory comes primarily from experiments in relatively simple quantum systems with limited entanglement. We have never tested quantum mechanics at the complexity frontier where thousands of particles maintain complex quantum correlations. As we push more deeply into this complexity frontier, how can we trust that our quantum devices are actually achieving what we think they are?

This concern is not merely theoretical. The fragile nature of quantum systems means that errors can accumulate and propagate in subtle ways, potentially leading to incorrect outputs that might nevertheless appear plausible. Moreover, as we venture into regimes of high entanglement that have never been experimentally explored, our intuition and standard verification techniques may fail us. This challenge motivates our fourth criterion for meaningful quantum advantages:

Definition 4 (Verifiability). A quantum advantage is verifiable if it can be checked efficiently that the quantum technology is producing the correct output.

Verification serves a dual purpose in quantum technology. First, it ensures our quantum devices are functioning as intended. For some quantum advantages, this verification is straightforward. In quantum communication protocols demonstrating Bell inequality violations, the correlations between measurement outcomes directly verify the quantum advantage. Similarly, for quantum algorithms like Shor's factoring, the output can be efficiently verified by classical multiplication.

Second, and perhaps more profoundly, verification protocols can serve as stringent tests of quantum mechanics itself in previously unexplored regimes. When a highly complex quantum system behaves exactly as quantum theory predicts, especially in scenarios where classical simulation is provably impossible, we gain confidence in the theory's validity at higher scales of complexity. Conversely, any deviation from quantum predictions in these regimes could signal the discovery of new physics. However, verification becomes considerably more challenging for tasks involving complex quantum states or dynamics.

Example 12 (Classical verification of quantum computation [20]). A breakthrough in verifiable quantum advantage came with protocols that enable a classical computer to verify any quantum computation. These protocols allow a classical client to delegate quantum computations to a quantum server while maintaining both privacy and verifiability. The verification relies on cryptographic assumptions, such as the quantum hardness of the learning with errors problem. Through an elegant combination of cryptographic primitives and measurement-based quantum computation, these protocols ensure that any deviation from the prescribed quantum computation is detected with high probability. Most remarkably, the verification can be performed entirely by classical computers, providing a bridge between classical and quantum computation that enables rigorous testing of quantum advantages. This result demonstrates that a classical system can verify a quantum computation even when it cannot efficiently simulate the quantum computation.

While universal verification schemes exist for quantum computation, the landscape of verifiable quantum advantages in other domains remains largely unexplored. In quantum sensing, verification often relies on indirect evidence or statistical consistency checks. For quantum machine learning, verifying that a quantum model has truly learned better features or representations than alternative classical models remains challenging. This leads to several open challenges.

Open question 2: Verification beyond computation. Can we develop general verification schemes for quantum advantages in communication, learning, and sensing that go beyond computation? How can we verify outputs from quantum technology when we venture into new physical regimes?

E. Usefulness

The ultimate test of any technological advantage lies not in its theoretical power but in its practical impact. While demonstrating a quantum advantage that is predictable, generic, robust, and verifiable is itself a remarkable scientific achievement, these properties alone do not

guarantee that the advantage will be valuable in real-world applications. This observation leads to our final criterion:

Definition 5 (Usefulness). A quantum advantage is useful if the advantage provides practical value to a user who seeks an effective method for a significant application (and who does not care whether the underlying technology is classical or quantum).

The emphasis on practical value for users who do not care about the underlying physics is crucial. Many early quantum algorithms focused on problems specifically constructed to showcase quantum advantages, like random circuit sampling [44] or boson sampling [64]. While these demonstrate fascinating quantum-classical separations, their practical utility remains limited. In contrast, truly useful quantum advantages should address problems we genuinely want to solve, regardless of how we solve them.

Example 13 (Quantum simulation of materials). Quantum simulation of strongly correlated quantum materials exemplifies a potentially useful quantum advantage. Understanding and predicting the properties of complex materials like high-temperature superconductors, exotic quantum magnets, or catalyst molecules has immense scientific and technological value, independent of the methods used. Classical computers face fundamental challenges in simulating these systems due to the intricate quantum correlations present in many-body systems, while quantum computers can naturally represent and evolve these quantum states. However, the landscape of simulation methods continues to evolve rapidly. Classical algorithms, particularly those leveraging machine learning and tensor networks [65], have shown remarkable progress in describing previously intractable quantum systems. It remains challenging to predict whether future classical technological breakthroughs will enable efficient simulation of the quantum materials most relevant to practical applications.

The path to useful quantum advantages often involves identifying problems where quantum effects are inherent to the challenge at hand, rather than forcing quantum approaches onto classical problems. Moreover, usefulness must be evaluated in the context of rapidly evolving classical capabilities and the practical constraints of quantum technology. It might turn out that the most impactful quantum advantages are not those with the most impressive theoretical speedups, but instead the applications where quantum approaches offer qualitatively new capabilities or insights.

Finding a new quantum advantage that achieves all five keystone properties simultaneously remains a grand challenge for the field. This is stated in the following open question.

Open question 3: Ideal quantum advantage. What quantum advantages are simultaneously predictable, typical, robust, verifiable, and useful? Can we develop new approaches to discover quantum advantages that satisfy all these criteria?

III. REALMS OF QUANTUM ADVANTAGE

The vast world of quantum advantage extends far beyond computational speedups and presents a rich tapestry of opportunities where quantum resources enable capabilities beyond those achievable with classical resources alone. We can organize these quantum advantages into distinct realms, each characterized by fundamentally different sources of quantum superiority over classical approaches. These realms differ not only in their applications but in their fundamental nature. Some rely on unproven complexity assumptions, while others are guaranteed by the laws of physics themselves. For comprehensive surveys of quantum algorithms and protocols across these realms, we refer readers to existing resources [66–69]. See Fig. 1(b) for an illustration of the four realms.

A. Computational advantages

The computational realm constitutes the most extensively studied realm of quantum advantages, characterized by algorithms that promise faster solutions to computational problems. Yet this realm has a complex landscape where genuine advantages coexist with subtle illusions, and where the foundation of each advantage rests on computational complexity assumptions.

Complexity-theoretic foundation. Unlike other realms we will explore, computational quantum advantages fundamentally depend on unproven conjectures about the computational difficulty of certain problems. The canonical assumption $\mathsf{BPP} \neq \mathsf{BQP}$ asserts that quantum computers can efficiently solve some problems that are intractable for classical computers. While this assumption is supported by substantial evidence, it remains a conjecture, one that underpins virtually every claimed computational quantum advantage.

Oracle-based advantages. The theoretical power of quantum computation was first established in oracular settings [70,71], where an oracle for computing a function is given as a black box. Simon's algorithm [71] provided the first exponential advantage and directly inspired the development of Shor's algorithm. Other foundational results include Grover's search [72] and quantum walks [73–75] exhibiting polynomial to exponential oracular advantages. While oracle-based advantages provide important conceptual foundations for understanding quantum computational power, they are not computational advantages in the usual sense. The black box assumption means that when the oracle's internal structure is revealed, the quantum advantage could disappear as classical algorithms may find efficient shortcuts that exploit the revealed structure.

Shor's algorithm. Building on the theoretical insights from Simon's algorithm, Shor's algorithm [18] exemplifies genuine computational quantum advantage. It exploits the hidden subgroup structure that quantum computers can efficiently solve, providing a concrete path from

oracle-based advantage to real-world impact. As discussed in earlier sections, it satisfies our keystone criteria and demonstrates how quantum advantage can have profound real-world implications through its threat to cryptographic infrastructure worldwide. Even decades after its discovery, no efficient classical algorithm has emerged to challenge its speedup.

Quantum simulation. Beyond number-theoretic problems, quantum simulation represents a particularly natural computational application, as the task of simulating quantum many-body systems inherently involves quantum mechanical effects that classical computers struggle to capture efficiently [76,77]. The intuition is compelling: Quantum systems should naturally simulate quantum physics. Yet this intuitive advantage faces relentless competition from advancing classical methods. Sophisticated tensor network algorithms [65], machine learning approaches [78], and quantum-inspired classical methods [79] continue to push the boundaries of classical simulation. While quantum simulators have the potential to achieve advantages for specific problems, the landscape remains dynamic. What appears quantum advantageous today may be challenged by classical innovations tomorrow.

Power of data. A subtle challenge to computational quantum advantages emerges with the recognition of the power of data [80]. When quantum experiments generate data that cannot be efficiently produced classically, even classical machine learning algorithms trained on such data can solve computational problems beyond the reach of classical algorithms operating in isolation [81,82]. This reveals a deeper vulnerability: As we accumulate more data about our quantum universe through quantum simulations, quantum sensors, and quantum experiments, classical machine learning algorithms, by exploiting that data, might learn to approximate quantum phenomena without requiring additional quantum computation. However, assuming $\mathsf{BPP} \neq \mathsf{BQP}$, classical algorithms cannot universally simulate all quantum experiments. This suggests that important applications can be realized by quantum computing and classical AI working together; for example, quantum computers might collect data in previously unexplored physical regimes, empowering classical AI to predict properties of exotic quantum systems that have not yet been observed.

The realm of computational advantages presents a nuanced picture: While the quantum advantage exhibited by foundational algorithms like Shor's remains well grounded, the broader landscape is characterized by ongoing competition between quantum and classical computers. Quantum computational advantages require constant vigilance as the classical state of the art continues to advance.

B. Learning and sensing advantages

Moving beyond pure computation, the learning and sensing realm offers quantum advantages that emerge when quantum technology interfaces directly with the physical world. Here, advantages arise not from abstract algorithmic improvements, but from the fundamental quantum nature of information gathering, the physical systems we seek to understand, and the ability to process quantum information in ways that classical systems cannot. Crucially, many advantages in this realm can be established without relying on computational complexity assumptions.

Physics-based foundation. Unlike computational advantages that rely on unproven complexity-theoretic conjectures, nearly all known learning and sensing advantages rest on fundamental physical principles. The standard quantum limit in sensing, Heisenberg's uncertainty principle, and the inevitable disturbance of quantum states upon measurement all derive from the rigorous mathematical foundations of quantum mechanics, foundations that have been extensively validated over the past century. These laws of nature enable the development of learning and sensing advantages without requiring any assumptions about computational hardness. This physical grounding endows many learning and sensing advantages with an intrinsic robustness that computational advantages inherently lack.

Quantum sensing fundamentals and limits. Quantum sensing harnesses intrinsically quantum objects, such as atoms, ions, molecules, or photons, to measure physical quantities with enhanced performance. This quantum enhancement requires sophisticated quantum states, exemplified by Greenberger-Horne-Zeilinger (GHZ) states for magnetic field sensing [83,84]. As described in Sec. II C, achieving this quantum advantage faces a fundamental adversary: noise. The Hamiltonian-not-in-Lindbladian-span (HNLS) criterion reveals that a quantum enhancement survives in the presence of noise only when the target signal and noise processes can be distinguished at the operator level [54]. In generic scenarios, this condition fails, rendering the asymptotic quantum sensing advantage fundamentally unattainable even with quantum error correction [85]. This stands in stark contrast to quantum computation, where universal fault tolerance provides a clear path to noise robustness.

Practical quantum sensing applications. Despite theoretical limitations, quantum sensing has already achieved remarkable real-world successes. While finding examples where entanglement enhances sensing capabilities in a real-world setting has been famously elusive, a notable exception is LIGO's detection of gravitational waves, employing squeezed light to reduce quantum noise below the standard quantum limit [86–89]. In contrast, advantages exploiting the innate sensitivity of quantum sensors or ensembles thereof have been far more common. For example, nitrogen-vacancy centers in diamond enable quantum-enhanced magnetometry with nanoscale spatial resolution [90,91] or under extreme temperature or pressure conditions [92,93], while quantum-enhanced atomic clocks push temporal precision to unprecedented levels [94–97]. In sensing applications, even constant factor improvements over existing

technology can provide transformative practical utility. Unlike computational tasks, where performance can often be enhanced by distributing workloads across multiple processors, sensing tasks frequently face fundamental constraints imposed by their physical infrastructure. For instance, achieving nanoscale spatial resolution in magnetometry requires the sensor to be physically positioned at the target location, while gravitational wave detection demands interferometers with enormous 4-km-long arms. These physical constraints make it prohibitively expensive or physically impossible to simply deploy additional sensors to enhance performance. Quantum enhancements thus offer a fundamentally different value proposition than computational speedups.

Quantum learning agents. Recent developments have found that the most significant advantages emerge when we consider quantum machines that can see with quantum sensors, think with quantum computers, and remember with quantum memory [58–60,63,98–101]. Such quantum learning agents signify a paradigm shift beyond traditional quantum sensing approaches by enabling coherent processing of quantum information across multiple experimental runs.

Quantum-computing-enhanced sensing. When applied to sensing classical oscillating fields such as electromagnetic signals or gravitational waves, quantum agents could achieve beyond-quadratic polynomial speedups. While quantum algorithms like Grover search [72] were not traditionally considered relevant to sensing applications, recent work [101] demonstrates how quantum computers can be integrated with quantum sensors in a framework termed quantum-computing-enhanced sensing to learn properties of classical fields substantially faster than quantum sensors controlled by classical computers alone.

Exponential learning advantages. While the aforementioned advantages are polynomial, one may wonder whether exponential quantum advantages exist in the learning and sensing realm. The answer is proven to be affirmative when quantum learning agents are applied to study many-body physical systems and dynamics [58–60,63,98–100]. These quantum agents can learn from exponentially fewer experiments than classical approaches for tasks including predicting physical properties and learning approximate models of quantum dynamics. The exponential improvement emerges from the ability to preserve quantum correlations across multiple experiments that classical approaches must necessarily destroy through intermediate measurements, enabling quantum learning agents to extract exponentially more information per experiment than their classical counterparts.

Rather than relying on unproven complexity assumptions, quantum advantages in the learning and sensing realm often emerge from the quantum nature of physical reality itself. Looking ahead, the most transformative opportunities may emerge from quantum learning agents that seamlessly integrate quantum sensors and quantum computers. Some of these advantages have been proven to persist under realistic noise [60–63], but the general theory of fault-tolerant quantum agents remains to be developed.

C. Cryptographic, communication, and strategic game advantages

In contrast to both computational and learning and sensing advantages, the cryptographic realm encompasses quantum advantages that emerge from quantum mechanics' fundamental properties, such as entanglement, no cloning, and measurement disturbance, applied to security, communication, and coordination tasks. Unlike computational advantages that depend on unproven complexity assumptions, these advantages derive their power directly from the mathematical structure of quantum mechanics itself, providing unconditional security guarantees that remain inviolable as long as quantum mechanics remains an accurate description of our Universe.

Physics-based security. The defining characteristic of this realm is that quantum mechanics itself provides security guarantees impossible to achieve classically. The no-cloning theorem ensures that unknown quantum states cannot be accurately copied without detection. The measurement postulate guarantees that extracting information from a quantum system necessarily disturbs it, leaving unmistakable evidence if someone has been eavesdropping. Quantum entanglement enables correlations that fundamentally surpass classical limitations, as demonstrated by Bell's inequality. Together, these principles create a foundation for cryptographic protocols whose security is guaranteed by the laws of physics.

Quantum encryption and certified randomness. Quantum mechanics enables parties to establish secure communication channels and verify information authenticity with guarantees that transcend computational security. Quantum key distribution exemplifies this revolutionary capability: The BB84 protocol [102] enables Alice and Bob to establish a shared cryptographic key via quantum communication with security guaranteed by the laws of physics. Any eavesdropper necessarily disturbs the quantum transmission, revealing their presence through detectable changes in the quantum states. This represents a paradigmatic shift from classical cryptographic protocols, where security relies on the assumption that computations that could break the protocol in principle are too difficult to execute in practice. Such computational assumptions could potentially be undermined by future algorithmic breakthroughs or advances in computational power. Certified randomness generation [103,104] exploits these principles to create truly random bits that have not previously been generated prior to the protocol, whose unpredictability is guaranteed by quantum mechanics. Unlike classical pseudorandom number generators that are fundamentally deterministic, quantum randomness provides bits that cannot be spoofed by a secret seed or other back door. Furthermore,

the randomness is guaranteed even if the user does not trust the inner workings of the quantum device used in the protocol. Such device-independent randomness certification has profound implications for cryptographic applications requiring secure random bits.

Unforgeability protocols. The no-cloning theorem enables the creation of objects with inherent copy protection, opening new classes of applications impossible in classical physics. Quantum money schemes [105] aim to create unforgeable currency by using quantum states to encode serial numbers that can be efficiently verified but cannot be perfectly duplicated. Certified deletion protocols [106] enable Alice to send quantum-encrypted data to Bob, who can later provide cryptographic proof that he has permanently and verifiably destroyed the information. This capability is impossible with classical physics, where information can always be copied before deletion. Copy-protected software [107] and one-time programs [108] extend these principles to software distribution, enabling programs that cannot be duplicated while preserving their computational functionality. These protocols may revolutionize intellectual property protection, data privacy, regulatory compliance, and rights management by making such violations physically impossible rather than merely legally prohibited.

Quantum entanglement and strategic games. Quantum entanglement transcends classical communication limitations through nonlocal correlations that exhibit behaviors impossible in any local classical theory. As explored in Puzzle 1, entangled quantum systems demonstrate correlations that cannot be reproduced by any classical mechanism. In strategic scenarios requiring spatially separated parties to act simultaneously in a coordinated manner, quantum protocols could provide fundamental advantages over classical protocols. For example, Ref. [109] has illustrated how entangled quantum systems shared between spatially separated stock exchanges can provide coordination advantages in some high-frequency trading applications, where small time improvements can translate into significant competitive value.

The cryptographic, communication, and strategic game realm uses quantum mechanics to provide security guarantees that are categorically impossible in the classical world. Unlike computational advantages that may be undermined by unexpected algorithmic discoveries, these advantages are guaranteed by the mathematical structure of quantum mechanics. As quantum technologies mature, this realm could revolutionize how we approach security, privacy, and trust in our interconnected world.

D. Space advantages

The realm of space advantages addresses a fundamental challenge of our data-driven era: processing data streams that vastly exceed available memory. While many classical algorithms require memory proportional to input size, quantum algorithms can sometimes achieve exponential space savings by encoding classical information in the exponentially large Hilbert space of a quantum device.

Quantum compression for data streams. As demonstrated in Puzzle 3, quantum states can serve as exponentially compressed representations of classical data for specific computational tasks. While Holevo's bound [7] forbids the extraction of an exponential number of classical bits from a small quantum memory, certain computational tasks can be performed on the compressed quantum representation without full data reconstruction. This quantum compression becomes particularly powerful in streaming scenarios where massive data streams must be processed with limited memory. Classical algorithms must make irreversible decisions about which information to retain as each data element arrives. In contrast, quantum algorithms can maintain superpositions over many possible configurations simultaneously, effectively postponing these irreversible decisions until more information becomes available.

Applications and limitations. Recent work on directed maximum cut problems [110] demonstrates how quantum algorithms can solve natural graph problems using exponentially less memory than classical approaches when the input size vastly exceeds memory capacity. The vector-in-subspace problem studied by Raz [8] provides another example: Determining which of two subspaces contains a given vector requires exponential classical communication but only linear quantum communication. However, space advantages face important limitations. They typically apply only when input size significantly exceeds memory capacity; furthermore, inputting a large amount of data into a quantum system can be costly, and the computational task performed must be compatible with Holevo's bound. Despite these constraints, space advantages represent a promising direction as data generation accelerates.

This realm highlights how quantum superposition can address fundamental resource limitations in data processing and storage, offering more efficient approaches to existing computational tasks. While this realm remains the least explored, it holds significant potential for making previously infeasible large-scale computations tractable through exponential memory reductions.

IV. FUTURE OF QUANTUM ADVANTAGE

The past decades have witnessed remarkable progress in predicting quantum advantages through theoretical approaches. Using mathematical proofs and complexity-theoretic arguments, researchers have established quantum advantages for various computational, sensing, learning, and communication tasks. Yet this theoretical framework, while powerful, captures only part of the story. The history of classical computing offers a cautionary tale: Algorithms with the strongest provable advantages rarely become the most useful in practice, and the most transformative

applications of modern computers were often unpredicted or achieved without rigorous performance guarantees. This observation suggests we need a broader framework for identifying and understanding genuine quantum advantages.

Envision a future, perhaps 100 years from now, where quantum technology has become as ubiquitous as classical computers are today. A complete quantum technological infrastructure would enable us to probe the world with quantum sensors, preserve quantum information in reliable quantum memories, process quantum data with fault-tolerant quantum computers, and transmit quantum states through quantum networks. In this future, the competition between quantum and classical approaches would shift from theoretical guarantees to head-to-head empirical competition, much as classical algorithms compete today based on practical metrics rather than asymptotic bounds. How should we expand our framework for understanding such future quantum advantages?

In this section, we explore three perspectives that we expect to become increasingly more relevant in the future. First, we examine how empirical evidence, rather than theoretical guarantees, might drive the discovery and validation of quantum advantages. Second, we consider how quantum approaches could offer conceptual advantages, new ways of thinking about problems that turn out to be more natural or insightful than classical approaches. Finally, and most profoundly, we show that some quantum advantages are fundamentally unpredictable using classical resources alone, suggesting that the landscape of quantum advantages may be far richer than we can currently envision.

A. Empirical quantum advantages

First, we emphasize that empirical performance, rather than theoretical guarantees, will likely drive the adoption of quantum methods in practice. Key metrics will include empirical run-time, cost of implementation, and total time to solution. This shift mirrors the evolution of classical computing, where widely used algorithms often differ from those with the strongest asymptotic guarantees. When focusing on run-time, empirical advantages can manifest in multiple ways. The most straightforward are constant-factor improvements on practical problem instances; these may not grab the attention of theorists but may be very important to users. It is also notable that the classical algorithms achieving the best asymptotic scaling might not be the most practically useful due to prohibitively large constant factors. Multiplication of $n \times n$ matrices provides a classic example. Although theoretically superior algorithms exist with better asymptotic complexity, the simple algorithm with $\mathcal{O}(n^3)$ steps remains standard in practice due to its smaller constants and cache efficiency.

Quantum advantage may also emerge from faster total time to solution starting from the problem or concept going all the way to the solution, which we term "concept to solution." Concept to solution includes not only the algorithm's run-time, but the entire process from problem formulation to working solution. This encompasses the human time spent designing experimental protocols, implementing algorithms, and iteratively refining them. A practical quantum advantage might occur if quantum technology enables faster overall solution time, even if it turns out that classical algorithms with similar run-time exist but are not discovered until much later.

Example 14 (Quantum advantage in concept to solution). Consider a scenario where designing a quantum algorithm for a new computational problem takes an hour and the quantum computation itself solves the problem in about one second. Even if a talented computer scientist could eventually discover a classical algorithm with a one-second classical computational time, if that discovery requires a year of dedicated research, the quantum approach still offers a significant advantage in concept to solution: an hour versus a year. This form of advantage, where quantum methods provide a more direct path to efficient solutions, may prove particularly valuable as quantum technology matures.

Empirical quantum advantage may also arise from generality. A single general-purpose quantum algorithm is often immediately applicable to an entire class of systems, whereas classical approaches typically require bespoke algorithms that must be individually designed, implemented, and tuned for each instance. Even when an efficient classical algorithm exists for every system in the class, discovering and deploying it demands time and expertise that must be reinvested for each new instance. A user who needs solutions across many systems may therefore find the general-purpose quantum approach to offer a significant overall advantage, even if it is not asymptotically better for any single instance. Quantum simulation offers a vivid illustration of this point: A natural classical competitor is simply building a physical experiment in the laboratory to study the system of interest. While such an experiment is in some sense "constant time" per system, it cannot be reprogrammed to simulate a different system the way a quantum computer can. This programmability and generalizability across a broad class of systems may thus be a key source of practical value that pure run-time comparisons fail to capture.

Large-scale quantum computing might also reveal new empirical benefits that transcend theoretical guarantees. Modern machine learning offers a preview of such possibilities. Large language models have demonstrated capabilities far beyond theoretical predictions. Similarly, as quantum systems scale up in size and architectural complexity, they may exhibit unexpected emergent behaviors that, while difficult to characterize theoretically, provide transformative practical advantages.

Such empirical quantum advantages represent a significant departure from how we typically evaluate quantum

technology today. Rather than focusing solely on asymptotic complexity and provable speedups, we envision a future where quantum methods might be preferred for their practical benefits: faster solution development, greater generality across problem classes, simpler implementation, or emergent capabilities that arise from scale. While such advantages may be harder to prove rigorously, they could ultimately drive the widespread adoption of quantum technology.

B. Conceptual quantum advantages

Second, we highlight the value of looking beyond physical resources like time, space, cost, and energy to consider broader forms of advantage. Modern philosophy of science recognizes that empirically equivalent theories can differ dramatically in their conceptual value. The heliocentric model of the Solar System, while making essentially identical predictions to certain geocentric models, provided a simpler framework that transformed our understanding of the cosmos and catalyzed new scientific directions.

We propose that quantum technology might offer similar conceptual advantages. Consider two solutions to a problem, one classical, one quantum, with similar resource requirements. Even if they offer comparable performance, the quantum solution might be conceptually simpler, providing a more natural framework for understanding the problem. This conceptual clarity could manifest in multiple ways: easier reasoning about correctness, more straightforward integration with other components, or more natural extensions to related problems.

The value of such conceptual advantages extends beyond mere aesthetic preference. Simpler conceptual frameworks often reflect deeper underlying structure, potentially leading researchers to new theoretical insights and technological directions. Just as quantum mechanics itself provided new ways of understanding physical phenomena, quantum approaches to practical problems might offer new perspectives that catalyze further advances. While such advantages might seem secondary in today's resource-constrained quantum era, they could become crucial in a future with ubiquitous quantum technology, forming cornerstones for new theoretical and technological developments.

C. Unpredictable quantum advantages

Our third and most profound perspective challenges a core assumption of our current framework: that we can predict quantum advantages using classical resources and mathematical reasoning. We prove that this assumption fails even for well-defined computational problems, revealing fundamental limits to our ability to identify quantum advantages using classical approaches alone. This result illustrates concretely the more general principle that some emergent quantum advantages may be quite unexpected when they are discovered.

To establish this limitation rigorously, we consider a concrete computational problem regarding whether a specified computational task admits a quantum advantage or not. This problem is inspired by the field of metacomplexity, the study of the complexity of computational problems which address questions about complexity theory itself [111,112]. Earlier, we discussed how Pauli propagation provides an efficient classical method for simulating most quantum circuits, though large families of circuits remain classically intractable. This naturally leads us to ask, given a specific quantum circuit, can we determine whether executing it on a quantum computer would outperform this classical simulation method?

Theorem 1 (Classical hardness in predicting quantum advantages, informal). Consider the decision problem of predicting whether executing a particular quantum circuit on a quantum computer has a computational advantage over the Pauli propagation classical simulation method on that same circuit. Assuming $\mathsf{BPP} \neq \mathsf{BQP}$, this decision problem exhibits a quantum advantage: While a quantum computer can efficiently solve this problem, no classical algorithm can solve it efficiently.

This theorem reveals a profound circularity in our quest for quantum advantages: The very act of predicting quantum advantages is itself a problem requiring quantum computation. Since Pauli propagation represents only one classical simulation method, predicting quantum advantage against any classical method, including those not yet conceived, is even more challenging. The proof, provided in Appendix D, shows that any classical algorithm that can determine efficiently whether or not running a given quantum circuit on a quantum computer has a quantum advantage over a classical heuristic can be adapted to create a classical algorithm that is able to efficiently solve any problem in BQP. Thus, detecting quantum advantage is classically hard, assuming that $\mathsf{BPP} \neq \mathsf{BQP}$. Furthermore, determining whether executing a quantum circuit exhibits quantum advantage is quantumly easy; it suffices to run both the quantum circuit and the classical heuristic on a polynomial number of randomly chosen instances and then check whether or not the results typically agree.

This establishes a fundamental barrier: Classical technology cannot reliably distinguish genuine quantum advantages from pseudoadvantages, confirming that quantum advantages exist beyond the predictive power of classical computers. This leads us to a profound implication:

> There are plenty of problems with genuine quantum advantages such that we cannot predict the advantage using only classical technology.

Paradoxically, to fully map out the landscape of quantum advantages, we must use the very quantum technologies whose power we are trying to characterize.

This unpredictability extends beyond computational problems to the frontiers of scientific discovery. Quantum technologies are already kindling new insights into the behavior of highly entangled quantum many-body systems; recent examples include quantum many-body scars and deep thermalization [113–116]. What marvelous phenomena are yet to be found? We cannot possibly know if our classical technologies are intrinsically incapable of predicting the phenomena. Thus a rich panoply of surprises may be unveiled as we probe the physical universe with ever more advanced quantum devices.

D. Path forward

We look forward to a future where a broad spectrum of quantum advantages come to light: some empirically discovered, others conceptually transformative, and many fundamentally unpredictable using current technology. Just as the most impactful applications of classical computers were unimaginable to their early pioneers, the true potential of quantum technology might only reveal itself through a journey of development and discovery.

We are not in that future yet. Today's quantum technology is limited, and mathematical analysis provides our best guide for identifying genuine quantum advantages. While mathematical rigor will always be essential, we have argued that theoretical analysis alone cannot capture the full landscape of potential quantum advantages awaiting discovery. For now, we face the challenge of developing new ideas and methods to navigate the vast world of quantum advantage, guided by, but not confined to, what is mathematically provable.

ACKNOWLEDGMENTS

The authors thank Sergio Boixo, Ryan Babbush, Edward Farhi, Sam Gutmann, and Stephen Jordan for a careful reading and helpful comments. We would also like to thank Robbie King and Mikhail Lukin for illuminating discussions, and Chi-Yun Claudia Cheng for discussions and assistance regarding the main figure. H.-Y. H. acknowledges support from the U.S. Department of Energy, Office of Science, National Quantum Information Science Research Centers, Quantum Systems Accelerator and the Broadcom Innovation Fund. S. C. acknowledges the support from the Center for Ultracold Atoms, an NSF Physics Frontiers Center (Grant No. 2317134). J. P. acknowledges support from the U.S. Department of Energy, Office of Science, Accelerated Research in Quantum Computing, Quantum Utility through Advanced Computational Quantum Algorithms (QUACQ) and Fundamental Algorithmic Research toward Quantum Utility (FAR-Qu), from the U.S. Department of Energy, Office of Science, National Quantum Information Science Research Centers, Quantum Systems Accelerator, and from the Institute for Quantum Information and Matter, an NSF Physics Frontiers Center (Grant No. PHY-2317110).

DATA AVAILABILITY

There are no publicly available research data or software supporting this manuscript. Requests for further information or data should be sent to the authors.

APPENDIX A: SOLUTIONS TO PUZZLES

1. Puzzle 1: Is this spooky?

Let us analyze this puzzle in two steps: first considering measurements in a single basis, then examining what changes when we allow measurements in multiple bases.

First scenario: Single-basis measurements. When Alice and Bob measure only in the standard basis $\{|\uparrow\rangle, |\downarrow\rangle\}$, their quantum state,

$$\frac{1}{\sqrt{2}}(|\uparrow\rangle \otimes |\uparrow\rangle - |\downarrow\rangle \otimes |\downarrow\rangle), \tag{A1}$$

produces perfectly correlated outcomes. This correlation appears mysterious. When Alice measures her spin and obtains $|\uparrow\rangle$, Bob's spin instantly collapses to $|\uparrow\rangle$, and similarly for $|\downarrow\rangle$. However, this behavior can be perfectly simulated using classical objects. Consider the following classical protocol.

(1) Take a pair of socks that are either both black or both white with equal probability.
(2) Place each sock in an opaque gift box.
(3) Give one box to Alice and one to Bob.
(4) Let them separate to arbitrary distances.

When Alice and Bob open their boxes, they see perfectly correlated colors instantaneously. The classical protocol exactly reproduces the quantum statistics: Each party sees a random outcome (up-down or black-white) with 50-50 probability, and the outcomes are always perfectly correlated.

Second scenario: Multiple-basis measurements. The situation changes dramatically when Alice and Bob can each randomly choose between measuring in the standard basis $|\uparrow\rangle, |\downarrow\rangle$ or the rotated basis $|\nearrow\rangle, |\swarrow\rangle$. Now they observe correlations that cannot be reproduced by any classical system. While the measurement procedures themselves are quantum, the remarkable feature is that the resulting statistical patterns in their classical outcomes, i.e., the sequences of bits they record, exhibit correlations that no classical mechanism can generate. The classical impossibility was first proven by Bell [1] and has been experimentally verified multiple times [117,118]. The correlations exhibited by quantum entanglement in multiple measurement bases violate Bell's inequality, providing rigorous proof that quantum systems can achieve statistical correlations among classical outcomes that are forbidden to any classical theory. This elucidates a genuine quantum

advantage. We refer readers interested in a detailed pedagogical treatment of Bell's inequality and its implications to Chap. 4 of Ref. [119].

2. Puzzle 2: Exponential quantum advantage in analyzing big data?

The solution to Puzzle 2 reveals how classical techniques can match quantum computation for this specific task. The key insight is that Bob's classical RAM can be enhanced to provide two crucial capabilities, each taking only $\mathrm{poly}(n)$ time:

(1) *Query:* Given any indices k and i, output the value $x_{k,i}$.
(2) *Sample:* Given index k, sample i with probability $p_k(i) = |x_{k,i}|^2$.

The construction of efficient sampling access from standard RAM is a well-known technique in randomized algorithms [120,121]. The goal is to sample from a discrete probability distribution where outcome i has probability $p_k(i) = |x_{k,i}|^2$. The naive approach would require $\mathcal{O}(2^n)$ time per sample by iterating through all entries, but preprocessing enables much faster sampling.

For each vector $\vec{x}_k$, we first compute the cumulative distribution function:

$$F_k(j) = \sum_{i=1}^{j} |x_{k,i}|^2. \tag{A2}$$

Since $\vec{x}_k$ is normalized, we have $F_k(2^n) = 1$. We then store these cumulative probabilities in a binary tree data structure where each leaf corresponds to an index $i \in \{1, 2, \ldots, 2^n\}$, and each internal node stores the sum of probabilities in its subtree. To sample from this distribution, we generate a uniform random number $r \in [0, 1]$ and use binary search to find the unique index i such that $F_k(i-1) < r \leq F_k(i)$. The binary tree structure allows us to traverse from the root to the appropriate leaf in $\mathcal{O}(\log 2^n) = \mathcal{O}(n)$ time: At each internal node, we compare r with the left subtree's total probability and branch left or right accordingly. This preprocessing requires $\mathcal{O}(2^n)$ time per vector to compute all cumulative probabilities and construct the tree, similar to the overhead required by QRAM. Once completed, this approach enables $\mathcal{O}(n)$ time sampling, making the sampling step efficient.

Using these capabilities, Bob can employ importance sampling to estimate the inner product efficiently. The key idea is to define a random variable $X_{k,\ell}$ as follows: First, sample i according to probability $p_k(i)$, then let $X_{k,\ell} = (x_{\ell,i}/x_{k,i})$. A straightforward calculation shows that

$$\mathbb{E}[X_{k,\ell}] = \sum_{i=1}^{2^n} p_k(i) \cdot \frac{x_{\ell,i}}{x_{k,i}} = \sum_{i=1}^{2^n} |x_{k,i}|^2 \cdot \frac{x_{\ell,i}}{x_{k,i}} = \vec{x_k} \cdot \vec{x_\ell}, \tag{A3}$$

$$\mathrm{Var}[X_{k,\ell}] = \sum_{i=1}^{2^n} p_k(i) \cdot \left(\frac{x_{\ell,i}}{x_{k,i}}\right)^2 - (\mathbb{E}[X_{k,\ell}])^2 \leq 1. \tag{A4}$$

The bounded variance is crucial: It means Bob can estimate the inner product to precision $1/\mathrm{poly}(n)$ by averaging $\mathrm{poly}(n)$ independent samples of $X_{k,\ell}$. This classical randomized algorithm matches the performance of the quantum algorithm exactly.

This solution exemplifies a broader pattern in quantum algorithm analysis. Similar classical techniques have appeared across different fields, including neural quantum states [78], distributed computing for inner product estimation [122], randomized linear algebra [123], and the dequantization of quantum algorithms [5]. These examples demonstrate how classical probabilistic techniques can sometimes achieve the same performance as quantum algorithms when operating under comparable data-access assumptions.

It is worth noting that the picture changes if we consider a different communication model. Suppose Alice and Bob do not communicate directly with each other, but instead each sends a message to a referee whose task is to compute the inner product $\vec{x}_k \cdot \vec{x}_\ell$. Since any protocol for inner product estimation can be used to solve equality testing, the $\Theta(2^{n/2})$ classical lower bound for equality testing established by Buhrman *et al.* [124] carries over: Any classical protocol requires $\Omega(2^{n/2})$ bits of communication, whereas quantum fingerprints of only $\mathcal{O}(n)$ qubits suffice. However, this quantum advantage disappears once Alice and Bob are permitted to share a random key, in which case classical protocols of constant communication cost suffice. This further illustrates how the presence or absence of a quantum advantage can depend sensitively on the precise communication model and available resources.

3. Puzzle 3: Can quantum systems encode exponential classical information?

The answer to this puzzle is a nuanced "yes," that reveals a subtle but powerful form of quantum advantage. The apparent conflict between exponential encoding and Holevo's bound is resolved by understanding that the usefulness of an encoding depends entirely on the task for which the information is used. Holevo's bound [7] states that it is impossible for Bob to learn all 2^n classical values x_i from the following n-qubit state:

$$|x\rangle = \sum_{i=1}^{2^n} x_i |i\rangle. \tag{A5}$$

Any quantum measurement strategy he employs can reveal at most $\mathcal{O}(n)$ bits of classical information about the 2^n-dimensional vector $\vec{x}$. If Alice's goal were to transmit the entire classical vector to Bob, then quantum state

encoding offers no asymptotic advantage over sending n classical bits.

However, the story changes if Bob's goal is not to reconstruct the entire vector, but to perform a specific computation on it. For some computational tasks, the quantum state encoding provides a genuine, exponential advantage in communication. A landmark result by Raz [8] established this rigorously. He defined a relational problem where Alice is given a unit vector $v \in \mathbb{R}^{2^n}$ and Bob is given one of two possible subspaces, P_0 or P_1, each of dimension $2^n/2$. The vector v is promised to lie in one of these two subspaces, and their task is to determine which one. Raz proved that any classical communication protocol for this problem requires exchanging at least $\Omega(2^{n/2})$ bits. In contrast, a quantum protocol can solve the problem with exponentially less communication: Alice simply encodes her vector into an n-qubit state $|v\rangle$ and sends it to Bob, who can then perform a measurement to identify the correct subspace. This establishes an exponential separation between the $\mathcal{O}(n)$ quantum communication complexity and the $\Omega(2^{n/2})$ classical communication complexity for this particular task.

This example demonstrates that for certain problems, the quantum state $|x\rangle$ acts as an exponentially compressed quantum representation of the classical data. While this compressed encoding does not allow full reconstruction of the original data, it contains enough information to answer specific questions about it exponentially more efficiently than any classical representation of similar size. Therefore, Alice's quantum compression scheme is not an illusion. It provides a genuine quantum advantage, but this advantage is in *communication complexity* for specific relational problems, not in general-purpose data transmission. This principle has found applications in quantum machine learning, for example, in protocols for learning quadratic functions of large datasets, where quantum communication can be exponentially more efficient than classical communication [11]. Alice and Bob did indeed gain an advantage, provided their future data processing tasks are of the right kind.

APPENDIX B: FRAGILITY OF QUANTUM ADVANTAGE IN NOISY SENSING

Here we illustrate the claims in Example 11 in a simple model of quantum sensing. We will explain how entanglement can be exploited to improve sensitivity in the noiseless case and prove that this sensitivity enhancement does not survive in the presence of noise.

1. Entanglement-enhanced sensitivity

In the noiseless version of this model, our task is to distinguish between two single-qubit unitary channels. $\mathcal{C}_\theta$ rotates the qubit about the Z axis by the positive angle θ, and $\mathcal{C}_0$ acts trivially:

$$\mathcal{C}_\theta[\rho] = e^{-i(\theta/2)Z}\rho e^{i(\theta/2)Z} \quad \text{and} \quad \mathcal{C}_0[\rho] = \rho. \tag{B1}$$

The rotation angle θ is the strength of the signal we wish to detect. Our goal is to detect the signal as quickly as possible, that is, to distinguish between the two cases using as few resources as possible. We will measure sensing efficiency in terms of the *sensing time* T, defined as the wall-clock time required to complete the protocol assuming one channel use takes one unit of time. Alternatively, we may fix the sensing time T and seek the smallest signal strength θ that can be distinguished from zero. This minimal detectable signal strength is called the *sensitivity*. We are most interested in the regime $0 < \theta \ll 1$, which will be assumed below.

When only a single sensing qubit is available, the optimal sensing protocol is well understood: Prepare the qubit in a superposition state $|+\rangle = (1/\sqrt{2})(|0\rangle + |1\rangle)$, apply the channel $T \approx \pi/\theta$ times sequentially, and measure the qubit in the X basis. Under $\mathcal{C}_0$, the qubit remains in the initial state, and the measurement yields $|+\rangle$ deterministically. Under $\mathcal{C}_\theta$, the qubit rotates by the total angle $T\theta \approx \pi$, and the measurement produces $|-\rangle$ with high probability, enabling us to distinguish the two cases. This procedure can be repeated a few times to amplify our confidence. Thus, one can resolve a small angle θ within a sensing time of $T = \mathcal{O}(1/\theta)$.

This analysis directly generalizes to N sensing qubits working in parallel. We now consider the channel $\mathcal{C}_x^{\otimes N}$ with $x = 0$ or θ, which applies the single-qubit channel $\mathcal{C}_x$ to all N qubits simultaneously. In this case, one can readily shorten the sensing time to $T = \mathcal{O}[1/(\sqrt{N}\theta)]$ by performing N copies of the single-particle protocol in parallel. Here, T represents the number of times the parallel channel $\mathcal{C}_x^{\otimes N}$ is applied, which equals the wall-clock time since all N qubits are processed simultaneously in each time step. More explicitly, each particle is prepared in the $|+\rangle$ state, the channel $\mathcal{C}_x^{\otimes N}$ is applied $T = \lceil \pi/(\sqrt{N}\theta) \rceil$ times, and then all qubits are measured. Recall that T, the number of (parallel) channel uses, is an integer; we are implicitly considering the case where $\sqrt{N}\theta \ll 1$. The measurement of each particle yields $|-\rangle$ with probability $\Omega(1/N)$ when $x = \theta$. Since a single $|-\rangle$ outcome heralds the presence of the signal, the two channels can be distinguished with high confidence after a constant number of repetitions. Thus, this protocol achieves a sensitivity scaling as

$$\theta^{-1} = \mathcal{O}(\sqrt{N}T), \tag{B2}$$

which one can show is optimal for any protocol that does not use entangled states.

Using entangled quantum states, the sensing time can be further improved to $T = \mathcal{O}[1/(N\theta)]$, where we now assume $N\theta \ll 1$. Consider the GHZ state $|\mathrm{GHZ};+\rangle = (|0\rangle^{\otimes N} + |1\rangle^{\otimes N})/\sqrt{2}$ evolved under the parallel channel

$\mathcal{C}_\theta^{\otimes N}$ repeated $T = \lceil \pi/(N\theta) \rceil$ times. The resulting state has a large overlap with $|\mathrm{GHZ}; -\rangle$, which is orthogonal to the initial state. Thus, by performing measurements in the $\{|\mathrm{GHZ}; \pm\rangle\}$ basis, one can distinguish the two channels within time $T = \mathcal{O}[1/(N\theta)]$. The reduction in sensing time by the factor $\sqrt{N}$ translates to an improved sensitivity arising from the multipartite entanglement in the GHZ state:

$$\theta^{-1} = \mathcal{O}(NT). \tag{B3}$$

One can show that this is the optimal sensitivity scaling using N sensing qubits. This entanglement-enhanced sensitivity is an example of quantum sensing advantage.

2. Fragility of entanglement-enhanced sensitivity in noisy sensors

We now consider a simple model of noisy quantum sensing. In realistic settings, noise arises from a variety of physical sources such as thermal fluctuations in the environment, imperfections in laser beams or microwave control pulses, and intrinsic physical properties of atoms or molecules. In our simplified model, sensing qubits are subject to rotation by an angle $\theta + \theta_n$, where θ is the signal to be detected and θ_n represents noise. The signal strength θ is assumed to be small and has the same value every time $\mathcal{C}_\theta$ is applied. The noise θ_n afflicts both $\mathcal{C}_\theta$ and $\mathcal{C}_0$ and is randomly fluctuating. We assume that θ_n is drawn independently from a Gaussian distribution with mean zero and variance γ each time either channel is applied. Our task is to determine whether $\theta = 0$ or not.

More explicitly, we need to distinguish between two channels:

$$\tilde{\mathcal{C}}_\theta \equiv \mathbb{E}_{\theta_n}[\mathcal{C}_{\theta+\theta_n}] = [(1-p)\mathcal{I} + p\mathcal{D}] \circ \mathcal{C}_\theta, \tag{B4}$$

$$\tilde{\mathcal{C}}_0 \equiv \mathbb{E}_{\theta_n}[\mathcal{C}_{\theta_n}] = [(1-p)\mathcal{I} + p\mathcal{D}] \circ \mathcal{C}_0, \tag{B5}$$

where $p = 1 - e^{-\gamma/2}$ is the probability that the qubit dephases, $\mathcal{I}$ is the identity channel $\mathcal{I}[\rho] = \rho$, and $\mathcal{D}[\rho] = |0\rangle\langle 0|\rho|0\rangle\langle 0| + |1\rangle\langle 1|\rho|1\rangle\langle 1|$ is the completely dephasing channel in the Z basis. In other words, the density operator's diagonal elements in the Z basis are unaffected by the Gaussian dephasing noise, while off-diagonal elements are suppressed by the factor $e^{-\gamma/2}$. We are especially interested in the high-sensitivity regime where $\theta \ll \gamma$.

For this noisy quantum sensing model, a lower bound on the sensing time is captured by the following theorem.

Theorem 2 (Noisy small-angle detection). Any quantum algorithm that correctly distinguishes two quantum channels $\tilde{\mathcal{C}}_0^{\otimes N}$ and $\tilde{\mathcal{C}}_\theta^{\otimes N}$ with probability $> 2/3$ must query the channel at least T times, where

$$NT = \Omega\left[\max\left(\frac{\gamma}{\theta^2}, \frac{1}{\theta}\right)\right]. \tag{B6}$$

We note that quantum algorithms may involve arbitrarily many and arbitrarily fast gate operations as well as an unlimited number of ancilla qubits and unbounded memory coherence times.

Theorem 2 implies that for any small constant noise strength γ, the sensing time T scales inverse quadratically in the signal strength θ. Thus, for any constant γ, the asymptotic sensitivity scales as

$$\theta^{-1} = \mathcal{O}(\sqrt{NT}). \tag{B7}$$

Importantly, this optimal sensitivity scaling can be attained by a simple, fully separable protocol that requires no entanglement, demonstrating that entanglement does not offer an asymptotic sensitivity advantage in this noisy quantum sensing model. Intuitively, dephasing noise destroys the phase coherences that enable entangled states like GHZ to accumulate phase collectively, eliminating the quantum advantage observed in the noiseless case.

The simple, fully separable protocol works as follows. We focus on the regime $\theta, \gamma \ll 1$ and $\theta \ll \gamma$ to simplify our analysis. Each sensing qubit is initialized in the $|+\rangle$ state, undergoes evolution under the channel $\lceil 1/\gamma \rceil$ times, and is then measured in the Y basis, yielding one of two outcomes: $|\pm i\rangle = (1/\sqrt{2})(|0\rangle \pm i|1\rangle)$. This procedure is repeated K times for N qubits in parallel, producing a total of KN measurement outcomes within sensing time $T = K\lceil 1/\gamma \rceil$. Finally, one counts the fraction of $|+i\rangle$ outcomes, $\hat{f}$, among all measurements.

In the absence of the signal, both outcomes $|\pm i\rangle$ occur with exactly the same probabilities; hence the probability distribution of $\hat{f}$ is centered around $\frac{1}{2}$ with variance $1/4KN$. In the presence of the signal, the measurement yields $|\pm i\rangle$ with probabilities $\frac{1}{2} \pm \epsilon$ with positive bias $\epsilon = \sin(\theta\lceil 1/\gamma \rceil)e^{-\gamma\lceil 1/\gamma \rceil/2}/2$, and the distribution of $\hat{f}$ is centered around $\frac{1}{2} + \epsilon$ with variance $1 - 4\epsilon^2/4KN$. Therefore, by checking whether $\hat{f}$ is larger than a threshold $\frac{1}{2} + (\epsilon/2)$, one can distinguish the two cases. The failure probability can be bounded using Chebyshev's inequality and becomes small whenever the standard deviation is sufficiently small compared to the bias: $\epsilon \gg (1/\sqrt{4KN})$, which implies $NT = \mathcal{O}(e^{\gamma\lceil 1/\gamma \rceil}/\lceil 1/\gamma \rceil/\theta^2)$ is sufficient. This matches the optimal sensitivity scaling for any fixed constant γ.

We now proceed to the proof of Theorem 2.

Proof. We prove two lower bounds separately: (i) $NT = \Omega(\gamma/\theta^2)$ and (ii) $NT = \Omega(1/\theta)$. The second bound follows directly from the noiseless case. This can be shown by noting that access to noiseless channels is at least as powerful as access to noisy channels because the latter can be simulated from the former by artificially

adding dephasing channel after the noiseless channel. As a result, the sensing time for the noisy case cannot be shorter than the noiseless case, which requires $NT = \Omega(1/\theta)$.

To prove the first bound, we employ a reduction to classical hypothesis testing. We introduce a hypothetical agent called *Demigod*, who is at least as powerful as any quantum agent capable of running arbitrary quantum algorithms. We show that even the powerful Demigod requires $NT = \Omega(\gamma/\theta^2)$ channel uses to distinguish the two channels. Thus, quantum agents, who are no more powerful than Demigod, require at least as many channel uses as Demigod to perform the distinguishing task.

We define the hypothetical agent Demigod to have the following enhanced capabilities.

(i) Unlimited computational power: Demigod can perform arbitrary classical or quantum computations instantaneously and with unlimited memory.

(ii) Direct observation of rotation angles: While quantum agents are given the ability to apply noisy quantum channels to quantum states, Demigod directly observes the noisy rotation angles, which are either θ_n or $\theta + \theta_n$ for signal θ and noise θ_n.

More precisely, when the unknown channel is $\tilde{\mathcal{C}}_\theta^{\otimes N}$, each query provides Demigod with N real numbers $\{\theta + \theta_{n,1}, \theta + \theta_{n,2}, \ldots, \theta + \theta_{n,N}\}$, where each $\theta_{n,i}$ is independently drawn from $\mathcal{N}(0,\gamma)$ with mean zero and variance γ. When the unknown channel is $\tilde{\mathcal{C}}_0^{\otimes N}$, she receives the list $\{\theta_{n,1}, \theta_{n,2}, \ldots, \theta_{n,N}\}$.

Since Demigod can query the rotation angles, she can perfectly simulate any quantum agent. To simulate a query to the noisy quantum channel $\tilde{\mathcal{C}}_\theta^{\otimes N}$ or $\tilde{\mathcal{C}}_0^{\otimes N}$ made by the quantum agent, she uses her enhanced ability to directly observe the list of rotation angles and rotates each of her N sensing qubits about the Z axis by the corresponding observed angle. Because of her unlimited computational power, she can also simulate any quantum computation performed by the quantum agent. Thus, she can perfectly simulate any quantum agent regardless of the quantum state preparation, evolution, channel query, or measurement strategy employed by the original quantum algorithm. Therefore, Demigod is at least as powerful as any quantum agent who tries to distinguish the two noisy channels.

After T queries to the parallel channel, Demigod has collected NT rotation angles, each angle is independently and identically drawn from either $P = \mathcal{N}(0,\gamma)$ or $Q = \mathcal{N}(\theta,\gamma)$ depending on whether the noisy channel is $\tilde{\mathcal{C}}_0$ or $\tilde{\mathcal{C}}_\theta$, respectively. For Demigod, the channel distinguishing problem reduces to the classical statistical problem of distinguishing between these two Gaussian distributions. The fundamental limit for this classical hypothesis testing problem is well established in information theory. To distinguish between P and Q with constant success probability, the number of samples must satisfy

$$\text{Number of samples} = \Omega\left(\frac{1}{D_{\mathrm{KL}}(P\|Q)}\right), \tag{B8}$$

where $D_{\mathrm{KL}}(P\|Q)$ is the Kullback-Leibler divergence. For Gaussian distributions:

$$D_{\mathrm{KL}}[\mathcal{N}(0,\gamma)\|\mathcal{N}(\theta,\gamma)] = \frac{\theta^2}{2\gamma}. \tag{B9}$$

Therefore, even Demigod requires at least $\Omega(\gamma/\theta^2)$ samples, establishing $NT = \Omega(\gamma/\theta^2)$. This lower bound on the number of samples applies no matter how much classical or quantum computing power Demigod applies to her task. Taking the maximum of both bounds yields $NT = \Omega\{\max[(\gamma/\theta^2), (1/\theta)]\}$. Thus, we obtain the desired lower bound on NT, completing the proof of Theorem 2. ∎

APPENDIX C: CLASSICAL HEURISTIC FOR SIMULATING QUANTUM CIRCUITS

1. Overview and context

The low-weight Pauli propagation algorithm [35–43] provides a powerful classical heuristic for simulating quantum circuits. The key idea is to truncate high-weight Pauli terms during the simulation, maintaining only the most important contributions. Consider a quantum circuit with L layers:

$$U = U_L U_{L-1} \cdots U_1. \tag{C1}$$

Our goal is to approximate expectation values of the form:

$$f_U(O) = \mathrm{tr}[U\rho U^\dagger O], \tag{C2}$$

where ρ is an input state and O is an observable.

2. Algorithm description

The low-weight Pauli propagation algorithm proceeds in the Heisenberg picture, where the observable is evolved backward through the circuit.

(1) We begin by initializing the operator O_L as the projection of the observable O onto the subspace of Pauli operators with weight at most k. Here, the weight of a Pauli operator is the number of non-identity terms it contains.

(2) Working backward from $j = L$ down to 1, we iteratively apply the following two steps:

(a) Evolve the operator through the jth layer: $O'_{j-1} = U_j^\dagger O_j U_j$.

(b) Project the result back onto the low-weight Pauli basis to get the new operator:

$$O_{j-1} = \sum_{\substack{P \in \{I,X,Y,Z\}^{\otimes n} \\ |P| \le k}} \frac{\mathrm{tr}[O'_{j-1}P]}{2^n} P. \quad \text{(C3)}$$

(3) The final estimate for the expectation value is computed using the fully evolved and truncated operator O_0 and the initial state ρ:

$$\text{Estimate} = \mathrm{tr}[O_0\rho]. \quad \text{(C4)}$$

The larger k is, the more accurate the low-weight Pauli propagation algorithm is. In practice, one can typically consider k to be a small constant. Even $k = 1$ can perform quite well in many cases. This completes the full description of the low-weight Pauli propagation algorithm.

3. Theoretical guarantees

When each circuit layer U_j is drawn from a *locally scrambling* distribution [125–129], which corresponds to a probability distribution that remains invariant under single-qubit Clifford rotations, Ref. [42] has established the following theorem to provide a rigorous guarantee for the low-weight Pauli propagation algorithm. Intuitively, the locally scrambling property means each layer is generic on a local scale and mixes the local basis.

Theorem 3 (Time complexity). Let U be sampled from an L-layered locally scrambling circuit ensemble on n qubits. For any error tolerance ϵ and failure probability δ, there exists a classical algorithm with run-time $Ln^{\mathcal{O}[\log(\epsilon^{-1}\delta^{-1})]}$ that outputs an estimate α satisfying

$$|\alpha - f_U(O)| \le \epsilon \|O\|_{\mathrm{Pauli},2}, \quad \text{(C5)}$$

with probability at least $1-\delta$. $\|O\|_{\mathrm{Pauli},2} = (2^{-n}\mathrm{tr}[O^\dagger O])^{1/2}$ is the normalized Hilbert-Schmidt norm.

This theorem establishes three key properties of the algorithm:

(1) The run-time is polynomial in system size n and circuit depth L for any small constant ϵ, δ.
(2) The accuracy parameters ϵ and δ appear only logarithmically in the exponent.
(3) The error scales naturally with the observable's normalized Hilbert-Schmidt norm.

4. Applicability and scope

This theoretical result encompasses various quantum circuit architectures. For the computational complexity bounds to hold, we require that for any Pauli operator P of weight k, its Heisenberg evolution $U_j^\dagger P U_j$ contains at most $n^{\mathcal{O}(k)}$ distinct Pauli terms and can be computed classically in time $n^{\mathcal{O}(k)}$. Intuitively, this condition ensures each circuit layer U_j remains reasonably shallow. The analysis applies to several important classes of quantum circuits:

(1) Brickwork circuits of SU(4) gates in 1D, 2D, and 3D [130–132].
(2) Circuits with universal single-qubit rotations followed by entangling Clifford gates [133].
(3) Quantum convolutional neural networks without feedforward [134].

Despite this broad applicability, the fundamental limitations of this approach warrant careful consideration. The low-weight Pauli propagation method remains inherently a classical heuristic. Its theoretical guarantees extend only to *most* quantum circuits where each layer is drawn from a locally scrambling distribution, rather than to arbitrary quantum circuits.

When presented with a specific quantum circuit of interest, determining whether this classical heuristic will yield accurate simulation results is a nontrivial challenge. This uncertainty points to a fundamental insight: As we demonstrate in the next section, the very task of determining whether a given quantum circuit exhibits an advantage over this classical heuristic constitutes a problem that itself admits a quantum advantage.

APPENDIX D: CLASSICAL HARDNESS IN PREDICTING QUANTUM ADVANTAGES

There are many classical heuristics for simulating quantum circuits and quantum algorithms. Some families of quantum circuits can be efficiently simulated using classical methods, nullifying any potential quantum advantage. However, the task of determining whether a given problem exhibits a quantum advantage can itself have a quantum advantage. Hence, even if a quantum circuit is classically simulatable using a specific classical heuristic, we might still need a quantum computer to determine or verify that the classical heuristic can indeed simulate the quantum circuit accurately.

Definition 6 (A classical heuristic for simulating quantum circuits). A classical algorithm $\mathcal{A}$ is called a classical heuristic for simulating a class of polynomial-size quantum circuits if:

(1) It runs in classical polynomial time.
(2) It correctly approximates the output probability $C(x)$ for a subclass of polynomial-size quantum circuits C, but it is known to fail, or is not proven to succeed, for some circuits in this class and generally the set of circuits for which it fails is unknown.

The algorithm is considered a heuristic because there is a gap between its practical performance on many instances and the lack of a guarantee for universal correctness within its target problem class.

A powerful example is the low-weight Pauli propagation heuristic, LOWWEIGHTPAULIPROP, discussed in Appendix C. It is known to be effective for most circuits in a measure-theoretic sense (a large subclass) but is also known to fail for others, fitting our definition perfectly. Given that a classical heuristic may succeed for some circuits and fail

for others, we need a precise, quantitative way to distinguish these cases. We formalize this by defining what it means for a quantum circuit to have a "computational advantage" over a classical heuristic. This occurs when the classical heuristic fails to provide a good approximation for a significant fraction of possible inputs.

Definition 7 (Quantum advantage over a classical heuristic). We say a quantum computer executing a quantum circuit C exhibits a computational advantage over a classical heuristic $\mathcal{A}$ if

$$|\mathcal{A}_C(x) - C(x)| \geq 1/3, \tag{D1}$$

for at least a $2/3$ fraction of the inputs $x \in \{0,1\}^n$. Conversely, we say it has no advantage if $|\mathcal{A}_C(x) - C(x)| < 1/3$ for at least a $2/3$ fraction of inputs.

The specific constants $1/3$ and $2/3$ are chosen by convention, following standard practice in computational complexity theory to create a mathematically convenient definition. The threshold of $1/3$ for the error ensures that a heuristic's prediction is significantly wrong, not just slightly inaccurate. The requirement that this occurs for a $2/3$ fraction of inputs establishes that the failure is typical for the circuit, not a rare exception. Together, these values create a large, constant-sized gap between the "advantage" case, where the heuristic is wrong for a clear majority of inputs, and the "no advantage" case, where it is right for a clear majority.

In the following, we prove that deciding if a given quantum circuit yields a computational advantage over a classical heuristic is itself a problem with a quantum advantage, under the widely believed conjecture that there exist some quantum computations that are classically hard.

For simplicity, we focus on LOWWEIGHTPAULIPROP with $k = 1$. The statement holds for $k > 1$ using a similar proof.

Theorem 4 (Quantum advantage in detecting quantum advantage). Consider the following computational problem for detecting quantum advantage in a given quantum circuit:

DETECTINGQUANTUMADVANTAGE

Input: A description of a quantum circuit C

Promise: A quantum computer executing C either exhibits a computational advantage over LOWWEIGHTPAULIPROP or it does not (as defined above).

Output: $\begin{cases} 1, & \text{if } C \text{ exhibits a quantum advantage} \\ 0, & \text{if } C \text{ does not exhibit a quantum advantage} \end{cases}$

This problem can be solved efficiently on a quantum computer (it is in BQP) but is intractable for any classical algorithm (it is not in BPP), assuming BPP $\neq$ BQP.

Proof. We divide the proof into two parts: one focused on showing that the problem is quantumly easy, and the second one focused on showing that the problem is classically hard.

The problem is in BQP *(Quantumly easy).* A quantum algorithm can solve DETECTINGQUANTUMADVANTAGE efficiently.

(1) Given the circuit description C, choose a set of $s = \mathcal{O}(1)$ random input strings $\{x_1, \ldots, x_s\}$.
(2) For each x_i for $i = 1, \ldots, s$:
   (a) Compute the heuristic's prediction $\mathcal{A}_C(x_i)$ using the classical LOWWEIGHTPAULIPROP algorithm as described in Appendix C.
   (b) Estimate the true probability $C(x_i)$ on a quantum computer. This is done by preparing the state $C(|x_i\rangle|0^m\rangle)$ and measuring the first qubit. Repeating this a polynomial number of times yields an estimate of $C(x_i)$ to within a small error $\delta \ll 1/3$ with high probability.
(3) For each x_i for $i = 1, \ldots, s$, check if $|\mathcal{A}_C(x_i) - C(x_i)| \geq 1/3$.
(4) If the fraction of inputs x_i satisfying this condition is greater than $1/2$, output 1. Otherwise, output 0. By a standard Chernoff bound with an appropriate sample size s, this sampling procedure correctly distinguishes the case where the true fraction is $\geq 2/3$ from the case where it is $\leq 1/3$ with high probability.

The entire procedure involves a polynomial number of classical steps and a polynomial number of runs on a quantum computer, so the problem is in BQP.

The problem is not in BPP *(Classically hard).* We show this by contradiction. Suppose there exists a polynomial-time classical algorithm $\mathcal{D}$ that can solve DETECTINGQUANTUMADVANTAGE. We will show how to use $\mathcal{D}$ to solve any problem in BQP classically, which would imply BPP = BQP, contradicting our assumption.

Before diving into technicalities, we briefly explain the main idea underlying our argument. Given any polynomial-size quantum circuit $C_?$ that solves a decision problem in BQP, we will construct a new polynomial-size quantum circuit C_{new} with input x that has the following properties. (1) If the answer to the decision problem is YES, then the output of the classical heuristic disagrees with the output of the quantum circuit C_{new} for most values of the input x. (2) If the answer to the decision problem is NO, then the output of the classical heuristic agrees with the output of the quantum circuit C_{new} for most values of the input x. Therefore, the polynomial-time classical algorithm $\mathcal{D}$, by determining whether the classical heuristic is successful or not, solves the decision problem.

Now we proceed with the formal proof. Let $C_?$ be a polynomial-size quantum circuit that solves an instance of an arbitrary problem in BQP. The circuit $C_?$ takes the all-zero state $|0^m\rangle$ as input, and when we measure the first qubit of $C_?|0^m\rangle$ in the Z basis, we are promised that the probability of observing 1 is either $\geq 2/3$ (YES instance) or $\leq 1/3$ (NO instance). We now construct a new circuit such that applying $\mathcal{D}$ to the new circuit will enable one to accurately distinguish between the YES and NO instances.

Step 1: Random circuit construction. Let $U = U_L \cdots U_1$ be an n-qubit random quantum circuit where each U_j is a layer of random two-qubit gates, and the total depth L can be chosen arbitrarily as long as L is polynomial in n and m. By choosing L to be at least linear in n with a sufficiently large constant, it is known that U forms an ϵ-approximate unitary 2-design with $\epsilon = \mathcal{O}(1/2^n)$ [21,135,136]. This means the distribution over the random unitary U nearly matches the Haar measure over all n-qubit unitaries $\mathrm{SU}(2^n)$ up to the second moment.

Step 2: Amplified circuit construction. To amplify the success probability, we construct $C_{?,\mathrm{ext}}$ that runs $\ell = \mathcal{O}(\log n)$ independent copies of $C_?$ and performs a coherent majority vote:

(1) Run ℓ copies of $C_?$ on separate ancilla registers.
(2) Apply a coherent majority vote circuit to the first qubits of each copy, storing the result in a designated ancilla qubit q_{maj}.

By standard concentration inequalities, such as the Chernoff bound, this amplified circuit satisfies the following:

(i) If $C_?$ is a YES instance: Pr[Measuring q_{maj} in Z basis gives1] $\geq 1 - 2^{-\Omega(\ell)} \geq 1 - 1/\mathrm{poly}(n)$.
(ii) If $C_?$ is a NO instance: Pr[Measuring q_{maj} in Z basis gives1] $\leq 2^{-\Omega(\ell)} \leq 1/\mathrm{poly}(n)$.

Note that $C_{?,\mathrm{ext}}$ acts on $m\ell + 1$ qubits.

Step 3: New circuit construction. We construct the circuit C_{new} that operates on the main n qubits plus all ancilla qubits:

$$C_{\mathrm{new}} = (U_L \cdots U_1 \otimes \mathbb{I}) \cdot \text{Controlled-}(U_1^\dagger \cdots U_L^\dagger) \cdot (\mathbb{I} \otimes C_{?,\mathrm{ext}}), \tag{D2}$$

where Controlled-$(U_1^\dagger \cdots U_L^\dagger)$ applies $U^\dagger = U_1^\dagger \cdots U_L^\dagger$ to the main n qubits conditioned on the majority vote qubit q_{maj} being in the state $|1\rangle$.

Let our observable be the Pauli operator $O = Z_1$ acting on the first qubit of the main system, and our initial state over the main and the ancilla systems be $\rho = |x\rangle\langle x| \otimes |0^{m\ell+1}\rangle\langle 0^{m\ell+1}|$ for an input n-bit string $x \in \{0,1\}^n$. Let us analyze the true expectation value:

$$\langle Z_1 \rangle_{\mathrm{true},x} = \mathrm{tr}[O \cdot C_{\mathrm{new}} \cdot \rho \cdot C_{\mathrm{new}}^\dagger]. \tag{D3}$$

We separate the analysis into two cases.

Case 1: $C_?$ is a YES instance. With probability $\geq 1 - 1/\mathrm{poly}(n)$, the majority vote qubit $q_{\mathrm{maj}} = 1$, so the effective unitary applied to the main qubits is

$$U_{\mathrm{eff}} = U_L \cdots U_1 \cdot (U_1^\dagger \cdots U_L^\dagger) = I. \tag{D4}$$

Therefore, for any input bit string $x \in \{0,1\}^n$, we have the following:

$$|\langle Z_1 \rangle_{\mathrm{true},x} - (-1)^{x_1}| \leq 1/\mathrm{poly}(n), \tag{D5}$$

where x_1 denotes the first bit of the input string x.

Case 2: $C_?$ is a NO instance. With probability $\geq 1 - 1/\mathrm{poly}(n)$, the majority vote qubit $q_{\mathrm{maj}} = 0$, so the effective unitary applied to the main qubits is just U. Therefore,

$$|\langle Z_1 \rangle_{\mathrm{true},x} - \langle x|U^\dagger Z_1 U|x\rangle| \leq 1/\mathrm{poly}(n). \tag{D6}$$

Because U forms a $1/\mathrm{poly}(n)$-approximate unitary 2-design, we have $|\langle x|U^\dagger Z_1 U|x\rangle| \leq 1/\mathrm{poly}(n)$ with high probability over the choice of U. Hence, for any input bit string $x \in \{0,1\}^n$, we have

$$|\langle Z_1 \rangle_{\mathrm{true},x} - 0| \leq 1/\mathrm{poly}(n) \tag{D7}$$

with high probability.

Now, we analyze how the LOWWEIGHTPAULIPROP heuristic $\mathcal{A}$ simulates the new circuit C_{new} layer by layer. The LOWWEIGHTPAULIPROP heuristic evolves Z_1 backward through the entire circuit. The following lemma focuses on propagation over the main n-qubit system and the first L layers in C_{new}.

Lemma 1 (Decay of Frobenius norm). Consider a main system of n qubits and an ancillary system of n' qubits. Let $U = U_L \cdots U_1$ be a random quantum circuit on the main n-qubit system, where each layer U_j consists of $n/2$ Haar-random two-qubit gates acting on disjoint pairs of qubits. Let O_{heur} be the operator computed by evolving Z_1 acting on $n + n'$ qubits backward through $U \otimes \mathbb{I}$ using the classical LOWWEIGHTPAULIPROP heuristic. Then we have

$$\mathbb{E}_U[\|O_{\mathrm{heur}}\|_F^2] = \left(\frac{2}{5}\right)^L \cdot 2^{n+n'}. \tag{D8}$$

Proof. We prove this by analyzing the evolution layer by layer. We track how the squared Frobenius norm changes as we apply each layer of the heuristic evolution. Let $O^{(L)} = Z_1$. For $j = L$ down to 1, each step of the heuristic evolution is given by

$$O^{(j-1)} = \mathcal{P}_1(U_j^\dagger O^{(j)} U_j), \tag{D9}$$

where $\mathcal{P}_1$ projects onto Pauli operators of weight at most 1. Consider a single layer U_j consisting of Haar-random two-qubit gates on disjoint pairs. Because the unitary U_j acts only on the first n qubits and all Pauli operators that act on two or more qubits are truncated and the evolution preserves the trace of zero, the $(n + n')$-qubit operator $O^{(j)}$ can be written as

$$O^{(j)} = \sum_{i=1}^{n} \sum_{P \in \{X,Y,Z\}} \alpha_{P_i}^{(j)} P_i, \tag{D10}$$

where P_i acts as Pauli operator P on the ith qubit in the main n-qubit systems and acts as identity $\mathbb{I}$ on all other qubits. The squared Frobenius norm is

$$\|O^{(j)}\|_F^2 = 2^{n+n'} \sum_{i=1}^{n} \sum_{P\in\{X,Y,Z\}} |\alpha_{P_i}^{(j)}|^2. \tag{D11}$$

For Z_1, we have $\alpha_{Z_1} = 1$, and all other coefficients are zero, so $\|Z_1\|_F^2 = 2^{n+n'}$.

Since each qubit participates in exactly one two-qubit gate per layer, and the gates are independent Haar-random unitaries, we can formally compute the expected squared Frobenius norm after one layer. After applying a two-qubit gate layer U_j and projecting with $\mathcal{P}_1$, we get

$$\begin{aligned} O^{(j-1)} &= \mathcal{P}_1(U_j^\dagger O^{(j)} U_j) \\ &= \sum_{i'=1}^{n} \sum_{P'\in\{X,Y,Z\}} \alpha_{P'_{i'}}^{(j-1)} P'_{i'}, \end{aligned} \tag{D12}$$

where the new coefficients are

$$\begin{aligned} \alpha_{P'_{i'}}^{(j-1)} &= \frac{1}{2^{n+n'}} \mathrm{Tr}(P'_{i'} \cdot U_j^\dagger O^{(j)} U_j) \\ &= \frac{1}{2^{n+n'}} \sum_{i=1}^{n} \sum_{P\in\{X,Y,Z\}} \alpha_{P_i}^{(j)} \mathrm{Tr}(P'_{i'} U_j^\dagger P_i U_j). \end{aligned} \tag{D13}$$

The squared Frobenius norm is

$$\|O^{(j-1)}\|_F^2 = 2^{n+n'} \sum_{i'=1}^{n} \sum_{P'\in\{X,Y,Z\}} |\alpha_{P'_{i'}}^{(j-1)}|^2 \tag{D14}$$

$$= 2^{n+n'} \sum_{i'=1}^{n} \sum_{P'\in\{X,Y,Z\}} \left| \frac{1}{2^{n+n'}} \sum_{i=1}^{n} \sum_{P\in\{X,Y,Z\}} \alpha_{P_i}^{(j)} \mathrm{Tr}(P'_{i'} U_j^\dagger P_i U_j) \right|^2 \tag{D15}$$

$$= \frac{1}{2^{n+n'}} \sum_{i'=1}^{n} \sum_{P'\in\{X,Y,Z\}} \left| \sum_{i=1}^{n} \sum_{P\in\{X,Y,Z\}} \alpha_{P_i}^{(j)} \mathrm{Tr}(P'_{i'} U_j^\dagger P_i U_j) \right|^2 \tag{D16}$$

Taking expectation over the random two-qubit gate layer U_j:

$$\mathbb{E}_{U_j}[\|O^{(j-1)}\|_F^2] \tag{D17}$$

$$= \frac{1}{2^{n+n'}} \sum_{i'=1}^{n} \sum_{P'\in\{X,Y,Z\}} \mathbb{E}_{U_j} \left[\left| \sum_{i=1}^{n} \sum_{P\in\{X,Y,Z\}} \alpha_{P_i}^{(j)} \mathrm{Tr}(P'_{i'} U_j^\dagger P_i U_j) \right|^2 \right] \tag{D18}$$

$$= \frac{1}{2^{n+n'}} \sum_{i'=1}^{n} \sum_{i,i''=1}^{n} \sum_{\substack{P\in\{X,Y,Z\} \\ P'\in\{X,Y,Z\} \\ P''\in\{X,Y,Z\}}} \alpha_{P_i}^{(j)} \overline{\alpha_{P''_{i''}}^{(j)}} \mathbb{E}_{U_j} \left[\mathrm{Tr}(P'_{i'} U_j^\dagger P_i U_j) \overline{\mathrm{Tr}(P'_{i'} U_j^\dagger P''_{i''} U_j)} \right] \tag{D19}$$

$$= \frac{1}{2^{n+n'}} \sum_{i'=1}^{n} \sum_{i,i''=1}^{n} \sum_{\substack{P\in\{X,Y,Z\} \\ P'\in\{X,Y,Z\} \\ P''\in\{X,Y,Z\}}} \alpha_{P_i}^{(j)} \overline{\alpha_{P''_{i''}}^{(j)}} \mathbb{E}_{U_j} \left[\mathrm{Tr}(P'_{i'} U_j^\dagger P_i U_j) \mathrm{Tr}(P'_{i'} U_j^\dagger P''_{i''} U_j) \right]. \tag{D20}$$

Now we need to compute the expectation $\mathbb{E}_{U_j}[\mathrm{Tr}(P'_{i'} U_j^\dagger P_i U_j)\mathrm{Tr}(P'_{i'} U_j^\dagger P''_{i''} U_j)]$. Since each qubit participates in exactly one two-qubit gate per circuit layer, we have two cases.

(i) Case 1: Qubits i and i'' are involved in different two-qubit gates. Let gate V_1 act on qubits $\{i, j\}$ and gate V_2 act on qubits $\{i'', j''\}$, where $\{i, j\} \cap \{i'', j''\} = \varnothing$. Since V_1 and V_2 are independent Haar-random gates, and P_i acts only on the space of V_1 while $P''_{i''}$ acts only on the space of V_2,

$$\mathbb{E}_{V_1,V_2}[\mathrm{Tr}(P'_{i'} V_1^\dagger P_i V_1)\mathrm{Tr}(P'_{i'} V_2^\dagger P''_{i''} V_2)] = 0. \tag{D21}$$

(ii) Case 2: Qubits i and i'' are involved in the same two-qubit gate that acts on the two qubits $\{i, j\}$. For a Haar-random two-qubit unitary gate V and single-qubit Pauli operators P_i, $P''_{i''}$, $P'_{i'}$ acting within the same two-qubit space $\{i, j\}$, we have

$$\mathbb{E}_V[\mathrm{Tr}(P'_{i'}V^\dagger P_i V)\mathrm{Tr}(P'_{i'}V^\dagger P''_{i''}V)] = \frac{4^{n+n'}}{15}\cdot\delta_{P_i,P''_{i''}}, \tag{D22}$$

from the standard Weingarten calculation for Haar-random unitaries [135,136]. Here, $\delta_{P_i,P''_{i''}}$ is 1 if the two n-qubit Pauli operators $P_i, P''_{i''}$ are the same and 0 otherwise. If i' is not acted on by the two-qubit gate V, then because $P'_{i'}$ is traceless, we have

$$\mathbb{E}_V[\mathrm{Tr}(P'_{i'}V^\dagger P_i V)\mathrm{Tr}(P'_{i'}V^\dagger P''_{i''}V)] = 0. \tag{D23}$$

Now, collecting all the surviving terms in the sum, we have

$$\mathbb{E}_{U_j}[\|O^{(j-1)}\|_F^2] = \frac{4^{n+n'}}{2^{n+n'}}\sum_{i=1}^{n}\sum_{P\in\{X,Y,Z\}} 6\cdot\frac{1}{15}\cdot|\alpha_{P_i}^{(j)}|^2, \tag{D24}$$

where 6 comes from the possible choices of $P'_{i'}$ such that qubit i' is involved in the unique two-qubit gate in U_j that acts on qubit i. Using $\|O^{(j)}\|_F^2 = 2^{n+n'}\sum_{i=1}^{n}\sum_{P\in\{X,Y,Z\}}|\alpha_{P_i}^{(j)}|^2$, we have

$$\mathbb{E}_{U_j}[\|O^{(j-1)}\|_F^2] = \frac{2}{5}\|O^{(j)}\|_F^2. \tag{D25}$$

After L layers, we have

$$\begin{aligned}\mathbb{E}_U[\|O_{\mathrm{heur}}\|_F^2] &= \mathbb{E}_U[\|O^{(0)}\|_F^2] = \left(\frac{2}{5}\right)^L\mathbb{E}_U[\|O^{(L)}\|_F^2]\\ &= \left(\frac{2}{5}\right)^L\|Z_1\|_F^2 = \left(\frac{2}{5}\right)^L\cdot 2^{n+n'}.\end{aligned} \tag{D26}$$

This concludes the proof. ∎

Let $\mathcal{P}$ be the projection onto the subspace of weight-1 Pauli operators. Let $T_V(O) = \mathcal{P}(V^\dagger O V)$ be the map for one step of the heuristic's backward evolution through a unitary V.

Lemma 2 (Nonincreasing Frobenius norm under the heuristic evolution). For any unitary V acting on q qubits and any operator O on the same space, we have

$$\|T_V(O)\|_F \le \|O\|_F. \tag{D27}$$

Proof. Since V is unitary, conjugation by V preserves the Frobenius norm: $\|V^\dagger O V\|_F = \|O\|_F$. The projection $\mathcal{P}$ onto a subspace can only decrease the Frobenius norm, so $\|\mathcal{P}(A)\|_F \le \|A\|_F$ for any operator A. Therefore, we have the following:

$$\|T_V(O)\|_F = \|\mathcal{P}(V^\dagger O V)\|_F \le \|V^\dagger O V\|_F = \|O\|_F. \tag{D28}$$

This completes the proof. ∎

Now we analyze the classical heuristic's simulation of circuit C_{new}. The circuit C_{new} acts on the combined Hilbert space of n main qubits and $m\ell+1$ ancilla qubits, for a total of $n+m\ell+1$ qubits. The observable of interest is $Z_1\otimes\mathbb{I}^{\otimes(m\ell+1)}$, which we denote as Z_1^{ext} for the extended system. The classical heuristic evolves Z_1^{ext} backward through the circuit $C_{\mathrm{new}} = (U_L\otimes\mathbb{I})\cdots(U_1\otimes\mathbb{I})\cdot\mathrm{Controlled}\text{-}(U^\dagger)\cdot(\mathbb{I}\otimes C_{?,\mathrm{ext}})$. We analyze this evolution in reverse order:

(i) *Backward evolution through the random circuit layers* $(U_L\otimes\mathbb{I})\cdots(U_1\otimes\mathbb{I})$: We start by evolving Z_1^{ext} backward through the L layers of random gates. Each layer applies $T_{U_j\otimes\mathbb{I}}$ to the current operator. Let $O_L = Z_1^{\mathrm{ext}} = Z_1\otimes\mathbb{I}^{\otimes(m\ell+1)}$. After backward evolution through $j=L$ down to 1:

$$O_{j-1} = T_{U_j\otimes\mathbb{I}}(O_j) = \mathcal{P}[(U_j^\dagger\otimes\mathbb{I})O_j(U_j\otimes\mathbb{I})]. \tag{D29}$$

Because the random circuit layers U_j act only on the main n qubits, we have

$$O_{\mathrm{postrandom}} \coloneqq O_0 = O_{\mathrm{heur}} \tag{D30}$$

for the $(n+n')$-qubit observable O_{heur} defined in Lemma 1 with $n' = m\ell+1$. We have

$$\begin{aligned}\mathbb{E}_U[\|O_{\mathrm{postrandom}}\|_F^2] &= \mathbb{E}_U[\|O_{\mathrm{heur}}\|_F^2]\\ &= \left(\frac{2}{5}\right)^L\cdot 2^{n+m\ell+1}.\end{aligned} \tag{D31}$$

This implies that over the choices of the L-layer random quantum circuit U, the average Frobenius norm decays exponentially in L.

(ii) *Backward evolution through the controlled operation:* The controlled operation is

$$\begin{aligned}\mathrm{Controlled}\text{-}(U^\dagger) = {}&|0\rangle\langle 0|_{q_{\mathrm{maj}}}\otimes\mathbb{I}_{\mathrm{main}}\otimes\mathbb{I}_{\mathrm{rest}}\\ &+|1\rangle\langle 1|_{q_{\mathrm{maj}}}\otimes U_{\mathrm{main}}^\dagger\otimes\mathbb{I}_{\mathrm{rest}}.\end{aligned} \tag{D32}$$

After backward evolution, we have the observable:

$$O_{\mathrm{postcontrolled}} = T_{\mathrm{Controlled}\text{-}(U^\dagger)}(O_{\mathrm{postrandom}}). \tag{D33}$$

By Lemma 2, we have $\|O_{\mathrm{postcontrolled}}\|_F \le \|O_{\mathrm{postrandom}}\|_F$.

(iii) *Backward evolution through* $(\mathbb{I}\otimes C_{?,\mathrm{ext}})$: Finally, we evolve backward through $(\mathbb{I}\otimes C_{?,\mathrm{ext}})$:

$$O_{\text{final}} = T_{\mathbb{I}\otimes C_{?,\text{ext}}}(O_{\text{postcontrolled}}). \tag{D34}$$

By Lemma 2, we have $\|O_{\text{final}}\|_F \le \|O_{\text{postcontrolled}}\|_F$. Combining these inequalities, we obtain the following:

$$\mathbb{E}_U[\|O_{\text{final}}\|_F^2] \le \mathbb{E}_U[\|O_{\text{postcontrolled}}\|_F^2] \le \mathbb{E}_U[\|O_{\text{postrandom}}\|_F^2] \le \left(\frac{2}{5}\right)^L \cdot 2^{n+m\ell+1}. \tag{D35}$$

By Markov's inequality, with probability at least $1-2^{-n}$ over the choice of U, we have

$$\|O_{\text{final}}\|_F^2 \le \left(\frac{2}{5}\right)^L \cdot 2^{2n+m\ell+1}. \tag{D36}$$

Therefore, for any input bit string $x \in \{0,1\}^n$, we have

$$|\langle Z_1\rangle_{\text{heur},x}| = |\text{Tr}[O_{\text{final}} \cdot (|x\rangle\langle x| \otimes |0^{m\ell+1}\rangle\langle 0^{m\ell+1}|)]| \le \|O_{\text{final}}\|_F \le \sqrt{\left(\frac{2}{5}\right)^L \cdot 2^{2n+m\ell+1}}. \tag{D37}$$

We can choose L arbitrarily as long as L is polynomial in $n+m\ell+1$. By choosing L to be greater than $6(n+m\ell+1)$, we have $(\frac{2}{5})^L \cdot 2^{2n+m\ell+1} \le (1/4)^{n+m\ell+1}$. Hence, irrespective of the YES and NO instances, the output of the classical heuristic for any input bit string $x \in \{0,1\}^n$ satisfies

$$|\langle Z_1\rangle_{\text{heur},x}| \le \left(\frac{1}{2}\right)^{n+m\ell+1} \tag{D38}$$

for all except an exponentially small fraction of random quantum circuits U.

Now we establish the separation between YES and NO instances. For YES instances, for any input bit string $x \in \{0,1\}^n$, we have

$$|\langle Z_1\rangle_{\text{true},x} - \langle Z_1\rangle_{\text{heur},x}| \ge |(-1)^{x_1}| - |\langle Z_1\rangle_{\text{heur},x}| - \frac{1}{\text{poly}(n)} \ge 1 - \frac{1}{\text{poly}(n)} > \frac{2}{3} \tag{D39}$$

for any sufficiently large n and all except exponentially small fraction of random quantum circuit U. For NO instances, for any input bit string $x \in \{0,1\}^n$, we have

$$|\langle Z_1\rangle_{\text{true},x} - \langle Z_1\rangle_{\text{heur},x}| \le |\langle Z_1\rangle_{\text{true},x}| + |\langle Z_1\rangle_{\text{heur},x}| \le \frac{1}{\text{poly}(n)} < \frac{1}{3} \tag{D40}$$

for any sufficiently large n and all except exponentially small fraction of random quantum circuit U. This establishes the following:

(i) If $C_?$ is a YES instance, then a quantum computer executing quantum circuit C_{new} exhibits a computational advantage over LOWWEIGHTPAULIPROP, so $\mathcal{D}(C_{\text{new}}) = 1$.

(ii) If $C_?$ is a NO instance, then a quantum computer executing quantum circuit C_{new} does not exhibit a computational advantage over LOWWEIGHTPAULIPROP, so $\mathcal{D}(C_{\text{new}}) = 0$.

Therefore, we can solve the original BQP problem for $C_?$ by the following classical algorithm:

(1) Construct the circuit C_{new} from $C_?$ as described above.

(2) Apply the assumed classical algorithm $\mathcal{D}$ to determine if a quantum computer executing quantum circuit C_{new} does not exhibit a computational advantage over LOWWEIGHTPAULIPROP.

(3) Output the result of $\mathcal{D}$.

Since this procedure works for any BQP problem and runs in polynomial time, we would have BPP = BQP, contradicting our assumption that BPP ≠ BQP. Therefore, no polynomial-time classical algorithm can solveDetectingQuantumAdvantage, completing the proof. ∎

[1] J. S. Bell, *On the Einstein Podolsky Rosen paradox*, Phys. Phys. Fiz. **1**, 195 (1964).

[2] S. Jaques and A. G. Rattew, *Qram: A survey and critique*, arXiv:2305.10310.

[3] A. M. Dalzell, A. Gilyén, C. T. Hann, S. McArdle, G. Salton, Q. T. Nguyen, A. Kubica, and F. G. Brandão, *A distillation-teleportation protocol for fault-tolerant QRAM*, arXiv:2505.20265.

[4] A. W. Harrow, A. Hassidim, and S. Lloyd, *Quantum algorithm for linear systems of equations*, Phys. Rev. Lett. **103**, 150502 (2009).

[5] E. Tang, *A quantum-inspired classical algorithm for recommendation systems*, in *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing* (Association for Computing Machinery (ACM), New York, NY, 2019), pp. 217–228.

[6] E. Tang, *Quantum principal component analysis only achieves an exponential speedup because of its state preparation assumptions*, Phys. Rev. Lett. **127**, 060503 (2021).

[7] A. S. Holevo, *Bounds for the quantity of information transmitted by a quantum communication channel*, Probl. Inf. Transm. **9**, 177 (1973).

[8] R. Raz, *Exponential separation of quantum and classical communication complexity*, in *Proceedings of the Thirty-First Annual ACM Symposium on Theory of Computing* (Association for Computing Machinery (ACM), New York, NY, 1999), pp. 358–367.

[9] A. Montanaro, *Quantum states cannot be transmitted efficiently classically*, Quantum **3**, 154 (2019).

[10] W. Kretschmer, S. Grewal, M. DeCross, J. A. Gerber, K. Gilmore, D. Gresh, N. Hunter-Jones, K. Mayer, B. Neyenhuis, D. Hayes *et al.*, *Demonstrating an*

unconditional separation between quantum and classical information resources, arXiv:2509.07255.

[11] D. Gilboa, H. Michaeli, D. Soudry, and J. McClean, *Exponential quantum communication advantage in distributed inference and learning*, Adv. Neural Inf. Process. Syst. **37**, 30425 (2024).

[12] Z. Zimborás, B. Koczor, Z. Holmes, E.-M. Borrelli, A. Gilyén, H.-Y. Huang, Z. Cai, A. Acín, L. Aolita, L. Banchi *et al.*, *Myths around quantum computation before full fault tolerance: What no-go theorems rule out and what they don't*, arXiv:2501.05694.

[13] S. Aaronson, A. M. Childs, E. Farhi, A. W. Harrow, and B. C. Sanders, *Future of quantum computing*, arXiv:2506.19232.

[14] R. King, *Quantum algorithms: A call to action*, https://quantumfrontiers.com/2025/04/20/quantum-algorithms-a-call-to-action/ (2025).

[15] O. Lanes, M. Beji, A. D. Corcoles, C. Dalyac, J. M. Gambetta, L. Henriet, A. Javadi-Abhari, A. Kandala, A. Mezzacapo, C. Porter *et al.*, *A framework for quantum advantage*, arXiv:2506.20658.

[16] I. Kerenidis and A. Prakash, *Quantum recommendation systems*, in *Proceedings of the 8th Innovations in Theoretical Computer Science Conference* (Schloss Dagstuhl—Leibniz-Zentrum für Informatik GmbH, Wadern/Saarbrücken, Germany, 2017), pp. 49:1–49:21.

[17] L. J. Stockmeyer, *The polynomial-time hierarchy*, Theor. Comput. Sci. **3**, 1 (1976).

[18] P. W. Shor, *Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer*, SIAM Rev. **41**, 303 (1999).

[19] O. Regev, *On lattices, learning with errors, random linear codes, and cryptography*, J. Assoc. Comput. Mach. **56**, 1 (2009).

[20] U. Mahadev, *Classical verification of quantum computations*, in *Proceedings of the 2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)* (IEEE Computer Society, Los Alamitos, California, 2018), pp. 259–267.

[21] T. Schuster, J. Haferkamp, and H.-Y. Huang, *Random unitaries in extremely low depth*, Science **389**, 92 (2025).

[22] UN Conference on Trade and Development, Digital economy report 2024: Shaping an environmentally sustainable and inclusive digital future (2024), business e-commerce sales reached $27 trillion in 2022 across 43 countries.

[23] Forrester Research, Global digital economy forecast, 2023 to 2028 (2024), digital economy projectedto reach $16.5 trillion by 2028.

[24] National Institute of Standards and Technology, Digital signature standard (dss) (2013), federal standard specifying RSA, DSA, and ECDSA for digital signatures.

[25] C.-F. Chen, H.-Y. Huang, J. Preskill, and L. Zhou, *Local minima in quantum systems*, in *Proceedings of the 56th Annual ACM Symposium on Theory of Computing* (Association for Computing Machinery (ACM), New York, NY, 2024), pp. 1845–1858.

[26] E. Farhi, J. Goldstone, and S. Gutmann, *A quantum approximate optimization algorithm*, arXiv:1411.4028.

[27] J. Basso, E. Farhi, K. Marwaha, B. Villalonga, and L. Zhou, *The quantum approximate optimization algorithm at high depth for* MAXCUT *on large-girth regular graphs and the Sherrington-Kirkpatrick model*, arXiv:2110.14206.

[28] J. Basso, D. Gamarnik, S. Mei, and L. Zhou, *Performance and limitations of the QAOA at constant levels on large sparse hypergraphs and spin glass models*, in *Proceedings of the 2022 IEEE 63rd Annual Symposium on Foundations of Computer Science (FOCS)* (IEEE Computer Society, Los Alamitos, California, 2022), pp. 335–343.

[29] E. Farhi, S. Gutmann, D. Ranard, and B. Villalonga, *Lower bounding the* MAXCUT *of high girth 3-regular graphs using the QAOA*, arXiv:2503.12789.

[30] S. P. Jordan, N. Shutty, M. Wootters, A. Zalcman, A. Schmidhuber, R. King, S. V. Isakov, T. Khattar, and R. Babbush, *Optimization by decoded quantum interferometry*, arXiv:2408.08292.

[31] A. Y. Kitaev, *Quantum measurements and the Abelian stabilizer problem*, arXiv:quant-ph/9511026.

[32] L. Lin and Y. Tong, *Near-optimal ground state preparation*, Quantum **4**, 372 (2020).

[33] D. Wu, R. Rossi, F. Vicentini, N. Astrakhantsev, F. Becca, X. Cao, J. Carrasquilla, F. Ferrari, A. Georges, M. Hibat-Allah *et al.*, *Variational benchmarks for quantum many-body problems*, Science **386**, 296 (2024).

[34] S. R. White, *Density matrix formulation for quantum renormalization groups*, Phys. Rev. Lett. **69**, 2863 (1992).

[35] D. Aharonov, X. Gao, Z. Landau, Y. Liu, and U. Vazirani, *A polynomial-time classical algorithm for noisy random circuit sampling*, in *Proceedings of the 55th Annual ACM Symposium on Theory of Computing* (Association for Computing Machinery (ACM), New York, NY, 2023), pp. 945–957.

[36] Y. Shao, F. Wei, S. Cheng, and Z. Liu, *Simulating quantum mean values in noisy variational quantum algorithms: A polynomial-scale approach*, arXiv:2306.05804.

[37] N. A. Nemkov, E. O. Kiktenko, and A. K. Fedorov, *Fourier expansion in variational quantum algorithms*, Phys. Rev. A **108**, 032406 (2023).

[38] T. Begušić, J. Gray, and G. K.-L. Chan, *Fast and converged classical simulations of evidence for the utility of quantum computing before fault tolerance*, Sci. Adv. **10**, eadk4321 (2024).

[39] T. Begušić, K. Hejazi, and G. K. Chan, *Simulating quantum circuit expectation values by Clifford perturbation theory*, J. Chem. Phys. **162** (2025).

[40] E. Fontana, M. S. Rudolph, R. Duncan, I. Rungger, and C. Cîrstoiu, *Classical simulations of noisy variational quantum circuits*, npj Quantum Inf. **11**, 84 (2025).

[41] M. S. Rudolph, E. Fontana, Z. Holmes, and L. Cincio, *Classical surrogate simulation of quantum systems with LOWESA*, arXiv:2308.09109.

[42] A. Angrisani, A. Schmidhuber, M. S. Rudolph, M. Cerezo, Z. Holmes, and H.-Y. Huang, *Classically estimating observables of noiseless quantum circuits*, arXiv:2409.01706.

[43] N. Dowling, P. Kos, and X. Turkeshi, *Magic resources of the Heisenberg picture*, Phys. Rev. Lett. **135**, 050401 (2025).

[44] F. Arute *et al.*, *Quantum supremacy using a programmable superconducting processor*, Nature (London) **574**, 505 (2019).
[45] Q. Zhu, S. Cao, F. Chen, M.-C. Chen, X. Chen, T.-H. Chung, H. Deng, Y. Du, D. Fan, M. Gong *et al.*, *Quantum computational advantage via 60-qubit 24-cycle random circuit sampling*, Sci. Bull. **67**, 240 (2022).
[46] A. Morvan, B. Villalonga, X. Mi, S. Mandra, A. Bengtsson, P. Klimov, Z. Chen, S. Hong, C. Erickson, I. Drozdov *et al.*, *Phase transitions in random circuit sampling*, Nature (London) **634**, 328 (2024).
[47] D. A. Abanin, R. Acharya, L. Aghababaie-Beni, G. Aigeldinger, A. Ajoy, R. Alcaraz, I. Aleiner, T. I. Andersen, M. Ansmann, F. Arute *et al.*, *Constructive interference at the edge of quantum ergodic dynamics*, arXiv:2506.10191.
[48] D. Gao, D. Fan, C. Zha, J. Bei, G. Cai, J. Cai, S. Cao, F. Chen, J. Chen, K. Chen *et al.*, *Establishing a new benchmark in quantum computational advantage with 105-qubit Zuchongzhi 3.0 processor*, Phys. Rev. Lett. **134**, 090601 (2025).
[49] D. Boneh, *The decision Diffie-Hellman problem*, in *Proceedings of the International Algorithmic Number Theory Symposium* (Springer, New York, 1998), pp. 48–63.
[50] V. Shoup, *Lower bounds for discrete logarithms and related problems*, in *Proceedings of the International Conference on the Theory and Applications of Cryptographic Techniques* (Springer, New York, 1997), pp. 256–266.
[51] A. Schmidhuber, R. O'Donnell, R. Kothari, and R. Babbush, *Quartic quantum speedups for planted inference*, Phys. Rev. X **15**, 021077 (2025).
[52] D. Gottesman, *Fault-tolerant quantum computation with constant overhead*, Quantum Inf. Comput. **14**, 1338 (2014).
[53] Q. T. Nguyen and C. A. Pattison, *Quantum fault tolerance with constant-space and logarithmic-time overheads*, in *Proceedings of the 57th Annual ACM Symposium on Theory of Computing* (Association for Computing Machinery (ACM), New York, NY, 2025), pp. 730–737.
[54] S. Zhou, M. Zhang, J. Preskill, and L. Jiang, *Achieving the Heisenberg limit in quantum metrology using quantum error correction*, Nat. Commun. **9**, 78 (2018).
[55] M. A. Fischler and R. C. Bolles, *Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography*, Commun. ACM **24**, 381 (1981).
[56] R. Tibshirani, *Regression shrinkage and selection via the lasso*, J. R. Stat. Soc. Ser. B **58**, 267 (1996).
[57] N. Natarajan, I. S. Dhillon, P. K. Ravikumar, and A. Tewari, *Learning with noisy labels*, Adv. Neural Inf. Process. Syst. **26**, 1196 (2013).
[58] H.-Y. Huang, R. Kueng, and J. Preskill, *Information-theoretic bounds on quantum advantage in machine learning*, Phys. Rev. Lett. **126**, 190505 (2021).
[59] S. Chen, J. Cotler, H.-Y. Huang, and J. Li, *Exponential separations between learning with and without quantum memory*, in *Proceedings of the 2021 IEEE 62nd Annual Symposium on Foundations of Computer Science (FOCS)* (IEEE Computer Society, Los Alamitos, California, 2022), pp. 574–585.
[60] H.-Y. Huang, M. Broughton, J. Cotler, S. Chen, J. Li, M. Mohseni, H. Neven, R. Babbush, R. Kueng, J. Preskill, and J. R. McClean, *Quantum advantage in learning from experiments*, Science **376**, 1182 (2022).
[61] H.-Y. Huang, S. T. Flammia, and J. Preskill, *Foundations for learning from noisy quantum experiments*, arXiv:2204.13691.
[62] S. Chen, J. Cotler, H.-Y. Huang, and J. Li, *The complexity of NISQ*, Nat. Commun. **14**, 6001 (2023).
[63] Z.-H. Liu, R. Brunel, E. E. Østergaard, O. Cordero, S. Chen, Y. Wong, J. A. Nielsen, A. B. Bregnsbo, S. Zhou, H.-Y. Huang *et al.*, *Quantum learning advantage on a scalable photonic platform*, arXiv:2502.07770.
[64] S. Aaronson and A. Arkhipov, *The computational complexity of linear optics*, *Proceedings of the Forty-Third Annual ACM Symposium on Theory of Computing* (Association for Computing Machinery (ACM), New York, NY, 2011), 333.
[65] R. Orús, *A practical introduction to tensor networks: Matrix product states and projected entangled pair states*, Ann. Phys. (Amsterdam) **349**, 117 (2014).
[66] C. L. Degen, F. Reinhard, and P. Cappellaro, *Quantum sensing*, Rev. Mod. Phys. **89**, 035002 (2017).
[67] S. Pirandola, U. L. Andersen, L. Banchi, M. Berta, D. Bunandar, R. Colbeck, D. Englund, T. Gehring, C. Lupo, C. Ottaviani *et al.*, *Advances in quantum cryptography*, Adv. Opt. Photonics **12**, 1012 (2020).
[68] A. M. Dalzell, S. McArdle, M. Berta, P. Bienias, C.-F. Chen, A. Gilyén, C. T. Hann, M. J. Kastoryano, E. T. Khabiboulline, A. Kubica *et al.*, *Quantum algorithms: A survey of applications and end-to-end complexities*, arXiv:2310.03011.
[69] S. P. Jordan, *Quantum algorithm zoo*, https://quantumalgorithmzoo.org.
[70] E. Bernstein and U. Vazirani, *Quantum complexity theory*, in *Proceedings of the Twenty-Fifth Annual ACM Symposium on Theory of Computing* (Association for Computing Machinery (ACM), New York, NY, 1993), pp. 11–20.
[71] D. R. Simon, *On the power of quantum computation*, SIAM J. Comput. **26**, 1474 (1997).
[72] L. K. Grover, *A fast quantum mechanical algorithm for database search*, in *Proceedings of The Twenty-Eighth Annual ACM Symposium on Theory of Computing* (Association for Computing Machinery (ACM), New York, NY, 1996), pp. 212–219.
[73] E. Farhi and S. Gutmann, *Quantum computation and decision trees*, Phys. Rev. A **58**, 915 (1998).
[74] A. M. Childs, R. Cleve, E. Deotto, E. Farhi, S. Gutmann, and D. A. Spielman, *Exponential algorithmic speedup by a quantum walk*, in *Proceedings of the Thirty-Fifth Annual ACM Symposium on Theory of Computing* (Association for Computing Machinery (ACM), New York, NY, 2003), pp. 59–68.
[75] E. Farhi, J. Goldstone, and S. Gutmann, *A quantum algorithm for the Hamiltonian NAND tree*, arXiv:quant-ph/0702144.
[76] R. P. Feynman, *Simulating physics with computers*, Int. J. Theor. Phys. **21**, 467 (1982).

[77] S. Lloyd, *Universal quantum simulators*, Science **273**, 1073 (1996).
[78] G. Carleo and M. Troyer, *Solving the quantum many-body problem with artificial neural networks*, Science **355**, 602 (2017).
[79] S. Lee, J. Lee, H. Zhai, Y. Tong, A. M. Dalzell, A. Kumar, P. Helms, J. Gray, Z.-H. Cui, W. Liu *et al.*, *Evaluating the evidence for exponential quantum advantage in ground-state quantum chemistry*, Nat. Commun. **14**, 1952 (2023).
[80] H.-Y. Huang, M. Broughton, M. Mohseni, R. Babbush, S. Boixo, H. Neven, and J. R. McClean, *Power of data in quantum machine learning*, Nat. Commun. **12**, 2631 (2021).
[81] H.-Y. Huang, R. Kueng, G. Torlai, V. V. Albert, and J. Preskill, *Provably efficient machine learning for quantum many-body problems*, Science **377**, eabk3333 (2022).
[82] L. Lewis, H.-Y. Huang, V. T. Tran, S. Lehner, R. Kueng, and J. Preskill, *Improved machine learning algorithm for predicting ground state properties*, Nat. Commun. **15**, 895 (2024).
[83] V. Giovannetti, S. Lloyd, and L. Maccone, *Quantum-enhanced measurements: Beating the standard quantum limit*, Science **306**, 1330 (2004).
[84] V. Giovannetti, S. Lloyd, and L. Maccone, *Quantum metrology*, Phys. Rev. Lett. **96**, 010401 (2006).
[85] R. Demkowicz-Dobrzański, J. Kołodyński, and M. Guţă, *The elusive Heisenberg limit in quantum-enhanced metrology*, Nat. Commun. **3**, 1063 (2012).
[86] T. L. S. Collaboration, *A gravitational wave observatory operating beyond the quantum shot-noise limit*, Nat. Phys. **7**, 962 (2011).
[87] B. P. Abbott *et al.*, *Observation of gravitational waves from a binary black hole merger*, Phys. Rev. Lett. **116**, 061102 (2016).
[88] R. Schnabel, *Squeezed states of light and their applications in laser interferometers*, Phys. Rep. **684**, 1 (2017).
[89] D. Ganapathy, W. Jia, M. Nakano, V. Xu, N. Aritomi, T. Cullen, N. Kijbunchoo, S. Dwyer, A. Mullavey, L. McCuller *et al.*, *Broadband quantum enhancement of the LIGO detectors with frequency-dependent squeezing*, Phys. Rev. X **13**, 041021 (2023).
[90] J. M. Taylor, P. Cappellaro, L. Childress, L. Jiang, D. Budker, P. Hemmer, A. Yacoby, R. Walsworth, and M. Lukin, *High-sensitivity diamond magnetometer with nanoscale resolution*, Nat. Phys. **4**, 810 (2008).
[91] I. Lovchinsky, A. Sushkov, E. Urbach, N. P. de Leon, S. Choi, K. De Greve, R. Evans, R. Gertner, E. Bersin, C. Müller *et al.*, *Nuclear magnetic resonance detection and spectroscopy of single proteins using quantum logic*, Science **351**, 836 (2016).
[92] P. Bhattacharyya, W. Chen, X. Huang, S. Chatterjee, B. Huang, B. Kobrin, Y. Lyu, T. J. Smart, M. Block, E. Wang *et al.*, *Imaging the Meissner effect in hydride superconductors using quantum sensors*, Nature (London) **627**, 73 (2024).
[93] M. Wang, Y. Wang, Z. Liu, G. Xu, B. Yang, P. Yu, H. Sun, X. Ye, J. Zhou, A. F. Goncharov *et al.*, *Imaging magnetic transition of magnetite to megabar pressures using quantum sensors in diamond anvil cell*, Nat. Commun. **15**, 8843 (2024).
[94] A. D. Ludlow, M. M. Boyd, J. Ye, E. Peik, and P. O. Schmidt, *Optical atomic clocks*, Rev. Mod. Phys. **87**, 637 (2015).
[95] T. Bothwell, C. J. Kennedy, A. Aeppli, D. Kedar, J. M. Robinson, E. Oelker, A. Staron, and J. Ye, *Resolving the gravitational redshift across a millimetre-scale atomic sample*, Nature (London) **602**, 420 (2022).
[96] T. L. Nicholson, S. Campbell, R. Hutson, G. E. Marti, B. Bloom, R. L. McNally, W. Zhang, M. Barrett, M. S. Safronova, G. Strouse *et al.*, *Systematic evaluation of an atomic clock at* 2×10^{-18} *total uncertainty*, Nat. Commun. **6**, 6896 (2015).
[97] N. Schine, A. W. Young, W. J. Eckner, M. J. Martin, and A. M. Kaufman, *Long-lived Bell states in an array of optical clock qubits*, Nat. Phys. **18**, 1067 (2022).
[98] D. Aharonov, J. Cotler, and X.-L. Qi, *Quantum algorithmic measurement*, Nat. Commun. **13**, 1 (2022).
[99] S. Chen, C. Oh, S. Zhou, H.-Y. Huang, and L. Jiang, *Tight bounds on Pauli channel learning without entanglement*, Phys. Rev. Lett. **132**, 180805 (2024).
[100] C. Oh, S. Chen, Y. Wong, S. Zhou, H.-Y. Huang, Jens A. Nielsen, Z.-H. Liu, J. S. Neergaard-Nielsen, U. L. Andersen, L. Jiang, and J. Preskill, *Entanglement-enabled advantage for learning a bosonic random displacement channel*, Phys. Rev. Lett. **133**, 230604 (2024).
[101] R. R. Allen, F. Machado, I. L. Chuang, H.-Y. Huang, and S. Choi, *Quantum computing enhanced sensing*, arXiv:2501.07625.
[102] C. H. Bennett and G. Brassard, *Quantum cryptography: Public key distribution and coin tossing*, Theor. Comput. Sci. **560**, 7 (2014).
[103] R. Colbeck and R. Renner, *Free randomness can be amplified*, Nat. Phys. **8**, 450 (2012).
[104] M. Liu, R. Shaydulin, P. Niroula, M. DeCross, S.-H. Hung, W. Y. Kon, E. Cervero-Martín, K. Chakraborty, O. Amer, S. Aaronson *et al.*, *Certified randomness using a trapped-ion quantum processor*, Nature (London), **640**, 1 (2025).
[105] S. Wiesner, *Conjugate coding*, ACM SIGACT News **15**, 78 (1983).
[106] A. Broadbent and R. Islam, *Quantum encryption with certified deletion*, in *Proceedings of the Theory of Cryptography Conference*, Lecture Notes in Computer Science Vol. 12552 (Springer, New York, 2020), pp. 92–122.
[107] S. Aaronson, *Quantum copy-protection and quantum money*, in *Proceedings of the 24th Annual IEEE Conference on Computational Complexity* (IEEE Computer Society, Washington, DC, 2009), pp. 229–242.
[108] S. Goldwasser, Y. T. Kalai, and G. N. Rothblum, *One-time programs*, in *Proceedings of the Annual International Cryptology Conference* (Springer, New York, 2008), pp. 39–56.
[109] D. Ding and L. Jiang, *Coordinating decisions via quantum telepathy*, arXiv:2407.21723.
[110] J. Kallaugher, O. Parekh, and N. Voronova, *Exponential quantum space advantage for approximating maximum directed cut in the streaming model*, in *Proceedings of the 56th Annual ACM Symposium on Theory of Computing*

(Association for Computing Machinery (ACM), New York, NY, 2024), pp. 1805–1815.

[111] H. Ren and R. Santhanam, *A relativization perspective on meta-complexity*, in *Proceedings of the 39th International Symposium on Theoretical Aspects of Computer Science (STACS 2022)* (Schloss Dagstuhl—Leibniz-Zentrum für Informatik GmbH, Wadern/Saarbrücken, Germany, 2022), pp. 54–1.

[112] *Simons Institute for the Theory of Computing, Meta-complexity program, Online Program with Videos and Materials (2023)* (accessed July 3, 2025). https://simons.berkeley.edu/programs/meta23.

[113] H. Bernien, S. Schwartz, A. Keesling, H. Levine, A. Omran, H. Pichler, S. Choi, A. S. Zibrov, M. Endres, M. Greiner *et al.*, *Probing many-body dynamics on a 51-atom quantum simulator*, Nature (London) **551**, 579 (2017).

[114] C. J. Turner, A. A. Michailidis, D. A. Abanin, M. Serbyn, and Z. Papić, *Weak ergodicity breaking from quantum many-body scars*, Nat. Phys. **14**, 745 (2018).

[115] M. Ippoliti and W. W. Ho, *Solvable model of deep thermalization with distinct design times*, Quantum **6**, 886 (2022).

[116] J. S. Cotler, D. K. Mark, H.-Y. Huang, F. Hernández, J. Choi, A. L. Shaw, M. Endres, and S. Choi, *Emergent quantum state designs from individual many-body wave functions*, PRX Quantum **4**, 010311 (2023).

[117] A. Aspect, J. Dalibard, and G. Roger, *Experimental test of Bell's inequalities using time-varying analyzers*, Phys. Rev. Lett. **49**, 1804 (1982).

[118] S. J. Freedman and J. F. Clauser, *Experimental test of local hidden-variable theories*, Phys. Rev. Lett. **28**, 938 (1972).

[119] J. Preskill, *Lecture notes for physics 229: Quantum Information and Computation, Lecture notes, California Institute of Technology* (1998).

[120] L. Devroye, *Nonuniform random variate generation*, Handb. Oper. Res. Manag. Sci. **13**, 83 (2006).

[121] M. D. Vose, *A linear algorithm for generating random numbers with a given distribution*, IEEE Transactions on Software Engineering **17**, 972 (1991).

[122] H. Buhrman, R. Cleve, S. Massar, and R. de Wolf, *Nonlocality and communication complexity*, Rev. Mod. Phys. **82**, 665 (2010).

[123] M. W. Mahoney, *Randomized algorithms for matrices and data*, Found. Trends Mach. Learn. **3**, 123 (2011).

[124] H. Buhrman, R. Cleve, J. Watrous, and R. de Wolf, *Quantum fingerprinting*, Phys. Rev. Lett. **87**, 167902 (2001).

[125] W.-T. Kuo, A. A. Akhtar, D. P. Arovas, and Y.-Z. You, *Markovian entanglement dynamics under locally scrambled quantum evolution*, Phys. Rev. B **101**, 224202 (2020).

[126] H.-Y. Hu, S. Choi, and Y.-Z. You, *Classical shadow tomography with locally scrambled quantum dynamics*, Phys. Rev. Res. **5**, 023027 (2023).

[127] M. C. Caro, H.-Y. Huang, N. Ezzell, J. Gibbs, A. T. Sornborger, L. Cincio, P. J. Coles, and Z. Holmes, *Out-of-distribution generalization for learning quantum dynamics*, Nat. Commun. **14**, 3751 (2023).

[128] J. Gibbs, Z. Holmes, M. C. Caro, N. Ezzell, H.-Y. Huang, L. Cincio, A. T. Sornborger, and P. J. Coles, *Dynamical simulation via quantum machine learning with provable generalization*, Phys. Rev. Res. **6**, 013241 (2024).

[129] H.-Y. Huang, S. Chen, and J. Preskill, *Learning to predict arbitrary quantum processes*, PRX Quantum **4**, 040337 (2023).

[130] H.-K. Zhang, S. Liu, and S.-X. Zhang, *Absence of barren plateaus in finite local-depth circuits with long-range entanglement*, Phys. Rev. Lett. **132**, 150603 (2024).

[131] J. Napp, *Quantifying the barren plateau phenomenon for a model of unstructured variational ansätze*, arXiv:2203.06174.

[132] P. Braccia, P. Bermejo, L. Cincio, and M. Cerezo, *Computing exact moments of local random quantum circuits via tensor networks*, Quantum Mach. Intell. **6**, 54 (2024).

[133] A. Letcher, S. Woerner, and C. Zoufal, *Tight and efficient gradient bounds for parameterized quantum circuits*, Quantum **8**, 1484 (2024).

[134] A. Pesah, M. Cerezo, S. Wang, T. Volkoff, A. T. Sornborger, and P. J. Coles, *Absence of barren plateaus in quantum convolutional neural networks*, Phys. Rev. X **11**, 041011 (2021).

[135] A. W. Harrow and R. A. Low, *Random quantum circuits are approximate 2-designs*, Commun. Math. Phys. **291**, 257 (2009).

[136] F. G. Brandao, A. W. Harrow, and M. Horodecki, *Local random quantum circuits are approximate polynomial-designs*, Commun. Math. Phys. **346**, 397 (2016).