KerrisdaleCapital

April 2025

D-Wave Quantum Inc (QBTS)

Past-Gen Computing

We are short shares of D-Wave, a $2 billion quantum annealing company whose stock has surged more than 600% since investors began chasing anything remotely associated with next-gen computing last year. While the broader market looks to gate-model quantum systems as the industry's future, D-Wave continues to promote its fundamentally different annealing architecture. Shares currently trade at over ***57x*** consensus 2026E revenue – a ridiculous multiple for a company that has never generated more than *$9 million* in annual recurring revenue, has no clear path to profitability, and sees stagnating customer growth as its approach is increasingly recognized as a commercial dead end. D-Wave is riding quantum hype, but with a core technology that cannot stay afloat.

D-Wave's business is built around quantum annealing, a niche offshoot of quantum computing developed decades ago and largely abandoned by the industry. Despite marketing its systems as optimized for complex real-world optimization problems, quantum annealing has failed to demonstrate a clear performance edge over classical solvers. Encoding problems for D-Wave's hardware requires cumbersome and lossy reformulations that inflate problem size and introduce instability, while the system's limited qubit connectivity forces users to spread logical variables across fragile chains of physical qubits. Strip away carefully worded press releases and as a former D-Wave engineer admitted to us, "**there is no proof that any optimization problem is solved faster"** using D-Wave's quantum systems. Multiple academic studies have shown little to no scaling advantage over classical methods making it no surprise that a wide range of **D-Wave customers we interviewed in key verticals like logistics, manufacturing, and pharmaceuticals reported seeing zero benefit from the technology**.

To compensate for the shortcomings of its quantum technology, D-Wave leans heavily on so-called "hybrid" solutions which combine its annealer with classical hardware and algorithms to tackle industrial-scale problems. But in a glaring red flag, the company refuses to disclose to customers the relative contribution of each. Why? **Because according to former insiders who developed and deployed these solutions, "hybrid" in practice means "almost entirely classical."** We believe the quantum component is minimal – often cosmetic – and its added value, beyond marketing puff pieces, was debated even internally.

Meanwhile, D-Wave's gate-model pivot appears both reactionary and stalled. After years of dismissing gate-based systems as impractical, the company abruptly reversed course in 2021 amid a surge of investor interest in gate-model competitors like IBM and Google. Over three years later, D-Wave has released no detailed architecture papers, no fidelity data, and no performance benchmarks. D-Wave's absence of a gate-model roadmap leaves it clinging to a fading technology while the rest of the industry passes it by.

Even D-Wave's most recent highly publicized research lacks commercial relevance. Last month, the company declared quantum supremacy on a "useful, real-world" problem, an assertion dismissed by physicists we spoke with as misleading. The benchmark was a toy problem engineered to align with D-Wave's hardware constraints, bearing little resemblance to real-world magnetic materials. D-Wave is not a leading quantum company – it's a struggling provider of uncompetitive optimization solutions. As investors wake up to this reality, the stock's rally will unwind, and its valuation will collapse under the weight of physics, finance, and fact.

Table of Contents

Executive Summary

Quantum annealing is a commercial dead end. D-Wave's long-running struggle to commercialize its niche offshoot of quantum computing reflects fundamental technical limitations. Serious adoption hurdles tied to physical constraints, algorithmic limitations, and persistent underperformance against classical optimization software will only continue to produce anemic financial results and erode shareholder value.

Encoding and scaling issues severely undercut D-Wave's practical utility. Real-world optimization problems involve multiple overlapping constraints. Advanced classical solvers can manage these with relative ease, while D-Wave's annealer requires cumbersome and lossy reformulation into an obscure mathematical format. This process increases problem size, adds auxiliary variables, and consumes scarce quantum resources. On top of this, D-Wave's limited qubit connectivity forces logical variables to be spread across chains of physical qubits, further reducing capacity and introducing instability. Studies have consistently shown a lack of meaningful scaling advantage for D-Wave's systems. While marketed as a revolutionary tool, D-Wave's hardware behaves more like an overengineered accelerator with no proven edge in solving large or complex industrial problems.

D-Wave hides how hybrid solutions are almost entirely driven by classical algorithms. Given the severe limitations of its quantum systems, D-wave's commercial strategy revolves around selling "hybrid" quantum-classical solutions for industrial optimization problems like vehicle routing and workforce scheduling. Based on multiple interviews with D-Wave customers and former employees, the company deliberately obfuscates how its hybrid solutions function in practice. Based on comments from former D-Wave insiders, the hybrid approach is driven "almost entirely" by advanced classical algorithms, with the inclusion of quantum processing representing little more than a marketing gimmick.

D-Wave's gate-model progress appears stalled. For most of its history, D-Wave dismissed gate-model quantum computing as impractical. That conviction abruptly reversed in October 2021, when the company announced plans to build a gate-based system – widely seen as a defensive reaction to investor capital and technical momentum shifting toward gate-model leaders like IBM, Quantinuum, and Google. Over the past three years, however, D-Wave has released no performance metrics, published no peer-reviewed data, and provided no detailed product roadmap. Without a credible gate-model program, D-Wave's already niche relevance in the quantum landscape risks further narrowing.

Quantum supremacy claim is commercially overstated. D-wave recently declared it had achieved quantum supremacy on a useful, real-world problem. According to three quantum physics experts we interviewed, however, the benchmark was in fact a "toy problem" carefully designed to match D-Wave's strengths and lacks industrial relevance. While the result may be of interest for condensed matter physicists engaged in academic research, it falls well short of real-world quantum advantage.

Share price is divorced from fundamentals. Despite pulling back from recent highs, the stock remains up over 600% since gate-model quantum computing hype reignited last year. The excitement over quantum's AI-like potential centers on the promise of gate-model systems – *not* D-Wave's entirely different annealing architecture. Yet shares trade at a ludicrous **57x 2026E inflated consensus revenue (and 152x our estimate),** even as customer growth stalls and years of cash burn loom ahead. The rally reflects misplaced quantum hype, not a business grounded in commercial reality.

Company Overview

Capitalization and Financial Summary

$ Millions, Balances as of Dec. 31, 2024	
QBTS share price	$7.53
Fully-diluted shares [(1)]	307
Market capitalization	**$2,311**
Operating lease liabilities	8
Loans payable	30
Cash and investments [(2)]	324
Net debt (cash)	(286)
Total enterprise value	**$2,025**

Financial Summary ($ mm)

Fiscal year end Dec 31,	**2022A**	**2023A**	**2024A**	**2025E**	**2026E**	**2027E**
QCaaS	$6	$5	$7	$8	$9	$10
Professional services	1	4	2	3	4	5
System sales and other	0	0	0	11	1	1
Total revenue	$7	$9	$9	$22	$13	$15
Growth y/y		*22%*	*1%*	*150%*	*-39%*	*10%*
Cash operating costs						
Cost of revenue	$2	$3	$2	$5	$3	$4
Margin	*33%*	*30%*	*27%*	*24%*	*25%*	*24%*
Research & development	27	28	27	30	33	36
Margin	*383%*	*322%*	*304%*	*138%*	*246%*	*241%*
SG&A	26	32	36	40	43	43
Margin	*356%*	*368%*	*402%*	*181%*	*324%*	*293%*
Adj. EBITDA loss (reported)	(48)	(54)	(56)	(54)	(66)	(67)
Stock-based compensation	(9)	(22)	(16)	(26)	(29)	(33)
Adj. EBITDA (incl. SBC)	(58)	(76)	(72)	(79)	(95)	(100)
Capex	(0)	(1)	(2)	(2)	(4)	(4)
Free cash flow (Adj. EBITDA less capex)	(49)	(55)	(58)	(56)	(70)	(71)
Key Metrics						
Enterprise Value / Sales			229x	92x	152x	137x

Source: Kerrisdale forecast. Historical financials per QBTS SEC filings.

1. *287.8m common shares outstanding and 3.5m exchangeable shares as of March 12, 2025. Includes 15.8m net dilutive effect of in-the-money outstanding stock options, warrants and rights using the treasury stock method.*
2. *December 31, 2024, balance pro forma $146.2m net proceeds from ATM equity offering completed in 1Q25.*

Co-founded in 1999 by Dr. Geordie Rose (former CTO) and a team of physicists, D-Wave pioneered the commercialization of quantum computers through its specialized quantum annealing approach – a branch of quantum computing that is starkly different from the gate-model architectures pursued by all other leading competitors in the field such as Google and IBM. Unlike competitors, who are focused on building universal, programmable quantum computers capable of running diverse algorithms, D-Wave's machines have been designed from the outset to solve a specific class of optimization problems using quantum effects to identify low energy solutions. The company went public in 2022 via a SPAC merger and is headquartered in Burnaby, British Columbia, with additional offices in Palo Alto, California.

D-Wave occupies a contentious position in quantum computing circles. One former senior executive bluntly characterized the company as the "redheaded stepchild" of the industry. This reputation stems from longstanding controversies, beginning with early skepticism that D-Wave's annealing approach constituted genuine quantum computation. The company further alienated members of the scientific community through aggressive marketing claims – including premature declarations of speedups over classical computing in carefully selected benchmarks that later failed to hold up under scrutiny. Though Rose departed in 2014, multiple industry sources confirmed to us that D-Wave continues to grapple with the legacy of his hype-heavy marketing approach.

Quantum annealing is a method of quantum computing designed to solve optimization problems – situations where the goal is to find the best solution among many possibilities. D-Wave's quantum annealers approach this by modeling a problem as a network of binary variables and searching for the configuration with the lowest "energy," a concept borrowed from physics. Instead of trying every option one by one like classical computers might, quantum annealing uses quantum superposition and tunneling to explore many solutions at once and "tunnel" through barriers that would trap classical methods in suboptimal answers. While powerful in theory, quantum annealing is limited in scope: it works best for a narrow class of problems and must be carefully tuned to each instance. For a more detailed overview of quantum annealing, see Appendix I.

Hybrid Solutions

D-Wave's quantum annealing technology faces core limitations that restrict its ability to solve large-scale real-world problems (discussed further later in this report). To work around these constraints – and revive struggling commercial prospects – D-Wave introduced hybrid quantum-classical solvers in 2019. These systems combine advanced classical optimization algorithms with quantum sub-processing, enabling users to tackle problems that pure quantum annealing cannot handle effectively. When D-Wave advertises the ability to solve optimization problems with up to 2 million variables and constraints, it is exclusively through these hybrid solvers. These solutions now form the backbone of D-Wave's go-to-market strategy of helping customers "confront real-world problems of growing complexity."

Based on multiple conversations with D-Wave customers and former employees, the company does not disclose precisely how much of its hybrid solution is powered by classical computing versus quantum hardware – making the system, by design, a literal black box. This lack of transparency is not coincidental. In our view, it reflects an uncomfortable truth: while customers are charged a premium for the supposed advantage of quantum computing in solving complex problems, D-Wave's hybrid solvers rely overwhelmingly on advanced *classical* techniques, such as parallel tempering. The quantum component, according to multiple sources, contributes no differentiated value. One former D-Wave engineer we spoke with put it bluntly: apart from marketing, the quantum component's added benefit is "highly questionable" and a source of internal skepticism and debate.

(No) Advantage Systems

D-Wave released its first commercial quantum annealer in 2011 with 128 qubits. The company's current Advantage system now features over 5,000 qubits and supports hybrid quantum-classical solvers. Prototypes of its second generation Advantage2 system with around 1,200 qubits are currently available with commercial rollout of a full-production 4,400 qubit version targeted before year-end. Advantage systems are available via the company's Leap quantum cloud service and access can be purchased directly from D-Wave or through Amazon Web Services Marketplace. D-Wave removed access to its devices from Amazon Bracket in 2022, following a falling out between the two companies. According to a knowledgeable source, this stemmed directly from D-Wave's frustration with certain Amazon technical staff who believed D-Wave's systems provided no practical advantage over classical methods in solving customer problems.

Though impressive sounding, D-Wave's oft-cited qubit counts can be misleading when comparing D-Wave's technology to gate-model quantum processors like those from IBM or Google. The two are fundamentally different architectures designed for distinct purposes. Gate-model quantum computers employ flexible, programmable qubits capable of universal quantum computation, where each additional qubit exponentially expands the computational space.

In contrast, D-Wave's annealing qubits serve as fixed components in a physical optimization system, functioning more like specialized sensors than general-purpose computing elements. In D-Wave's architecture, qubits are not fully connected, meaning logical problems must be "embedded" into the hardware using extra qubits as bridges. According to experts consulted, this embedding process ties up large numbers of qubits – sometimes >90% of them – just to represent a problem, not solve it. When D-Wave promotes a 5,000-qubit+ Advantage system, only a fraction of those qubits typically contributes to meaningful computation after accounting for the substantial overhead.

In addition to classical optimization use cases, D-Wave promotes its annealing technology as having potential in quantum simulation, particularly for materials science applications. In March 2025, the company claimed to be the first in the world to demonstrate quantum supremacy on a "useful, real-world problem." However, as we detail later in the report, expert perspectives from the scientific community raise serious doubts about the true significance of this claim.

Lack of Gate-Model Progress

For most of D-Wave's corporate history, its leadership dismissed gate-model quantum computing as impractical in the near term, citing crippling challenges like qubit coherence, fidelity, and error correction. Founder Geordie Rose once derided gate-based systems as a "rotten idea," insisting quantum annealing was the only viable path to commercial applications. Yet in October 2021, D-Wave pivoted, announcing plans to develop its own gate-model quantum computer. The reversal came amid mounting pressure: customers demanded algorithms beyond annealing, rivals like IBM and Google had made great technical strides in materials engineering and error-correction, and investors had poured billions into gate-based approaches.

Three-and-half years after announcing the program, the company has published no gate fidelity measurements and offered no updates on error correction strategies – all while continuing to heavily promote its annealing systems. An industry source with direct connections to D-Wave told us he believed D-Wave's system had been "de-prioritized." While competitors like IBM and Quantinuum routinely demonstrate progress through peer-reviewed publications, detailed hardware roadmaps, and verifiable performance metrics, D-Wave maintains a conspicuous silence on substantive details. The absence of these standard technical disclosures suggests D-Wave's gate-model efforts remain in early exploration phases at best.

As the broader industry coalesces around gate-model architectures as the future of general-purpose quantum computing, D-Wave remains tethered to a modality with limited scalability, constrained applicability, and declining investor interest.

Financial Profile

D-Wave generates revenue from three main sources:

1. Cloud-based quantum computing as a service (QCaaS) (79% of 2024 revenue). Access to quantum annealing computers with more than 5,000 qubits and quantum-classical hybrid solvers that can solve problems with up to two million variables.
2. Professional services (19% of 2024 revenue). Consulting services to customers, assistance in identifying and implementing quantum computing applications.
3. On-premise system sales (0% of 2024 revenue). On-premise quantum annealers to

research, academic, and government customers.

Despite beginning commercial operations in 2011, D-Wave's financial results remain startlingly weak. Thirteen years of commercialization have never yielded more than $9m in annual revenue (excluding periodic system sales) and a consistent record of significant operating losses (accumulated deficit of $627m as of YE2024). IBM spends more on quantum research in a single quarter than D-Wave generates in revenue in a year.[1] In 2024, D-Wave reported an EBITDA loss of $(56)m, slightly below company guidance which called for an improvement over 2023's EBITDA loss of $(54)m. This dismal performance persists despite the company burning through over half a billion dollars in venture funding, government grants, SPAC proceeds, and hundreds of millions in dilutive equity issuance (p.134).

D-Wave's financial performance exemplifies the worst hype of the SPAC bubble era, with 2024 revenue coming in *-88%* below original projections. This shortfall has occurred despite marketing claims touting "industry-leading" quantum solutions and an expanding pipeline of real-world use cases. D-Wave's actual results reveal stagnation. 2024 total revenue was flat year-over-year, with a modest $1.9 million increase in QCaaS revenue coming entirely from price increases as the number of commercial customers *declined* year-over-year. We believe D-Wave's performance reflects a fundamentally poor customer value proposition, struggling to drive adoption even as the broader gate-model quantum market advances.

The increase in consensus revenue estimates for 2025E revenue to ~$22m from $8.8m reported last year is driven by the sale of an Advantage system to Forschungzentrum Jülich, a German high performance computing center. On the company's 4Q24 earnings call, QCaaS bookings in 4Q24 were described as only "incrementally" better than 3Q24, which were only $2.3m – which itself was a decline of 22% year-over-year from 3Q23. According to a supercomputing expert we interviewed, there are a handful of superconducting clusters globally that purchase experimental systems on multi-year cycles. These computing centers then typically make the device available on a grant-based system to academic researchers. D-Wave's system sale does not reflect a particularly novel strategy – neutral atom quantum computing company Pasqal delivered a computer to the same Julich supercomputing center as well in December 2024, and similar quantum/supercomputing integrations have taken place in Spain, Poland (twice), and Australia. Though any form of revenue is undoubtedly welcome, the deal represents a one-time sale of specialized hardware to an extremely niche market – not an inflection in commercial adoption.

Facing significant losses for the foreseeable future and failing to maintain NYSE minimum bid price compliance through much of 2H24, D-Wave was forced to massively dilute its share count by ~77% over the last 12 months at a weighted average price of just $2.89 ~62% below the current stock price.

---

[1] Total IBM R&D spending in 2024 was $7.5bn and quantum computing is a key pillar of IBM research (alongside AI and semiconductors). Even assuming 1% was spent on quantum research would equate to multiples more than D-Wave has ever generated in annual revenue.

Significant Dilutive Share Issuance: 2Q24-1Q25

	Common Shares Issued	Net Proceeds ($ Millions)	Implied Share Price
2Q24			
Lincoln Park Purchase Agreement	14.9	$20.3	$1.36
Needham / B. Riley / Roth ($100m ATM)	8.3	$9.1	$1.10
3Q24			
Lincoln Park Purchase Agreement	10.6	$11.9	$1.12
Needham / B. Riley / Roth ($100m ATM)	11.1	$11.6	$1.05
4Q24			
Lincoln Park Purchase Agreement	9.3	$12.1	$1.30
Needham / B. Riley / Roth ($100m ATM)	30.5	$76.5	$2.51
New Needham / B. Riley / Roth ($75m ATM)	15.6	$72.9	$4.68
1Q25			
Another Needham / B. Riley / Roth ATM ($150m)	24.6	$146.2	$5.94
Total issuance	124.9	$360.6	$2.89
Dilution since 2023 / % below current share price	*77%*		*-62%*

Source: Kerrisdale analysis, QBTS SEC form 10-Q and form 10-K filings.

Despite pulling back from recent highs, D-Wave's share price remains up over 600% since last October – a staggering rally that has occurred alongside largely unchanged consensus revenue estimates, as shown below.

Recent Share Price Performance vs. Consensus Revenue Estimates

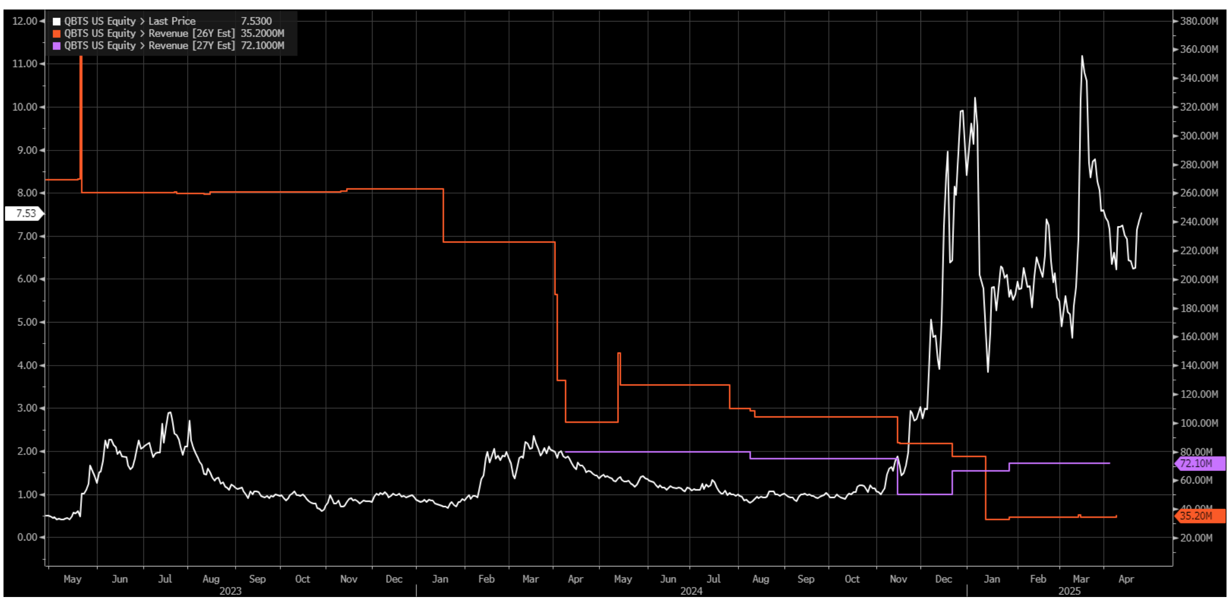

Source: Bloomberg

Shares currently trade at 57x consensus expectation of ~$35m in 2026E – an estimate that optimistically assumes significant year-over-year growth next year despite the non-recurring Jülich system sale in 2025E. This valuation appears especially tenuous given that the current wave of quantum enthusiasm is driven by the long-term potential of gate-model quantum computing – not D-Wave's fundamentally different annealing approach. Moreover, D-Wave has provided no meaningful updates on its own gate-model development, which, according to industry participants we spoke with, significantly trails competitors.

Quantum Annealing: Fundamental Limitations

> *"**There is no proof that any optimization problem is solved faster using [D-Wave's] system** compared to classical solvers."*
>
> — Former D-Wave senior data scientist
>
> *"**Adding quantum slows down the solution** because you have to send it to the [quantum computing] device…And there's a real nuance that you have to crush all your variables down in some semi-principled way so that you can fit a representation of the problem on the quantum computer. And then you have to reconstitute that solution into the larger landscape. And so, in all of that **there's never been any great evidence that D-Wave has provided internally, or externally to be honest, that that was doing any good.**"*
>
> — Former D-Wave engineer, technical consultant

For over a decade, D-Wave has positioned itself as a leader in practical quantum computing, promising to solve real-world optimization problems faster than classical machines. Yet today, the company still generates a pittance in revenue for a firm of its longevity and purported technological significance. The reason is not for lack of quality customer engagement or awareness – D-Wave's investor presentations contain a number of household names in the Fortune 500 (several of which we interviewed). The reason, we believe, is quantum annealing faces profound adoption challenges rooted in fundamental physical limitations, algorithmic constraints, and unfavorable comparisons with classical optimization methods.

Classical optimization solvers like Gurobi and CPLEX represent the culmination of decades of algorithmic refinement, software engineering, and hardware exploitation – making them difficult for quantum annealers to outperform. These solvers employ advanced algorithms like parallel tempering and techniques such as branch-and-bound, branch-and-cut, and cutting planes, all of which allow them to efficiently navigate massive, high-dimensional search spaces. They are also highly optimized for modern computing architectures, leveraging multi-core CPUs, GPUs, and distributed systems to accelerate performance. Crucially, classical solvers retain the native problem structure – something lost when converting problems into the mathematical form required for D-Wave's quantum annealer. As a result, Gurobi and CPLEX can routinely solve large-scale industrial problems to optimality or near-optimality with high reliability and speed, setting a performance bar that numerous academic studies show D-Wave's annealers, even when hybridized with classical methods, have yet to meaningfully clear.

Lost in Translation: Mapping the Real World to QUBO

When companies try to solve complex logistics tasks like vehicle routing for a fleet of delivery trucks, they must account for numerous, overlapping real-world constraints: delivery time windows, vehicle capacity, driver shifts, road restrictions, customer priorities, and more. Classical solvers like Gurobi can encode these constraints directly. They allow companies to describe and translate problems just as they exist in the real world, using flexible mathematical rules that the solver can understand natively.

For D-Wave to solve a logistics or manufacturing optimization problem, however, the entire problem must be translated into a mathematical form called Quadratic Unconstrained Binary Optimization (QUBO) – a requirement that introduces complexity, scalability, and performance challenges. First, all constraints in a real-world problem must be encoded as penalties within the QUBO objective function. For example, if a truck is not supposed to exceed its capacity, one must encode a penalty that increases the cost if too many deliveries are assigned to that truck. Similarly, time windows for deliveries cannot be encoded directly – they must be penalized if violated. A study by Lucas (2014) provides a detailed catalog of QUBO mappings for various constraint types and found that even simple constraint satisfaction problems grow rapidly in size and complexity when encoded for annealing. More recently, studies by Quintero and Zuluaga (2021), Quintero et al. (2021), and Zielinski et al. (2023) collectively highlight that reformulating constrained problems into QUBO form often leads to increased problem size and complexity, which then requires many extra variables (known as slack or auxiliary variables), which in turn consumes limited quantum resources.

Beyond the technical challenge of QUBO encoding lies a major practical drawback: most customers have no experience framing their business problems in abstract mathematical terms. This means D-Wave and its partners often need to devote significant consulting resources to help clients translate real-world use cases into a QUBO-compatible format. According to our research, it can take a dedicated team of D-Wave engineers several months of close collaboration just to reframe a problem in a way the system can attempt to solve. D-Wave has developed software libraries to reduce this consulting burden, but, in practice, customer problems are often highly specialized. The company's strategy of targeting an exceptionally diverse range of use cases – from aircraft loading for Airbus to workforce scheduling for a Canadian grocery chain to vehicle deployment for North Wales police – limits the ability to reuse code or apply learnings across engagements. While the company markets its offering as "Quantum Computing as a Service" (QCaaS), the reality is far from anything resembling a scalable, SaaS-like model. Each customer engagement is highly bespoke and resource intensive. In contrast, classical solvers can often tackle the same problem with far less effort and in most cases deliver competitive, if not superior performance.

Physical Limitations and Sparse Connectivity

In addition to encoding challenges, D-Wave's quantum annealing systems face several core physical limitations that reduce their usefulness in practice. For example, D-Wave's Pegasus and upcoming Zephyr architectures provide each qubit only 15-20 direct connections. This is relatively sparse connectivity compared to the fully connected interactions many real-world interactions require. As a result, logical variables must be mapped onto long chains of physical qubits to simulate the necessary connectivity.

This embedding process introduces significant overhead. First, it increases the number of physical qubits needed for each logical variable, reducing the effective problem size that can be solved. Second, it introduces the risk of "breaks," where long chains meant to represent a single logical variable fall out of sync, reducing solution accuracy. A study by Könz et al. (2021) showed this overhead to be severe. The authors showed that embedding problems into the sparse topology of D-Wave's quantum annealers incurs a "quadratic space overhead," (meaning the size of hardware needs grows much faster than the size of the problem) leading to a significant increase in the time to solution.

No Proven Scaling Advantage

While D-Wave promotes its annealers as accelerators for optimization problems, academic benchmarks consistently show that they fail to deliver meaningful speedups as problem sizes increase. Multiple studies have demonstrated that D-Wave's hybrid solvers provide, at best, constant-factor improvements in runtime – nowhere near the exponential speedups often associated with quantum computing, and far below what we were advised would be needed to galvanize customer adoption. Mandrà et al. (2017) found that even for problems known as deceptive cluster loops – a benchmark we were advised was specifically chosen to stack the deck in favor of annealing to help tease out quantum advantage – there was still zero scaling improvement found.

Similarly, Rønnow et al. (2014) evaluated D-Wave devices on benchmark problems and found performance curves nearly identical to classical algorithms, and specifically for larger problem sizes D-Wave did not surpass that of classical methods, indicating no scale advantage. More recently, Quinton et al. (2024) benchmarked D-Wave's hybrid quantum-classical solver against classical optimization tools like CPLEX, Gurobi, and IPOPT across various problem types. The study found as the number of variables and constraints in a problem increases, the computational time for D-Wave's hybrid solver grew significantly. Furthermore, while D-Wave's hybrid solver showed competitiveness in certain binary quadratic problems, its advantage diminished in more complex scenarios. In these cases, classical solvers like Gurobi outperformed the hybrid solver in both computational time and quality.

Quantum Price Tag for Classical Solution

Based on our research, D-Wave's quantum services are dramatically more expensive than traditional classical solvers – often by an order of magnitude or more. Access to D-Wave's hybrid solvers is priced via a cloud-based, usage-driven model that can run into thousands of dollars per hour depending on problem complexity and volume, with enterprise contracts typically ranging from the mid-six to low-seven figures annually. The standard price just to run a proof of concept? A staggering $350,000.

In contrast, best-in-class classical solvers like Gurobi, CPLEX, and Google OR-Tools are either free for academic or limited use, or available under commercial licenses at a fraction of the cost. These tools run efficiently on standard CPUs or GPUs and require no specialized hardware. For most real-world optimization problems, the classical approach is not only faster and more accurate – it's vastly cheaper. In our view, organizations that use D-Wave's quantum offering for classical optimization are paying a steep premium for the quantum "brand," not differentiated quantum performance. And as we outline in the next section, in many cases we believe customers who use D-Wave's hybrid solvers do not even realize they are paying quantum prices for what are, in essence, classically driven solutions.

Quantum Shell Game

> *"If you use the hybrid solver, the hybrid solver has to compose QUBO problems. And by design they design the algorithm in a way the construction of QUBO problems is efficient. Otherwise, you would pay a price for using the quantum solver.* ***However, you are still composing a classical problem and because there is no quantum advantage, you can substitute quantum with classical and still get better results. So, the entire hybrid framework could be beneficial, but it's not because of quantum****, it's because of its construct, and* ***the quantum part can be replaced with classical.***
>
> — Former D-Wave data scientist

Former D-Wave Engineer*: "One of the key issues is* ***[D-Wave's] hybrid solver doesn't really tell you how much quantum is in the hybrid****…*

If I were using the LEAP [hybrid] solver, I would not notice any difference between classical solver A, classical solver B, and that Leap hybrid because it's all classical, right? It's gonna return sort of a classical answer *somewhere in the middle. It's gonna do some quantum, but there's a fundamental issue there….****it's a strange situation that I never was very entirely comfortable with****…*

Kerrisdale: Which technology is actually solving the optimization problem?

Former D-Wave Engineer*: If I had to bet,* ***I would say almost entirely the classical system.***

Kerrisdale: What is the added benefit of including quantum in a hybrid solution?

Former D-Wave Engineer*: I would say,* ***the added benefit is we call it 'quantum' and this is a benefit, to some, in the marketing department.***

> *"D-Wave is selling* ***bullsh-t****"*
>
> — Executive at large multi-national technology and logistics company, former professor of quantum physics and quantum algorithms

When D-Wave launched its 5,000-qubit Advantage™ system in 2020 – billed as the "first quantum computer built for business" – the company explicitly acknowledged that solving real-world problems would require hybrid approaches. Despite ongoing upgrades to its quantum hardware, it remains hybrid solvers that enable D-Wave's ability to address large-scale optimization problems effectively.

Concerningly, based on multiple interviews with D-Wave customers and former employees, the company does not disclose how its hybrid solutions function in practice. Those same interviews consistently indicated that its hybrid approach is more accurately described as advanced *classical* algorithm driven – with quantum annealing only performing certain subroutines that provide no differentiating contribution to the outcome. D-Wave's hybrid solvers may indeed be solving optimization problems efficiently, but the involvement of quantum is, in essence, a gimmick.

We believe the conferences D-Wave hosts which showcase glowing customer testimonials highlight not only users who fail to grasp the true role quantum plays in D-Wave's hybrid solutions but also reflect a critical selection bias. According to a former senior D-Wave C-suite executive, most featured client engagements (e.g., Vinci Group, Mastercard) come through innovation teams tasked with exploring emerging technologies – not with identifying the most effective tools for solving their problems. These case studies almost never detail which classical methods were used as a baseline and typically omit whether D-Wave's hybrid approach was benchmarked against industry-leading classical solvers like Gurobi or CPLEX. For instance, Pattison Food Group, featured at Qubits2024, reported an 80% increase in scheduling efficiency over its previous method – which, according to our research, was simply an Excel spreadsheet. Nearly any modern optimization tool could have delivered that improvement. In our view, these testimonials are less about demonstrating quantum advantage and more about putting on a show of corporate innovation theater. As a former D-Wave engineer put it: "[Customers] are helping quantum computing, more than quantum computing is helping them."

Real-World Customer Perspectives

Rather than rely on D-Wave's carefully selected testimonials, we encourage investors to consider the candid feedback we gathered from D-Wave customers across manufacturing (Airbus), chemicals (BASF), pharmaceuticals (GlaxoSmithKline), logistics IT (Unisys), and from the world's largest logistics provider (Amazon). Taken together, these perspectives paint a clearer picture of why despite a steady stream of impressive-sounding press releases, D-Wave's customer growth has stagnated.

> *"I know that with D-Wave, with the Advanced system and so on, we did a lot on the, I would say "announcement side,"* ***but on the actual improvements it was for me a bit harder to see what we're actually improving because even if it did improve, I think it's still kind of a niche solution****…I am kind of in the belief, in the long term a lot of movements are not based on hype and short term research, but rather on the fundamentals of the businesses, on the actual measurables that they create.* ***And for me, I cannot see how [D-Wave] really are able to be a very sustainable business."***
>
> — Research and technology manager, **Airbus SE**

The Airbus manager confirmed that it began using D-Wave as part of experimental research geared at exploring quantum computing broadly. Airbus still tinkers with D-Wave for theoretical exercises but does not actively use D-Wave for real-world optimization problems. The customer also interestingly noted that contracts like the one with D-Wave are likely among the first to be cancelled within the research department in a recessionary environment.

> *"The project that we developed using the quantum solution was something small…During the execution and development phase [beyond initial proof-of-concept] you have a lot of different issues to solve, and when we were going through that,* ***we tried using some hybrid solution from [D-Wave].*** *And I can tell you* ***the results were 'okay,' not something amazing…in my opinion the combination of classical and quantum – it's not strong enough in comparison with other competitors."***

— Project Manager, **BASF SE**

The BASF representative later clarified how the relationship with D-Wave has stalled after deployment of one small steam generation system for a chemical plant and no new projects with D-Wave are planned. Like every other customer we spoke with, the BASF engineer said it was a "black box" as to the respective roles quantum versus classical hardware played in D-Wave's hybrid solutions.

*"**They're very coy about [the mix of classical and quantum]** because when we actually ran this, what we mapped to a true QUBO problem, so you could run it entirely quantum on their system **using the fully quantum annealer, it worked, but the noise was just horrendous. The error bars were always all over the place** and it would never really converge well. When we used their hybrid solver, all of those problems went away. What I could never get an answer from them though, and we asked this was, **'what exactly is that hybrid solver?"** I asked as much as, **"How do I know it's not just a GPU sitting in a big fancy box with D-Wave written on the side?" And they never gave us a suitable answer.***

Any of these optimization problems, I can't believe anybody in the field right now is truly doing something on a quantum computer that they can't just do using regular computing methods."

— Former senior executive in Scientific Computing, **GlaxoSmithKline**

*"The best way we've described it is it's great for school projects, but when you get to real business problems, it's not mature enough to be able to really find its niche…**I don't feel in the next five years annealing will find an effective business problem that it can solve in a way that it is viable financially**, 'cause it's expensive to run, it's expensive to access and it is actually slower…for us, **hybrid solving didn't actually give us any benefits**…fundamentally, it worked, **but it did not provide any benefit over classical computing."***

— Logistics IT Architect, **Unisys**

*"They were going with this approach called quantum annealing. It just did not seem to be bearing fruit at any point over the last like 20 years. So eventually they switched over to what they call **hybrid solvers, which are a mix of classical computing and some amount of quantum, but they won't tell you what exactly, or like what's going on under the hood.** So, the general understanding in the quantum community was, I mean, [hybrid] is probably just classical computing, **if they have like a quantum element to it, it's probably not doing anything substantial**."*

— Former Product Manager, **Amazon Bracket**

Quantum Supremacy Spin

> *“**The claim of quantum supremacy is a clear no**...I’m willing to bet a thousand bucks that in the next three months some science team is gonna simulate that diamond lattice and say, there you go…**Papers have refuted D-Wave’s claim and D-Wave is just holding onto a small technicality**…Ok, so they did a different lattice shape? Big deal. Let’s not forget the obvious fact that **this problem is very contrived,** and the real-world lattice would be a pyrochlore lattice and that they cannot do.”*
>
> — PhD Physics, expert in quantum physics and materials science

> *“**It’s a toy problem**…**it’s certainly <u>not a real-world problem</u>** in the sense of having practical applications that will lead to new discoveries and material science or something along those lines.”*
>
> — PhD Physics, 30 years of experience in quantum computing research, has worked closely with D-Wave quantum annealers since 2011.

> *“It’s almost like what they built is to solve this problem…They say it’s material science, but my thought would be, **in reality you want to make this more complex**, and as soon as you make this more complex – more variables or you make it a mixture or something – **as soon as you introduce that, their system won’t be able to solve it."***
>
> — PhD Computational Chemistry, 20+ years as scientist and leader of research teams engaged in quantum chemistry, lead developer of world-renowned molecular dynamics simulation software.

This report has so far focused on the limitations of D-Wave’s quantum annealer and hybrid solvers in tackling classical optimization problems – the foundation of the company’s commercial strategy. But quantum annealing is also being explored for quantum simulation, with potential applications in materials science and drug discovery. Recently, D-Wave announced what it described as a landmark achievement in this area.

In a press release highlighting its paper *“Beyond-Classical Computation in Quantum Simulation,”* the company claimed to have “unequivocally validate[d] the world’s first and only demonstration of quantum computational supremacy on a *useful, real-world problem”* [emphasis added]. D-Wave stated its annealing quantum computer solved a magnetic materials simulation exponentially faster than classical supercomputers. However, three quantum physics experts we interviewed – one of whom has worked with D-Wave annealers since their commercial debut – argued that the claims are highly misleading. The simulated problem, involving simplified spin glass lattices, was custom-tailored to play to D-Wave’s architectural strengths. In contrast, real magnetic materials – such as lithium holmium fluoride – exhibit fully connected, three-dimensional lattice structures that D-Wave’s hardware cannot efficiently represent or solve. The company’s “supremacy” claim, therefore, hinges on an artificial benchmark and experts noted that with further optimization, classical tensor network methods could likely match D-Wave’s results.

While the underlying research reflects meaningful progress in academic quantum simulation, framing it as a real-world breakthrough is a stretch. The lattice configurations used in the paper were carefully engineered to sidestep D-Wave's known limitations. Even the most complex instance cited – the biclique lattice – is an abstract problem with no direct industrial relevance. As all three experts quoted above emphasized, the result is academically impressive but falls well short of D-Wave's bold framing around "useful" quantum advantage.

This episode fits a familiar pattern in quantum computing: incremental progress repackaged as revolutionary breakthrough – most recently evident in Microsoft's overhyped Majorana qubit announcement. We believe that modeling complex real-world magnetic systems remains beyond D-Wave's capabilities. D-Wave's latest simulation results are not a paradigm-shifting advance as the company suggests. The problem was contrived, the competitive context questionable, and the real-world utility effectively nonexistent.

Conclusion

> *"They're trying to turn the business around. But to me, the technology is old and it's going to be outperformed by true gate-based quantum...they [D-wave] were great at starting off quantum, but they're in the ending arc of their business, they're not in the expanding arc."*
>
> — Industry consultant, 20 years in quantum computing business development for IBM, Quantinuum, and several others in the industry.

Quantum annealing is not the future of quantum computing. D-Wave's unwillingness to disclose the true role of quantum hardware in its hybrid solutions underscores a deeper reality: annealing isn't just niche – it's commercially unviable. Optimization customers have been sold a costly marketing gimmick rather than a groundbreaking tool, even as the broader industry advances toward fault-tolerant, gate-model systems capable of tackling problems annealing was never designed to handle. Without a successful pivot – unlikely as its gate-model effort stagnates while deep-pocketed competitors race ahead – D-Wave risks becoming a footnote in the very revolution it claims to lead.

Appendix I

Quantum annealing is a quantum heuristic designed to solve combinatorial optimization problems by leveraging the principles of quantum mechanics. In this model, the solution to a problem is encoded as the ground state (i.e., the lowest energy configuration) of a quantum system.

The process begins with qubits initialized in a uniform superposition – meaning each qubit exists in a state that represents both 0 and 1 simultaneously. The system is governed by a time-dependent Hamiltonian (an energy function), which gradually evolves from a simple initial form (where the ground state is known) to a final Hamiltonian that encodes the problem to be solved. This evolution follows a carefully tuned annealing schedule.

Throughout the process, two key quantum phenomena are at play:

1. Quantum Superposition allows the system to explore many possible solutions at once.
2. Quantum Tunneling enables the system to escape local minima in the energy landscape – solutions that are "good but not optimal" – by passing through energy barriers that classical algorithms would have to climb over.

If the annealing is performed slowly enough (a condition known as adiabatic evolution), the system ideally ends in the ground state of the final Hamiltonian, which corresponds to the optimal or near-optimal solution.

However, real-world quantum annealing faces several limitations. Noise, thermal effects, and imperfections in the hardware can prevent the system from reaching the true ground state. Additionally, not all problems can be naturally expressed in a form suitable for annealing, often requiring an intermediate step called minor embedding to map them onto the hardware's limited qubit connectivity.

In practice, D-Wave's systems solve problems represented in Quadratic Unconstrained Binary Optimization (QUBO) or Ising model form. These must be translated from high-level business or scientific use cases, a process that can be complex and time-consuming.

Full Legal Disclaimer

As of the publication date of this report, Kerrisdale Capital Management LLC and its affiliates (collectively "Kerrisdale") have short positions in the stock of D-Wave Quantum Inc. (the “Covered Issuer”). In addition, others that contributed research to this report and others that we have shared our research with (collectively with Kerrisdale, the “Authors”) likewise may have short positions in the stock of the Covered Issuer. The Authors stand to realize gains in the event that the price of the stock decreases.

This report is not a recommendation to short or sell shares of any company, including the Covered Issuer, and is only a discussion of why Kerrisdale is short the Covered Issuer. We are not your financial advisor and we do not owe a fiduciary duty to you. We don’t recommend that you do anything whatsoever – we don’t even know who you are.

Following publication of the report, the Authors will transact in the securities of the Covered Issuer. The Authors may buy, sell or short shares of the Covered Issuer and other securities covered herein subsequent to publication. The Authors will continue to transact in the Covered Issuers’ securities for an indefinite period, and such position(s) may be long, short, or neutral at any time thereafter regardless of the Authors’ initial position(s) and views as stated in this report. Kerrisdale will not update this report to reflect changes in its positions.

All content in this report represents the opinions of Kerrisdale. The Authors have obtained all information herein from sources they believe to be accurate and reliable. However, such information is presented “as is,” without warranty of any kind – whether express or implied. The Authors make no representation, express or implied, as to the accuracy, timeliness, or completeness of any such information or with regard to the results obtained from its use. All expressions of opinion are subject to change without notice, and the Authors do not undertake to update or supplement this report or any information contained herein.

This document is for informational purposes only. It is not intended as an official confirmation of any transaction. All market prices, data and other information are not warranted as to completeness or accuracy and are subject to change without notice. The information included in this document is based upon selected public market data and reflects prevailing conditions and the Authors’ views as of this date, all of which are accordingly subject to change. The Authors’ opinions and estimates constitute a best efforts judgment and should be regarded as indicative, preliminary and for illustrative purposes only.

This report discusses estimated fair values of securities and companies, utilizing valuation methodologies. Such estimated fair values are not price targets and the Authors will not hold securities until such estimated fair values are reached. The Authors may change their estimates of fair values at any time in the future without updating this report or disclosing the new fair values publicly. The Authors will also transact in the securities of the Covered Issuer and any companies covered herein for many reasons that have nothing to do with the Authors’ estimates of the securities’ fair values. The estimated fair values only represent a best efforts estimate of the potential fundamental valuation of a specific security, and are not expressed as, or implied as, assessments of the quality of a security, a summary of past performance, or an actionable investment strategy for an investor.

Any investment involves substantial risks, including, but not limited to, pricing volatility, inadequate liquidity, and the potential complete loss of principal.

This document does not in any way constitute an offer or solicitation of an offer to buy or sell any investment, security, or commodity discussed herein or of any of the affiliates of the Authors. Also, this document does not in any way constitute an offer or solicitation of an offer to buy or sell any security in any jurisdiction in which such an offer would be unlawful under the securities laws of such jurisdiction. To the best of the Authors' abilities and beliefs, all information contained herein is accurate and reliable. The Authors reserve the rights for their affiliates, officers, and employees to hold cash or derivative positions in any company discussed in this document at any time. As of the original publication date of this document, investors should assume that the Authors are short shares of the Covered Issuer and stand to potentially realize gains in the event that the market valuation of the company's common equity is lower than prior to the original publication date.

The Authors shall have no obligation to inform any investor or viewer of this report about their historical, current, and future trading activities. In addition, the Authors may benefit from any change in the valuation of any other companies, securities, or commodities discussed in this document.

Kerrisdale does not provide investment advice to the readers of its reports. You understand and agree that Kerrisdale does not have any investment advisory relationship with you and does not owe fiduciary duties to you. Giving investment advice requires knowledge of your financial situation, investment objectives, and risk tolerance, and Kerrisdale has no such knowledge about you. In no event shall Kerrisdale and the Authors be liable for any claims, losses, costs or damages of any kind, including direct, indirect, punitive, exemplary, incidental, special or consequential damages, arising out of or in any way connected with any information presented in any Kerrisdale report. This limitation of liability applies regardless of any negligence or gross negligence of Kerrisdale and the Authors. You accept all risks in relying on the information presented in this report.

The information contained in this document may include, or incorporate by reference, forward-looking statements, which would include any statements that are not statements of historical fact. Any or all of the Authors' forward-looking assumptions, expectations, projections, intentions or beliefs about future events may turn out to be wrong. These forward-looking statements can be affected by inaccurate assumptions or by known or unknown risks, uncertainties and other factors, most of which are beyond the Authors' control. Investors should conduct independent due diligence, with assistance from professional financial, legal and tax experts, on all securities, companies, and commodities discussed in this document and develop a stand-alone judgment of the relevant markets prior to making any investment decision.

You agree that any dispute between you and Kerrisdale or the Authors arising from or related to this report or viewing the material presented herein shall be governed by the laws of the State of Florida, without regard to any conflict of law provisions. The failure of Kerrisdale to exercise or enforce any right or provision of this Full Legal Disclaimer shall not constitute a waiver of that right or provision. If any provision of this Full Legal Disclaimer is found by a court of competent jurisdiction to be invalid, the parties nevertheless agree that the court should endeavor to give effect to the parties' intentions as reflected in the provision, and that the other provisions of this Full Legal Disclaimer remain in full force and effect, in particular as to the governing law and jurisdiction provision. You agree that regardless of any statute or law to the contrary, any claim or cause of action arising out of or related to this report or related material must be filed within one (1) year after the occurrence of the alleged harm that gave rise to such claim or cause of action, or such claim or cause of action be forever barred.