Beyond the breakthrough: Who will capture quantum's lasting value?

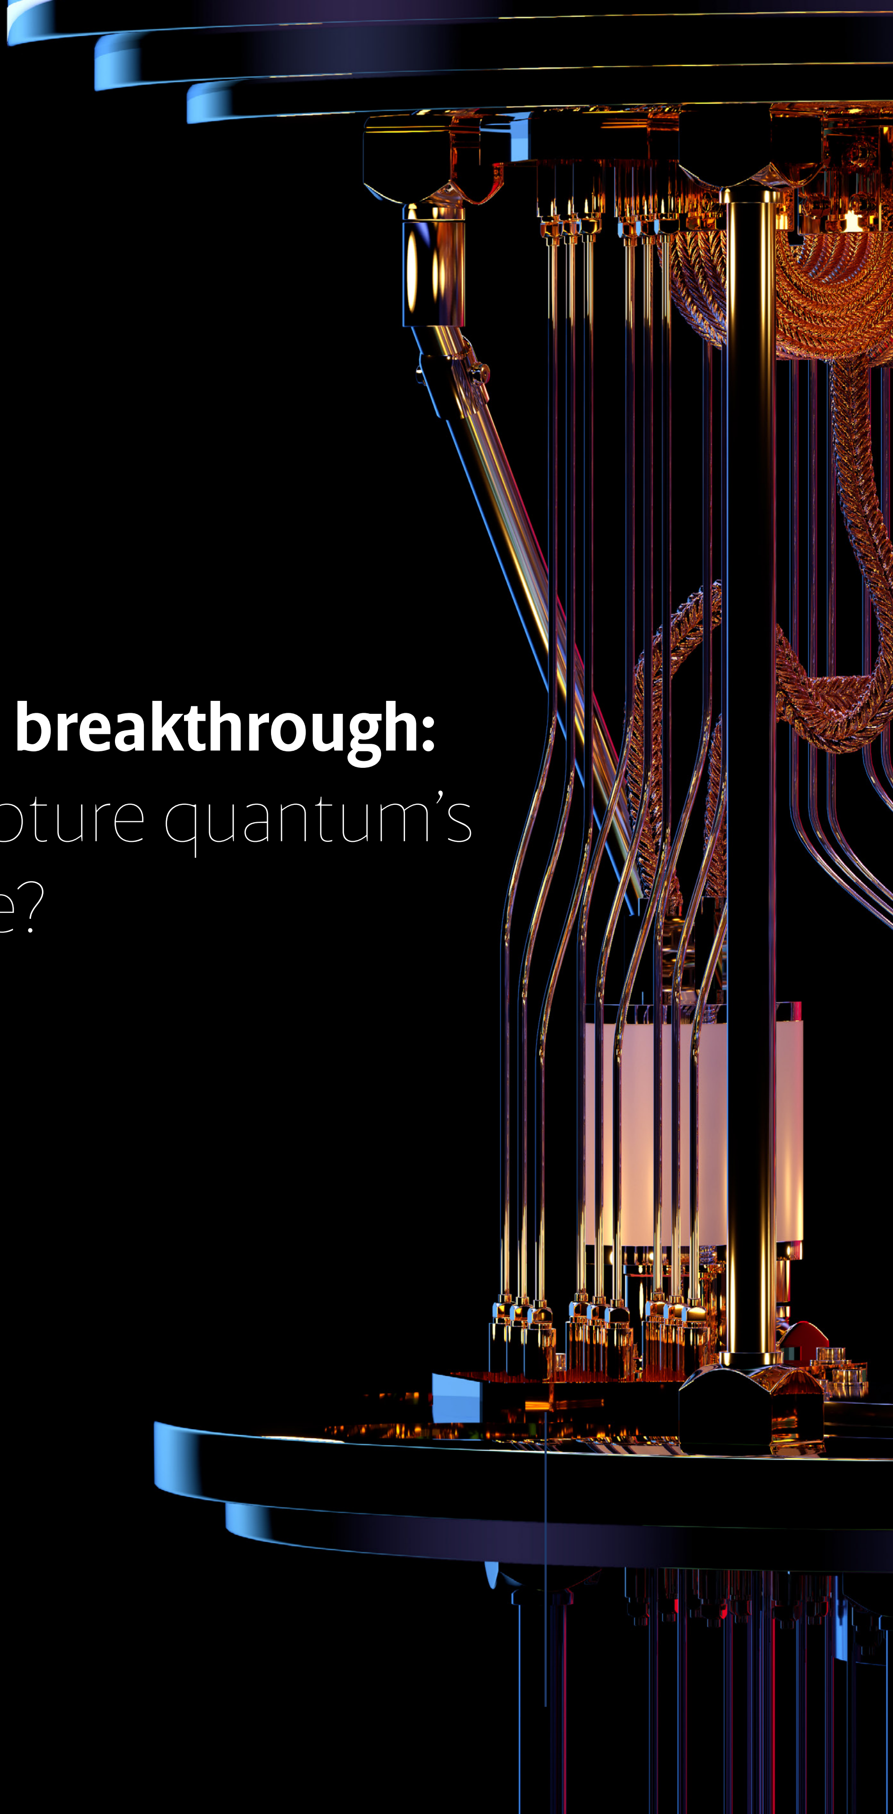

novo holdings

Contents

2 From capability to capture: where durable value will accumulate in quantum

5 Executive summary

8 The question: where will value accumulate in the quantum stack?

10 What the semiconductor industry teaches us: value formed at specific control points, not in any one layer

14 Mapping the quantum stack: where the winnable control points are

21 Life science: case study in value beyond hardware

27 Conclusion: where value will accumulate

30 References

From capability to capture:

Where durable value will accumulate in quantum

Life science as a worked case study in application-layer value creation

For most of the last decade, the hard question in quantum computing was scientific: can useful fault-tolerant machines be built? The questions that matter now are where value will form and who will capture it. These are being decided now, because the positions that will matter are taken before the technology matures, not after.

This paper is written by the Novo Holdings, Seed Investments – Quantum Team from the vantage point of an investor at the intersection of quantum and life science. It is an analysis, not a prospectus, and draws on public roadmaps, peer-reviewed work and market data, set against the landscape as we see it today. The paper stays conditional where the evidence is thin. It is not an investment thesis, a forecast or a defence of our portfolio.

Almost every quantum strategy rests on one assumption: that a company's altitude in the stack, hardware at the base and applications at the top, predicts the value it captures. This paper tests that assumption. It takes the historical precedent of semiconductors, examines where durable value formed and why, and applies the result to the quantum stack, and in more detail to life science. The result is a test, not a forecast. It puts no date on the outcome and does not claim quantum will unfold in the same order; it tells the reader what to look for.

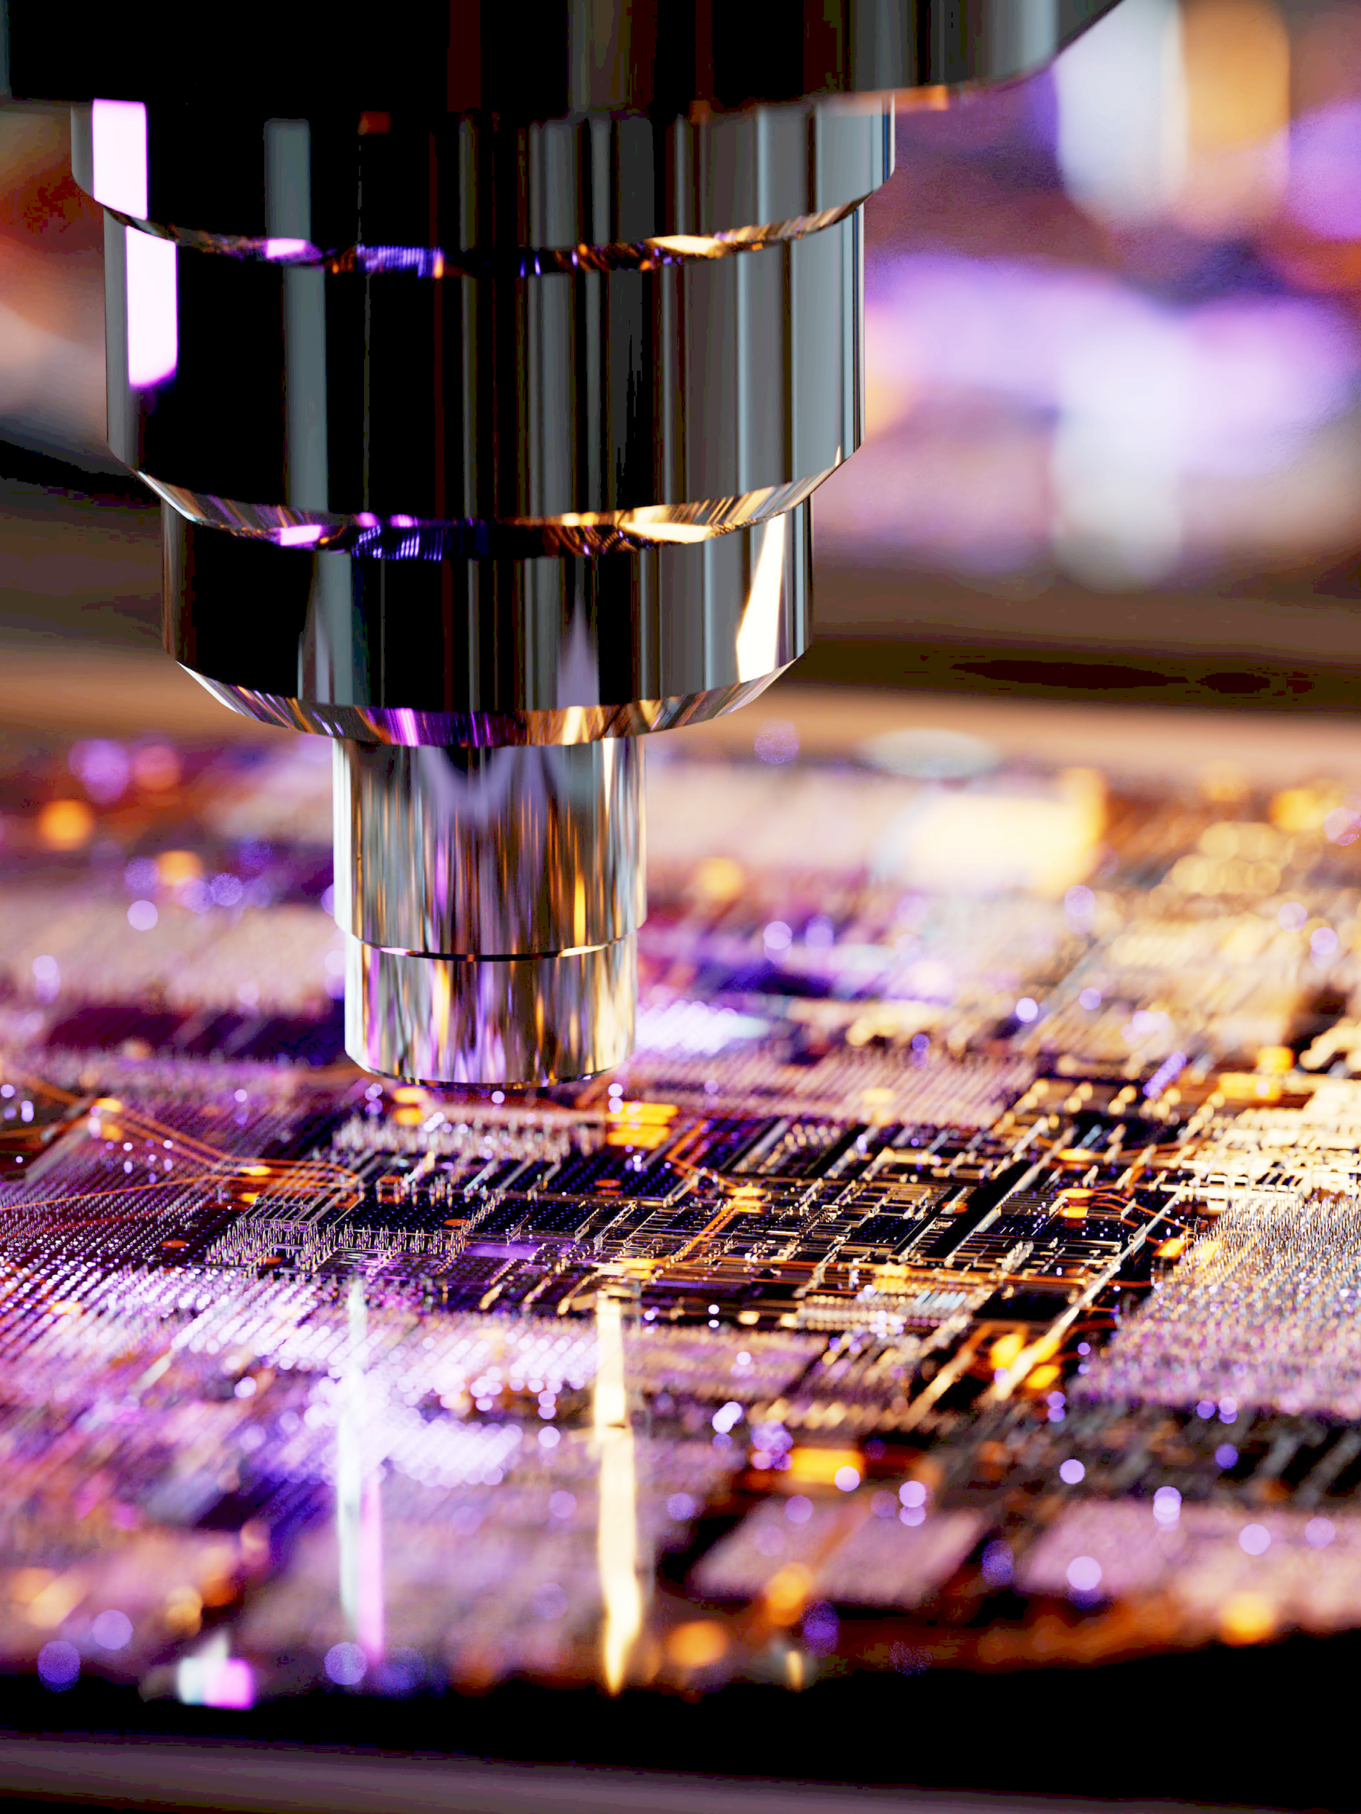

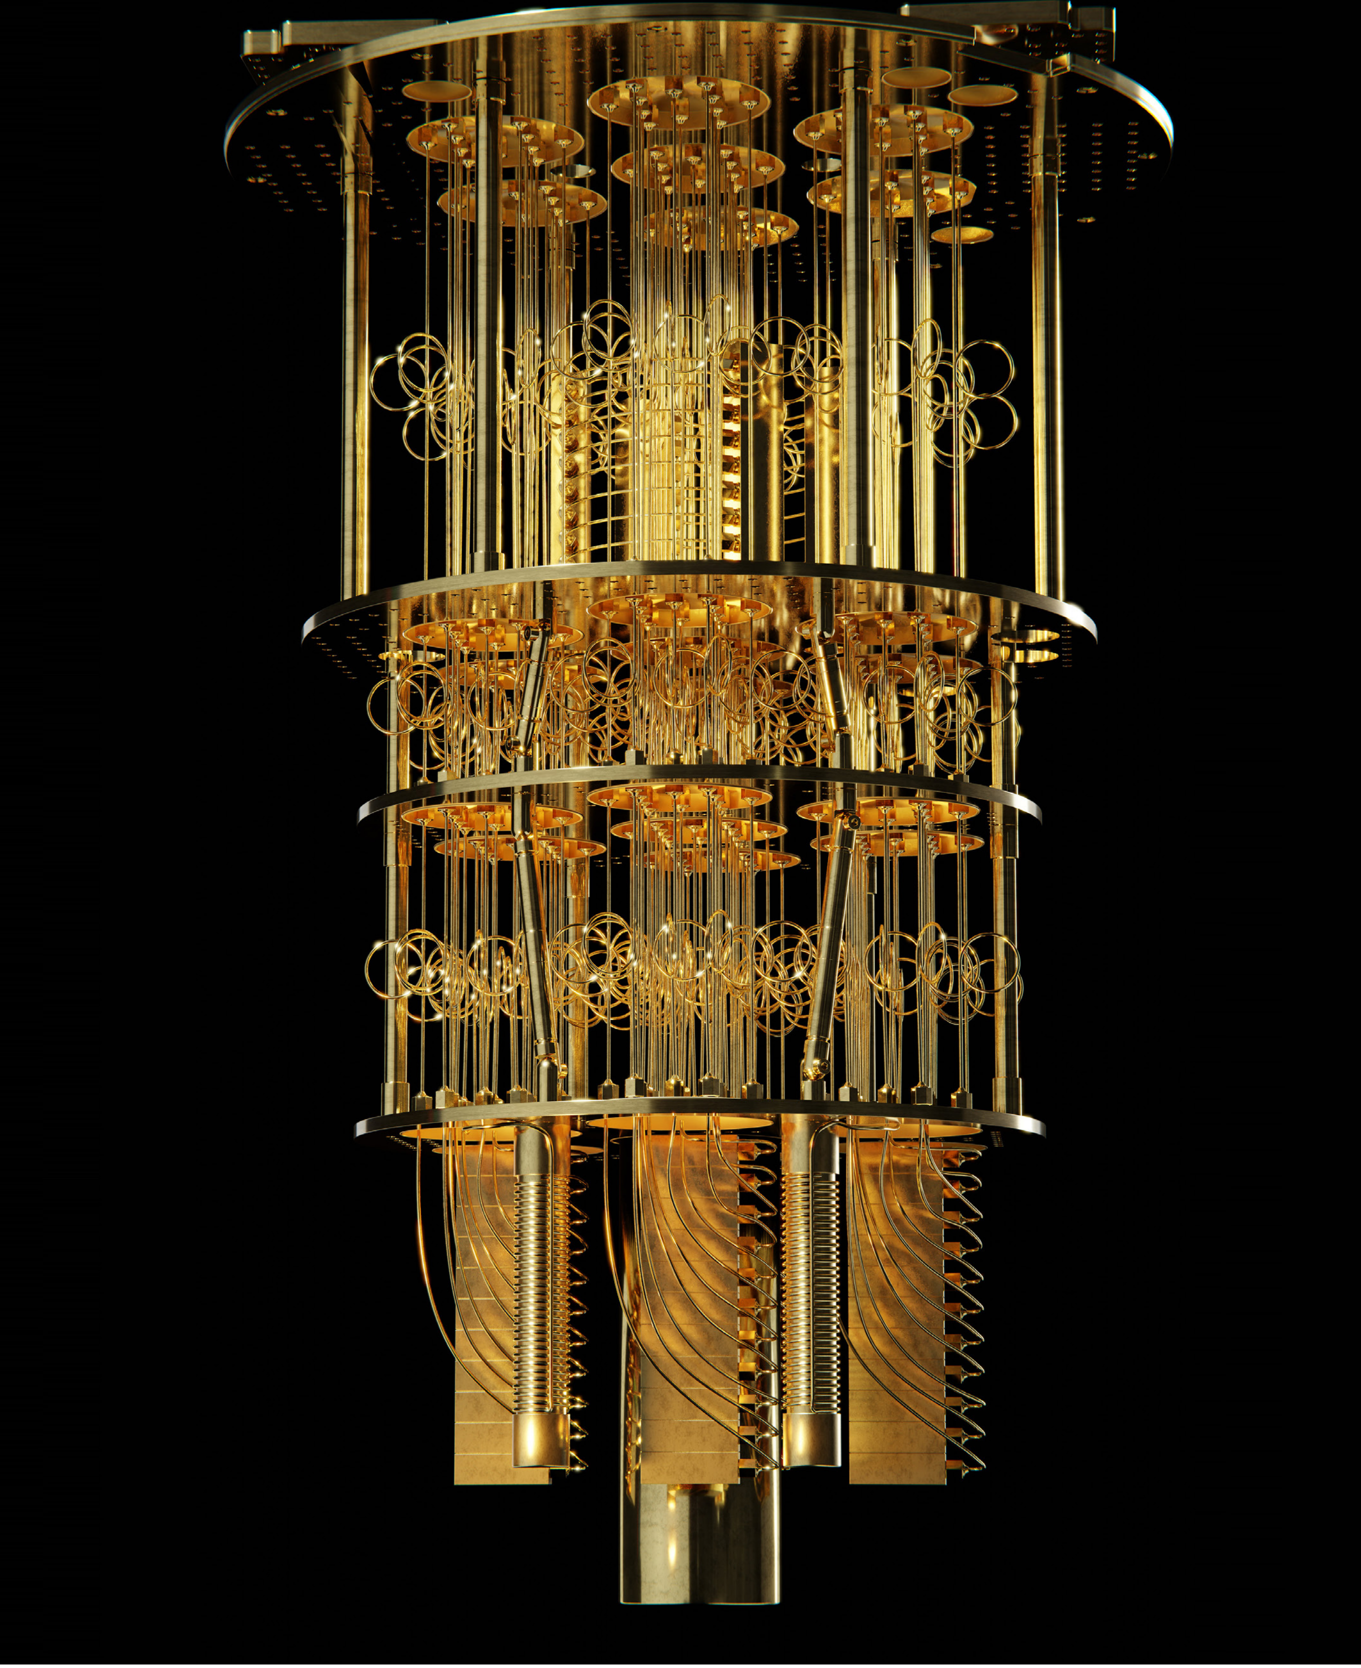

1 Executive summary

Durable value in quantum will not accrue to a layer of the technology stack but to control points: capabilities that are hard to reproduce, that customers build into their day-to-day decisions and that independent companies can own. Today the opportunity lies in the application and algorithms layer, with life science as the clearest case. That window exists now, before fault tolerant systems arrive, and narrows as they become widely used. The time to back these companies is during this window, not after it.

For investors, the objective is not simply to maximise exposure to quantum. It is also to identify companies whose capabilities are becoming harder to reproduce, whose outputs are becoming harder for customers to replace, and whose ownership of algorithms, data, validation and customer workflows strengthens as quantum hardware matures from emergent early fault tolerance to fault tolerance at scale. The case is conditional: it holds where the workloads prove technically feasible, commercially meaningful and ownable, and where classical methods do not close the gap first.

Five conclusions explain why:

1. Value will not follow stack position. In the one deep-technology stack that has run to maturity, durable returns accrued across different layers, never to any single one, and went to the companies that turned a scarce capability into a customer dependency they could own. Altitude in the stack did not predict capture. Control did, and which control points were winnable depended on the industry's phase and timing.

2. Future winners forge controllable positions and capture the key nodes an ecosystem cannot bypass. A position becomes a durable control point only where three conditions hold together: the capability is scarce, it becomes a dependency customers cannot easily replace, and a single company can own it and keep the economics. Scarcity, dependency and ownership are necessary and none is sufficient. This paper applies that test to quantum.

3. Hardware creates the capability, but applications turn it into value. Hardware holds some of the clearest technical scarcity today, but the value of these systems depends on what can be done with them. Domain-specific algorithms and verticalised platforms combine real technical scarcity with a dependency customers build into their decisions, stronger ownership potential and lower capital intensity. The window to carve out a leading position within quantum applications is opening now, as the first constrained fault-tolerant machines are targeted for the late 2020s to early 2030s, and narrows as workflows begin to embed and data and validation compound.

4. Life science is the case that shows the framework in full. Life science is a high value industry long accustomed to adopting new science, engineering and technology for competitive advantage. Its documented R&D productivity problem, combined with a demonstrated willingness to pay [1,2] for better outcomes, means the revenue for a better result exists before quantum advantage does. The investable unit is the specific bottleneck, not the sector, and where classical methods may close the gap first, the defensible asset is the validated workflow and proprietary data, not the algorithm primitive itself.

5. Capital and value have separated, and the gap is the opportunity. Roughly 70% of the last decade's quantum investment went into hardware and components [3], leaving the application layer that turns quantum resources into customer decisions underfunded relative to its option value. The next phase of the industry's development must finance application formation alongside hardware, backing companies that own the problem, the workflow, the validation and the data before fault tolerance arrives.

Hardware holds some of the clearest technical scarcity today, but the value of these systems depends on what can be done with them

---

Capital and value have separated, and the gap is the opportunity

2

The question:

Where will value accumulate in the quantum stack?

Where value forms in the quantum technology stack will not be decided by a company's position within it.

Every serious quantum strategy already embeds a view on where value will form, whether it is a national programme, a venture portfolio or a corporate partnership. For most of the field's history that view could hide behind a scientific question: can useful fault-tolerant machines be built? That question is already being worked on, and it is no longer the one that should occupy an investor. Quantum computers are improving, capital is already committed across the hardware stack, and returns will be decided long before the science is settled. Waiting for certainty is itself a position, but an expensive one, because the ownable positions are built before the technology matures, not after.

The economic question is the one that matters: as fault-tolerant systems arrive, where will value form, and which companies will capture it? The conventional answer reduces to a choice between layers. Hardware creates the capability and may stay scarce. Software promises the gross margins investors like. Applications sit closest to the customer. Each view has merit, and each rests on the same assumption: that a company's altitude in the stack is a reliable guide to the value it captures. This paper tests that assumption and provides an alternative framework for investors.

FIGURE 1

Quantum funding by category excluding public-market raises, 2014–2026.

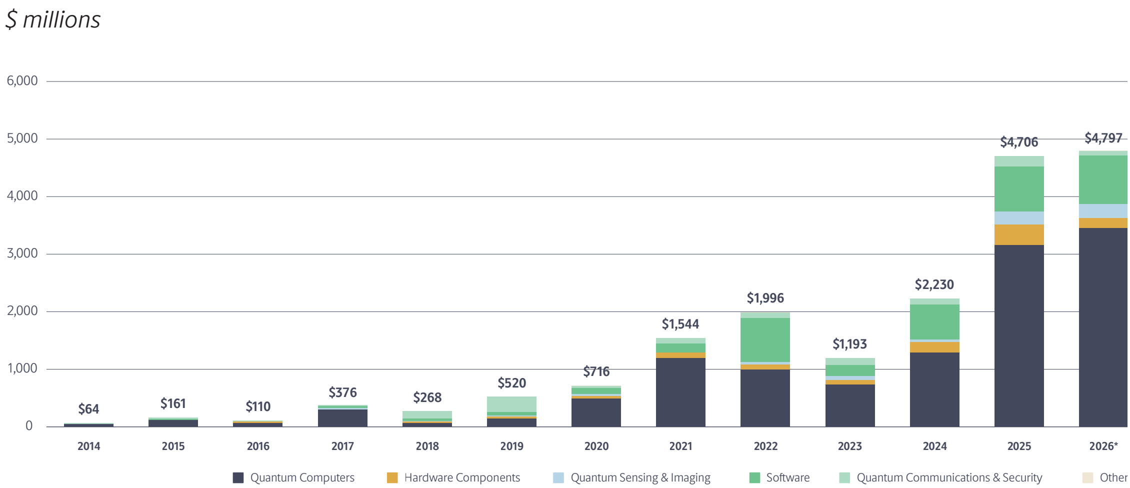

** 2026 is year to date (2 September 2026); the year is not complete.*

Disclosed private financings only. Post-IPO placements, ATM programmes, SPACs and IPOs are excluded. Government programme funding is excluded.

Quantum companies only. Companies whose funded business is photonic or AI computing, general semiconductor and telecom optics, display materials, classical security or AI are excluded.

Source: The Quantum Insider funding database, extracted 8 September 2026. Prior-year figures are restated as rounds are added retrospectively.

Of roughly $13.9 billion invested across quantum-computing categories between 2014 and 2025, around 70% went to hardware and components, with the bulk of the remainder to quantum sensing, networking & security, software, algorithms and applications, a concentration that has persisted into 2026 and understates the gap, since “software” also includes compilers, control tooling and error correction rather than the companies solving customer problems (Figure 1). Rightfully so, investment has focussed upstream. Whether that is where the value will settle is the question and one comparison frames it: by 2035 the economic value at stake for the end-user industries adopting quantum technologies is estimated at $1.3 to $2.7 trillion, an order of magnitude above the roughly $100 billion in revenue expected to accrue to quantum technology providers over the same period [4].

We think the sharper form of it is narrower. Not which layer wins, but which specific positions become control points: capabilities that are hard to reproduce, that customers cannot easily replace, and that a company can own to capture the economics. Some technologies are scientifically important yet never anchor a business to scale, diffusing through publication, bundled by a platform, or internalised by a customer. A smaller number become difficult to bypass. The aim is to identify those positions before the market prices them.

3 What the semiconductor industry teaches us:

Value formed at specific control points, not in any one layer

In the one deep-technology stack to reach maturity, durable value did not follow stack position, sophistication or margin. It accrued to positions the ecosystem could not bypass. That pattern is the test we carry into quantum.

Quantum is too young to show where its profit pools will settle, so we look to the closest matured precedent. Semiconductors share quantum's foundations, deep science, specialised supply chains, long development cycles and tight coupling between hardware, software and applications, but they have completed the journey from laboratory result to global infrastructure. That journey ran through phases, from vertically integrated manufacturers to a specialised and interdependent stack, to scaling and consolidation, and finally to a small set of strategic control points around which value has concentrated (Figure 2). Their history settles one question that matters for quantum: does value follow the layer? It does not. Value concentrated in a handful of positions the rest of the ecosystem could not work around, and those positions sat at very different points in the stack.

FIGURE 2

From vertical integration to strategic control points across the semiconductor stack [7].

A. **Industry Evolution:** How the Semiconductor Stack Emerged and Evolved

1960s – 1980s	1980s – 2000s	2000s – 2010s	2010s – present
Vertical Integration	**Specialisation**	**Scaling & Consolidation**	**Strategic Control Points**
Design, process, technology, manufacturing and packaging were largely internal to Integrated Device Manufacturers (IDMs)	Stable interfaces enabled Electronic Design Automation (EDA), IP, fabless, foundry and OSAT (packaging and testing) specialists to emerge	Rising complexity and fixed costs increased barriers to entry, encouraged consolidation and reduced the number of viable players	Value concentrated around a limited number of scarce, embedded, scaled or integrated capabilities that customers could not readily bypass.

Interface standardisation — Specialist scale economies — Consolidation — Strategic control points

B. Semiconductor stack with representative winners and winning archetypes

LAYER OF THE STACK	REPRESENTATIVE WINNER	WINNING ARCHETYPE	EVIDENCE OF WHY THE POSITION WINS
Equipment **Provide tools and manufacturing systems**	**ASML** Only commercial supplier of EUV lithography systems	**Indispensable Bottleneck:** A scarce technology or capability customers cannot readily bypass	• Leading-edge lithography is a technological choke point • Extreme complexity and the supplier ecosystem create high barriers to substitution • Installed base drives services and upgrade value
EDA **Electronic Design Automation**	**SYNOPSYS** $11.4bn backlog. 71% time-based and maintenance/ service revenue [5]	**Embedded Platform:** Proprietary design tools integrated into mission-critical workflows and processes	• Tools become embedded in mission critical workflows • Recurring licences, support and cloud access expand revenue over time • Contracted backlog provides substantial revenue visibility
Semiconductor IP **Reuseable intellectual property and architectures**	**arm** 350bn+ Arm-based chips shipped. Present in 99%+ of smartphones	**Standards / IP Tollbooth:** Foundational architecture and IP that is widely adopted across the industry	• Broad ecosystem adoption drives ubiquitous deployment • Licensing and per-chip royalties convert adoption into high-margin recurring revenue • Low marginal cost to license additional customers
Fabless Design **Design chips and outsource manufacturing**	**NVIDIA** 6m+ CUDA developers & 900+ CUDA-X libraries and AI models	**Ecosystem Orchestrator:** Owns a platform of hardware, software and developers reinforcing one another	• CUDA, networking and system software create a developer ecosystem • Better software attracts more developers; more developers drive demand for the platform • Expands value capture above the component level
Foundry **Manufacture chips for multiple customers**	**tsmc** 534 customers & 12,682 products across 305 technologies	**Neutral Scale Infrastructure:** Owns essential manufacturing capacity at scale with customer neutrality	• Scale and process leadership support low costs and best-in-class yields • Institutional knowledge, high utilisation and long-term customer trust reinforce scale • Neutrality attracts the industry's leading designs
IDM **Design and manufacture owned products**	**Texas Instruments** 80,000+ products, serving 100,000+ customers	**Focused Process-Product Model:** Applies deep expertise to defined processes and product categories	• Proprietary processes, a broad portfolio and scale at mature nodes. • Long product lives and strong distribution create durable cash flows • Enduring customer relationships and reliability are hard to replicate
Packaging & Assembly **Integrate, package and test semiconductor devices**	**ASE** Packaging and testing revenue +20% YoY. Services reached $1.6bn in 2025 [6]	**Integration Specialist:** Controls complex integration, packaging, testing and lifecycle services at scale	• Advanced packaging and integration increase system performance and efficiency • Execution history, qualification and scale create high customer dependence • Turnkey services embed the company more deeply in customer roadmaps and workflows

Key Takeaway: Value capture is not inherent to a layer or to an asset-light or capital-intensive model. It emerges when a company converts its position into a scarce capability, embedded workflow, adopted standard, trusted scale platform or difficult-to-replace integration function.

The winners prove the point precisely because they have so little else in common. ASML owns the lithography step without which leading-edge chips cannot be made. TSMC turned manufacturing scale and process leadership into a position its customers design around. Arm made its architecture the standard that an industry licenses. Asset-light and capital-heavy, upstream and midstream, high- and low-margin: what they share is not their place in the stack or their financial profile, but a single progression. A scarce capability became a dependency customers built into their operations, and each company kept control of the interface, standard or process through which it was expressed.

That progression is the test. A position becomes a durable control point only where three conditions hold together (Figure 3): the capability is scarce and hard to reproduce within a commercially relevant time and cost; it becomes a dependency customers cannot easily or safely replace; and a single company can own it, retaining the economics rather than seeing them diffuse to customers, platforms or open ecosystems.

The logic is cumulative. Scarcity without dependency stays peripheral; dependency without ownership enriches everyone but the company that built it; only where all three hold does value compound. The precedent adds one further lesson. These positions were won while their interfaces were still forming, and became difficult to contest once those interfaces set: ASML's and Arm's positions could not be created from scratch today, because the switching costs are now prohibitive. Timing, not position alone, decides what is winnable.

FIGURE 3

How technical capability becomes durable value capture. A capability must be difficult to reproduce, become sufficiently embedded that customers find it difficult or risky to replace it, and be controlled by a company able to retain part of the resulting economics. Durable value emerges only where scarcity, dependency and ownership reinforce one another.

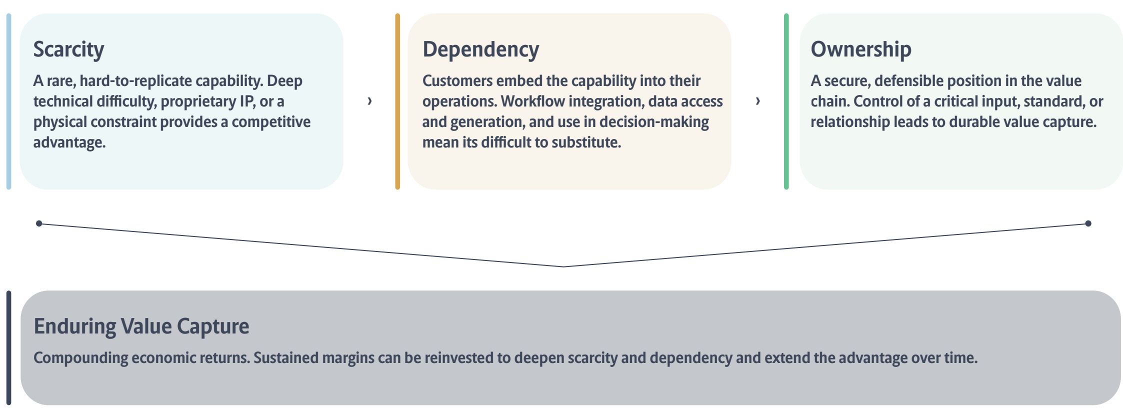

The test tells you what to look for. It does not forecast how quantum will unfold and it does not tell you whether a position is a good investment. On the first, quantum is far earlier than semiconductors were, its interfaces unformed and its architecture contested. We do not claim it will disaggregate in the same order. On the second, showing a defensible position can exist is not the same as showing it is worth backing. That needs a second screen: market size, capital intensity, time to revenue, entry valuation, and the risk of the position being internalised by an incumbent or platform. A company can hold an excellent position and still be a poor investment at the wrong price, terms or timing.

4

Mapping the quantum stack:

Where the winnable control points are

Run the same test on quantum today and it points to the application layer, the one position that is scarce, ownable and still unclaimed.

The control points that endured in semiconductors: lithography, foundry, the architectural standard, were upstream. It would be easy to read that as a forecast that value in quantum will settle upstream too. That is a misread of the lesson. What those positions shared was not their place in the stack but that each was scarce, became a dependency, and could be owned. Semiconductors concentrated value upstream because that is where those three conditions happened to converge as the industry matured. The lesson is the test, not the outcome.

It is worth pausing on one objection. Part of what drove value out of chip design and assembly was commoditisation: as those layers standardised, their margins thinned and value migrated to the positions that had not standardised. Quantum may not commoditise for a long time, if ever, so it is fair to ask whether the lesson survives. It does, because the control points were created by scarcity, not commoditisation: ASML and TSMC captured value because their positions could not be reproduced or bypassed, which can happen in a market that never commoditises. What the argument requires from quantum is therefore narrow: not that any layer becomes a commodity, but that a position becomes a scarce, unbypassable dependency. Where quantum remains differentiated rather than commoditised, that does not weaken the test; it changes only which positions the three conditions identify.

Applied to quantum in its current phase, the same test lands in a different place. Beyond the coarse hardware-software-applications split, we divide the stack into eight positions and evaluate each against present scarcity, the dependency it could build, whether an independent company could own it, and the capital this would require (Figure 4). This is a directional judgement over a three-to-five-year horizon. Read against those axes, three groups behave differently.

FIGURE 4

Directional value-capture map of the quantum stack. The x-axis shows present technical scarcity, y-axis shows potential customer dependency. Bubble fill indicates independent ownership potential and bubble size is capital requirement.

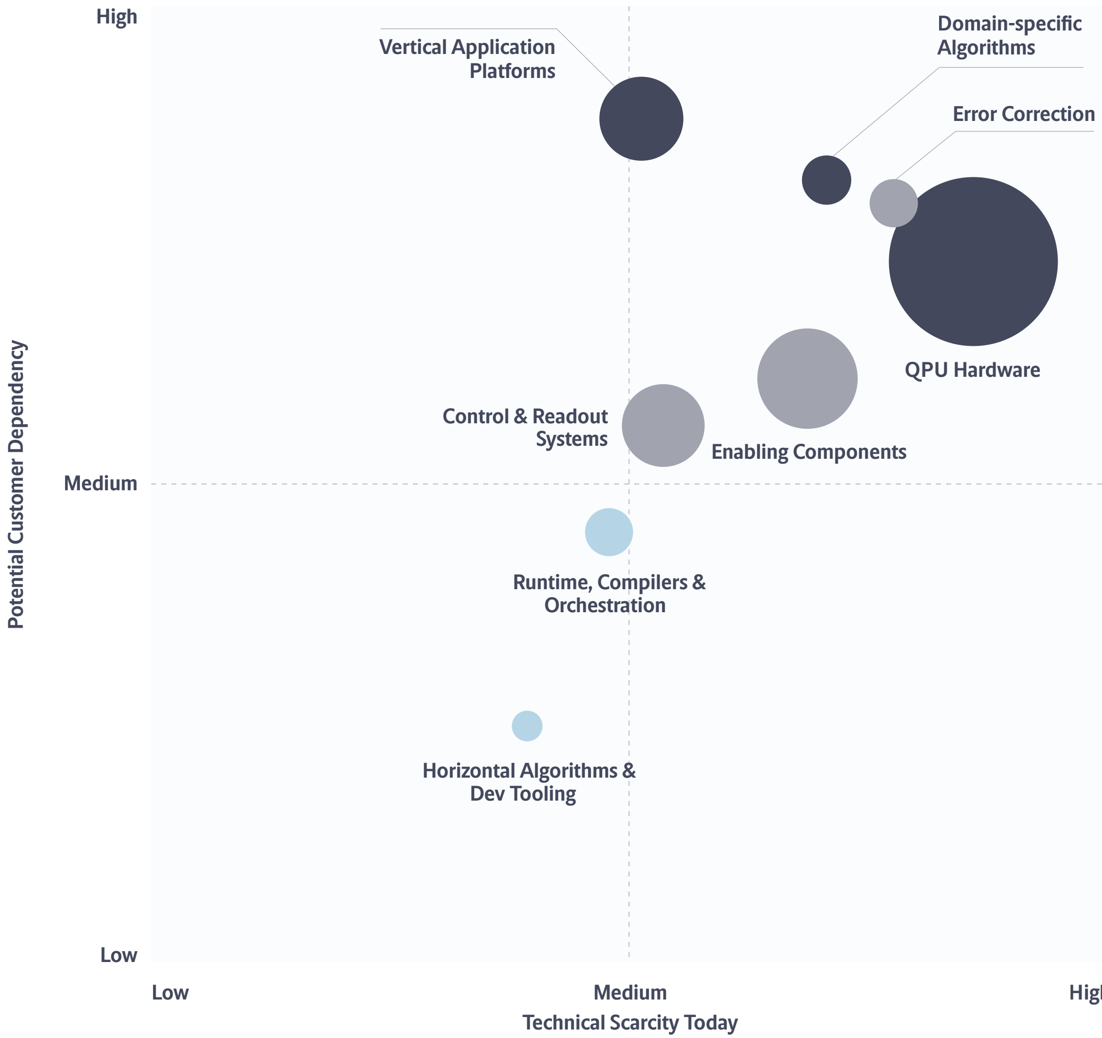

Bubble fill is startup ownership potential: light-blue (low) - mid-blue (medium) - navy (high)

Bubble size is indicative capital intensity: small (low) - medium - large (high)

**Bubble size indicates capital requirements, not market size or investment attractiveness*

Methodology Note: Positions are our directional judgements as of September 2026. Technical scarcity today is rated low, medium or high based on replication difficulty, specialised talent, proprietary know-how, time and capital requirements, and demonstrated performance. Potential customer dependency is a three-to-five-year forward judgement based on workflow integration, technical qualification, switching risk, data accumulation and decision criticality. Bubble fill indicates the potential for an independent company to retain the resulting economics through intellectual property, infrastructure, contracts, data or customer integration. The positions are not measured based on market data.

The deepest technical scarcity today sits upstream, in enabling components and fabrication, quantum processing hardware, and control and systems engineering. These are formidable positions, capital-intensive and still contested: the winners yet to be settled; new architectures continue to emerge; and value will concentrate in a few suppliers that turn scientific performance into manufacturable systems. Whether those capabilities remain modular or are absorbed into vertically integrated platforms is an open question, and an important one.

The second group carries a genuine commodity risk because error correction, compilation, runtimes and horizontal developer tools can be essential to the ecosystem without being ownable by anyone. The ecosystem may depend on compilation while several interchangeable compilers compete, or while the function is bundled into hardware and cloud platforms as a feature. Essential is not the same as ownable. On present evidence one could expect error correction and compilation to become dependencies without a durable independent owner, creating important markets, weak individual positions. The exception the same test allows is a toolchain that automates one of these functions well enough to become the standard on which others build. At that point it stops being an interchangeable feature and starts to satisfy the scarcity and dependency the test requires. This is a judgement about where ownership is hard to establish.

The third group is domain-specific algorithms and vertical application platforms, which combine real technical scarcity with a dependency customers build into their decisions, stronger ownership potential, and materially lower capital intensity than hardware. The scarce asset is not an algorithm primitive, which publication reproduces, but the repeatable ability to turn a valuable problem into a better decision, carrying problem formulation, resource reduction, hardware-aware implementation, benchmarking and validation. Algorithmic efficiency also compounds the timing argument: a reduction in gate count, circuit depth or state-preparation cost can pull a valuable workload into an earlier hardware regime, so the best algorithm companies move applications forward along the hardware roadmap. Compilation is entangled with this, since how well a workload is compiled decides how much a given device delivers, and the hand-tuning algorithm teams rely on today will not scale as qubit counts grow. The ownable edge is therefore less any single algorithm than the tooling a company builds to automate its own algorithm design and compilation at scale, which compounds across problems so that each new one starts from a stronger base.

Several positions could become durable control points, but for an independent new entrant one stands out. The application layer, domain-specific algorithms and vertical platforms offer the strongest current combination of scarcity, a dependency it can build, ownership through validation and data, and is far less capital intensive than the rest of the stack. It is also largely unclaimed, because the industry is only now approaching early fault-tolerant regimes and no company has yet embedded deep enough to capture it.

FIGURE 5

Published quantum hardware roadmaps and indicative application windows [8]. Published vendor targets suggest that the first constrained fault-tolerant systems may begin to emerge around the end of the decade, with larger-scale systems following during the 2030s. These targets are directional rather than forecasts and are not directly comparable across architectures. Application readiness depends on the full resource envelope, not logical-qubit count alone (see Table 1 note) [9].

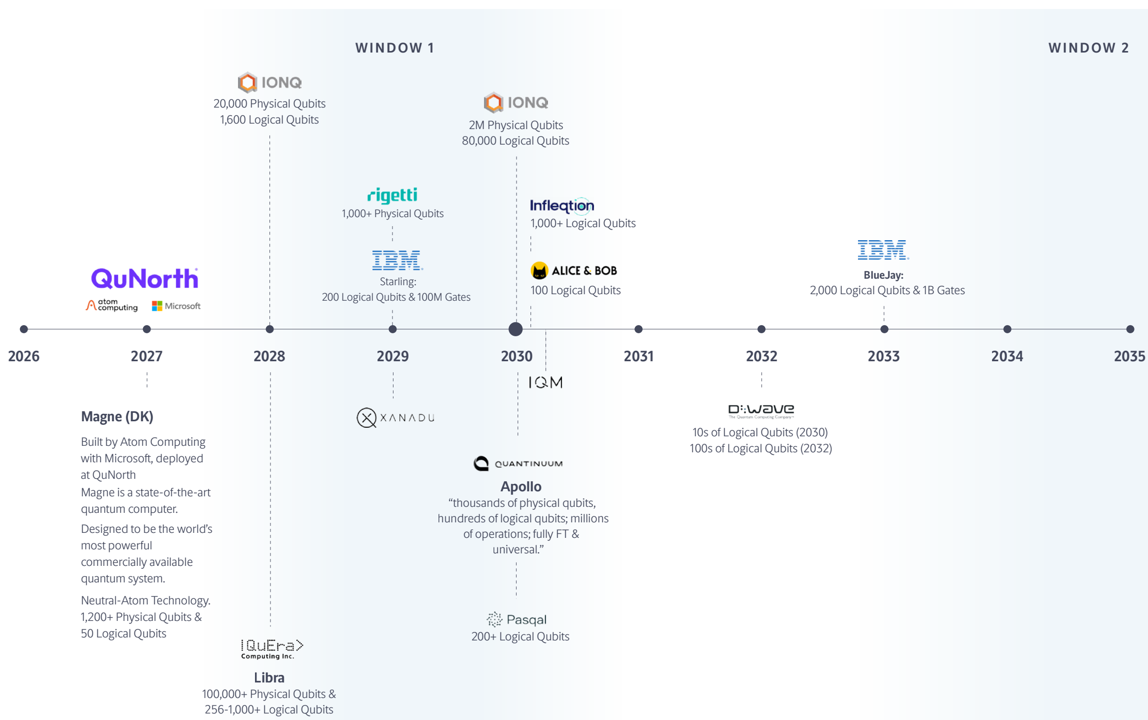

Timing decides when that position pays off, and quantum sits early in the phase sequence from Section 3: the interfaces are unstable, architectures are contested, and delivering a useful system often requires some vertical integration.

Table 1 therefore uses two scenario windows rather than a single forecast threshold.

Window 1 represents an early fault-tolerant regime in which logical resources, circuit depth, precision and throughput remain tightly constrained; the credible contribution is therefore narrow: a high-accuracy quantum subroutine handling a strongly correlated sub-problem inside a larger classical workflow, not an end-to-end quantum calculation.

Window 2 represents a scaled fault-tolerant regime with materially greater logical capacity, depth, repetition and throughput, enabling broader and more repeated use. The consequence for underwriting is direct: a business resting solely on Window 2 is underwriting a hardware roadmap, while a stronger one uses the period before and through Window 1 to build the customer relationships, workflows and platform or product it will later come to rely on.

For an investor, early is not a reason to defer; it is the reason to move. The interface has not yet formed, and that is precisely the condition under which the position can still be taken: a company that embeds now and compounds validation and data as the interface sets will hold it, while one that waits for the interface to settle before committing risks losing it. The window is widest before interfaces harden and switching costs rise, so early is when the position is most winnable.

TABLE 1

Two potential windows for fault-tolerant quantum value. The timings are indicative, and commercial readiness will depend on the full resource requirement, not logical-qubit count alone.

	Window 1: Early fault-tolerant scenario	Window 2: Scaled fault-tolerant scenario
Illustrative hardware envelope	Hundreds to low thousands of logical qubits Deep but tightly constrained workloads	Thousands to low tens of thousands of logical qubits. Materially greater depth, repetition and throughput of workloads.
Illustrative timing band	Late 2020s to early 2030s	Mid-2030s and beyond
Life-science reach	Selected active spaces, reaction centres, excited states and high-accuracy reference calculations within hybrid workflows.	Larger molecular systems, broader reaction networks, more extensive dynamics and repeated portfolio-scale workflows.
Commercial signals	Co-development revenue, first paid quantum-enabled outputs improving a defined scientific decision. Early indications of platform breadth.	Recurring platform use across programmes and deeper integration into pharma workflows.
Investment posture	Investable today where bridge revenue compounds the same assets required for scaled fault-tolerant quantum advantage.	Significant longer-term upside but should not be the sole underwriting case today.

These are scenario bands, not forecasts or application thresholds. Feasibility depends on the full resource envelopes including logical qubits, logical-gate volume, circuit depth, logical error rate, cycle time, runtime, target precision, state preparation, measurement and classical processing. A logical-qubit count is therefore necessary but not sufficient to place a workload in either scenario.

Recent implementation-driven work reinforces this caution: Quantinuum and SoftBank frame their use-case roadmap as an assumption-driven envelope rather than a deterministic forecast, while their quantum-chemistry benchmark shows that algorithm-level error-corrected execution has not yet reached break-even against the unencoded physical implementation [10].

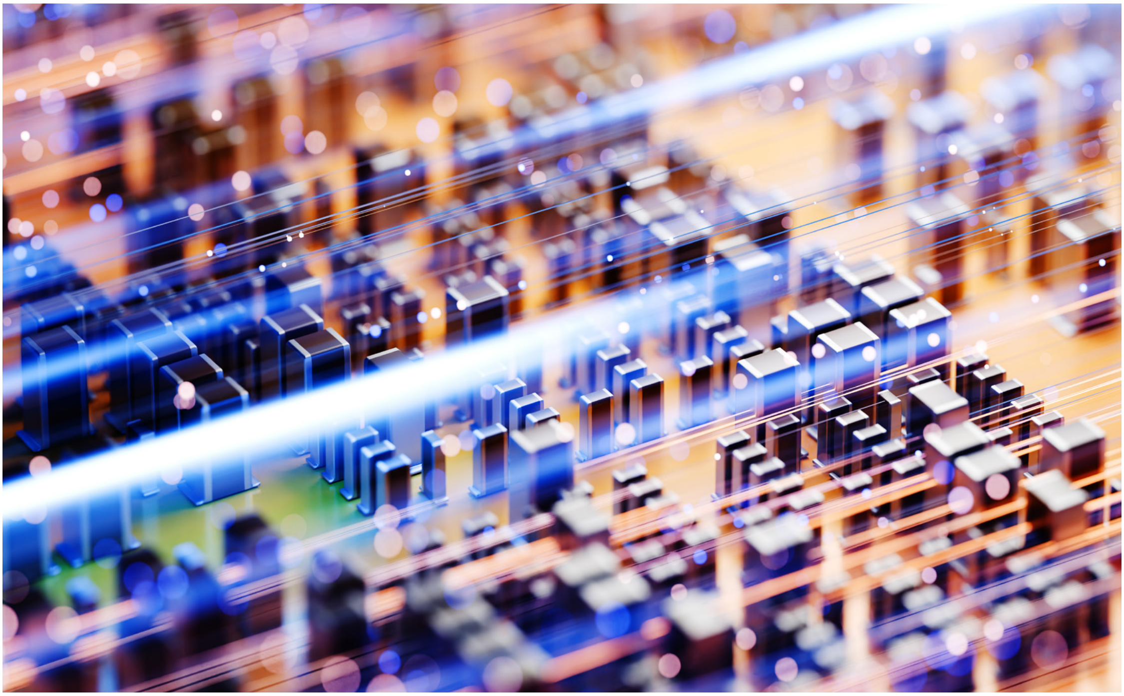

This is where the business model stops being a labelling question and becomes the mechanism that allows startups to capture the control point. Many early pitches assume software-like economics: proprietary algorithms sold through licences or usage-based APIs. That model may arrive, but today it is ahead of a technology still coupled to hardware that does not exist at commercial scale, and a company cannot price a research-intensive, hardware-dependent capability as mature software before its performance is demonstrated. The choice is not consultancy versus SaaS. A company can sell a validated scientific outcome while progressively encoding problem formulation, computation, validation and reporting into a repeatable, software-enabled workflow. The right question is not what the contract is called but whether each engagement reduces marginal delivery effort, strengthens reusable IP and data, and grows the share of future work delivered without proportional growth in expert headcount. There is a discovery dimension to this as well: which specific application is worth owning is rarely obvious at the outset, so the early consulting is how a company finds it, and the signal to back is the point at which bespoke work begins converging on a repeatable product. That is how an early position is converted into an owned one: not by waiting for fault tolerance, but by building the dependency, workflow embedding, validation history and proprietary data, before the hardware scales. Pre-fault-tolerant revenue counts only when it compounds the same assets. Revenue from disconnected bespoke projects does not.

5 Life science:

Case study in value beyond hardware

Life science is where the test holds most completely: high-value business decisions, a demonstrated willingness to pay, and workflows a company can own.

The strongest investable position and the earliest technical benefit are different questions, and life science answers the first. This section takes simulation as the near-term case: of the three routes by which quantum could contribute to life science, simulation of the quantum-mechanical systems, optimisation of combinatorial choices and machine learning across biological data, simulation is the strongest, while machine learning and optimisation are real but adjacent[1] and sit outside the argument here.

Optimisation is the weaker of the two. Most real-world optimisation problems, such as planning routes or schedules, have to be rewritten into a mathematical form a quantum computer can accept, and that translation usually makes the problem harder rather than easier. The few cases that suit quantum are the rare problems that already fit that form, or ones that would need a genuinely new method to solve. Machine learning faces a different limit: the bottleneck is getting large amounts of ordinary data into the machine in the first place, so it fits best where the data is small or highly structured, such as networks or graph problems, rather than large biological datasets.

Materials, energy and catalysis are plausibly higher-priority and nearer-term beneficiaries on technical grounds, owing to smaller active spaces, strong-correlation problems that map naturally onto near-term machines and clear relevance in energy (batteries, photovoltaics, nuclear and fusion) and other industrial applications. But the investable question is different: where is the value of a better decision highest, where will customers pay for it and where can the company delivering it build something it can own? On that axis, life science is where all three conditions hold together most clearly, because customers already pay to improve these calculations and there is a clear line of sight for a startup to own that workflow.

[1] The adjacent space includes quantum generative models such as quantum GANs and quantum circuit Born machines. These are earlier-stage still and sit outside the near-term investable case considered here.

The claim is not that biology is a quantum-computing problem. Most life-science challenges are limited by experimental throughput, biological complexity and clinical execution, not computation.

The claim is narrower: a concentrated set of high-value molecular decisions relies on electronic behaviour, how bonds form and break, how charge moves, how excited states behave, that leading classical methods model poorly, and these calculations sit inside established workflows, can be tested experimentally, and compound through proprietary data and repeated use. Life science also carries the decisive commercial feature: a severe R&D productivity problem, high-value products, and a demonstrated willingness to pay for anything that improves a molecular decision, and ultimately a commercial outcome. It is also an industry long accustomed to decade-long development horizons, which makes it unusually tolerant of a technology whose payoff builds over time. That willingness is not abstract. The budgets that already fund computational chemistry, molecular modelling and contract research are attached to exactly these decisions, candidate selection, mechanism, property prediction, so a better calculation competes for money that already exists and is already pointed at the decision, rather than depending on a new quantum budget being created.

TABLE 2

How the major commercial verticals compare across technical fit, the value of a better result and the potential to build a defensible platform. No vertical dominates every axis. Life science is the case worked in full here because it is where strong quantum relevance, valuable decisions, experimental validation and recurring workflows come together at once, which is the combination that makes a position ownable and worth paying for.

Industry	Technical Fit	Economic Impact	Platform Defensibility	Investment View
Life Science	●●●●●	●●●●●	●●●●○	Intrinsic quantum problems, high-value decisions, experimental validation and persistent workflows
Materials, Energy & Chemicals	●●●●●	●●●●○	●●●○○	Strong technical fit, but often fewer recurring customer workflows
Financial Services	●●○○○	●●●●●	●●●○○	High willingness to pay, but weaker evidence that quantum is uniquely suited to the underlying problems
Logistics & Optimisation	●○○○○	●●●○○	●●○○○	Clear customer need, but intense competition from classical and AI methods
Cybersecurity	●●●○○	●●●●●	●●●●○	Large strategic value, but application timing and independent platform ownership remain uncertain

Methodology Note: ratings are our directional judgements as of September 2026 on a five-point scale. Technical fit assesses whether commercially relevant problems map to quantum structures for which credible algorithms and resource studies exist. Economic impact assesses the value, frequency and budget ownership of the decision being improved. Platform defensibility assesses the potential to accumulate proprietary algorithms, data, validation, workflow integration and repeat customer use. The investment view is a qualitative assessment.

That combination is the point (Table 2). No vertical dominates every axis: materials, energy and chemicals rival life science on technical fit, and financial services can match it on willingness to pay. What distinguishes life science is that strong quantum relevance, high decision value, experimental validation, proprietary learning and persistent workflows come together in one place. It is not first on every measure. It is the vertical where the specific combination that makes a position ownable and worth paying for holds at once, which is the combination that matters for an investor.

For an investor looking at a company that claims a quantum opportunity, the first question is what is actually being solved. The investable unit is the problem, not the sector or the technology: a company is not investable because the technology is novel, the sector is large, or it has a "Q" in its name. A credible opportunity starts with a specific problem and clears four tests:

A genuine classical failure point:

» Current methods cannot deliver the accuracy, cost or robustness the decision requires, which is mostly true for strongly correlated systems (e.g. OLED materials, battery chemistries, catalysts and metallo-organic drugs, including some oncology candidates).

» The opportunity is not to improve on a weak baseline, but to reach where existing methods break down.

Decision leverage: A better calculation must change a real decision: which molecule advances, which experiment runs, which process is selected, where capital is allocated.

A workload a real machine can run: The estimate must cover the full resource envelope (e.g. number of gates, depth, error-correction overhead, cycle time, state prep etc), not just the number of logical-qubits, while still delivering the precision the decision needs.

Commercial ownership: The company must build reusable algorithms, validation, data generation pipelines and deep workflow integration, rather than delivering a series of isolated one-off results.

The first two tests are the ones most often skipped, and both must hold: an accurate quantum result that changes no decision is not investable, however impressive the calculation. Clearing the classical-failure test without the decision-leverage test is the most common way a plausible-looking company fails.

TABLE 3

Life-science problem classes assessed against quantum fit, economic impact and the earliest scenario window in which a credible role exists. Ratings are our directional judgements as of September 2026; windows refer to the scenarios in Table 1.

Problem Class	Quantum Fit	Economic Impact [29,30]	Earliest Plausible Window	Earliest Credible Role
Strongly correlated electronic structure and active-site chemistry	●●●●● [11,12,13,14]	●●●●○	Window 1	High-accuracy treatment of selected regions within a broader classical workflow. [15,16,17]
Reaction mechanisms, bond breaking and catalytic barriers	●●●●○ [13,18]	●●●●○	Window 1	Targeted treatment of local reaction centers where classical approximations break down [18,19]
Excited states, photochemistry and spectroscopy	●●●●○ [20]	●●●◐○	Window 1	Calculations tied to measurable properties such as spectra, transition energies and photostability. [20]
Protein–ligand binding and free-energy prediction	●●●○○ [21,22,23,24]	●●●●○	Window 1 for local corrections; Window 2 for the full workflow	Improved treatment of local electronic effects within broader binding and free-energy workflows. [21,22]
Molecular dynamics and conformational sampling	●●○○○ [23,25]	●●●●○	Window 2	Hybrid treatment of selected quantum-mechanical regions within larger molecular simulations. [15,25]
Molecular generation, docking and screening	●●○○○ [23,26,27,11]	●●●○○	Not a near-term target	Selected subroutines only, which must outperform rapidly improving classical and AI methods [23,11]
ADMET, omics and clinical optimisation	●○○○○ [23,28]	●●○○○	Not a near-term target	Indirect or optimisation-based roles rather than a core quantum-computing application [23,28]

The strongest problems are specific bottlenecks, not entire sectors. The broad landscape of life-science use cases, across quantum computing, simulation and sensing, is mapped in the companion Nordic Quantum Life Science white paper [32]. Applying the four tests to that landscape narrows it sharply. The strongest near-term problems are not categories like "drug discovery" or "molecular simulation" but specific computational bottlenecks inside them: active-site electronic structure, reaction barriers and selected excited-state calculations, where strong correlation makes classical methods structurally, not just temporarily, limited (Table 3). The common thread is specificity. Early fault-tolerant quantum computers will not model a biological process end to end; they will address the small, decisive part of a workflow where classical approximations break down, and the investable company is the one that identifies that bottleneck, formulates it, reduces the resource requirement, and integrates the result so it changes a decision. Protein-ligand binding shows both the opportunity and its limit: it is economically important, but binding depends on conformational sampling, solvent, entropy and kinetics, not just electronic structure, so the credible early role is a high-accuracy correction inside the wider binding workflow, not a replacement for it.

The defensible asset is the workflow, not the algorithm. This carries a direct consequence for ownership, and it is the heart of what an investor should underwrite. The near-term problems where quantum is most justified are also where classical and machine-learning methods are improving fastest. Where classical accuracy is fundamentally constrained (strong correlation, near-degeneracies, certain excited states), quantum has a durable role. Where classical methods are merely catching up, the window can close before application-relevant

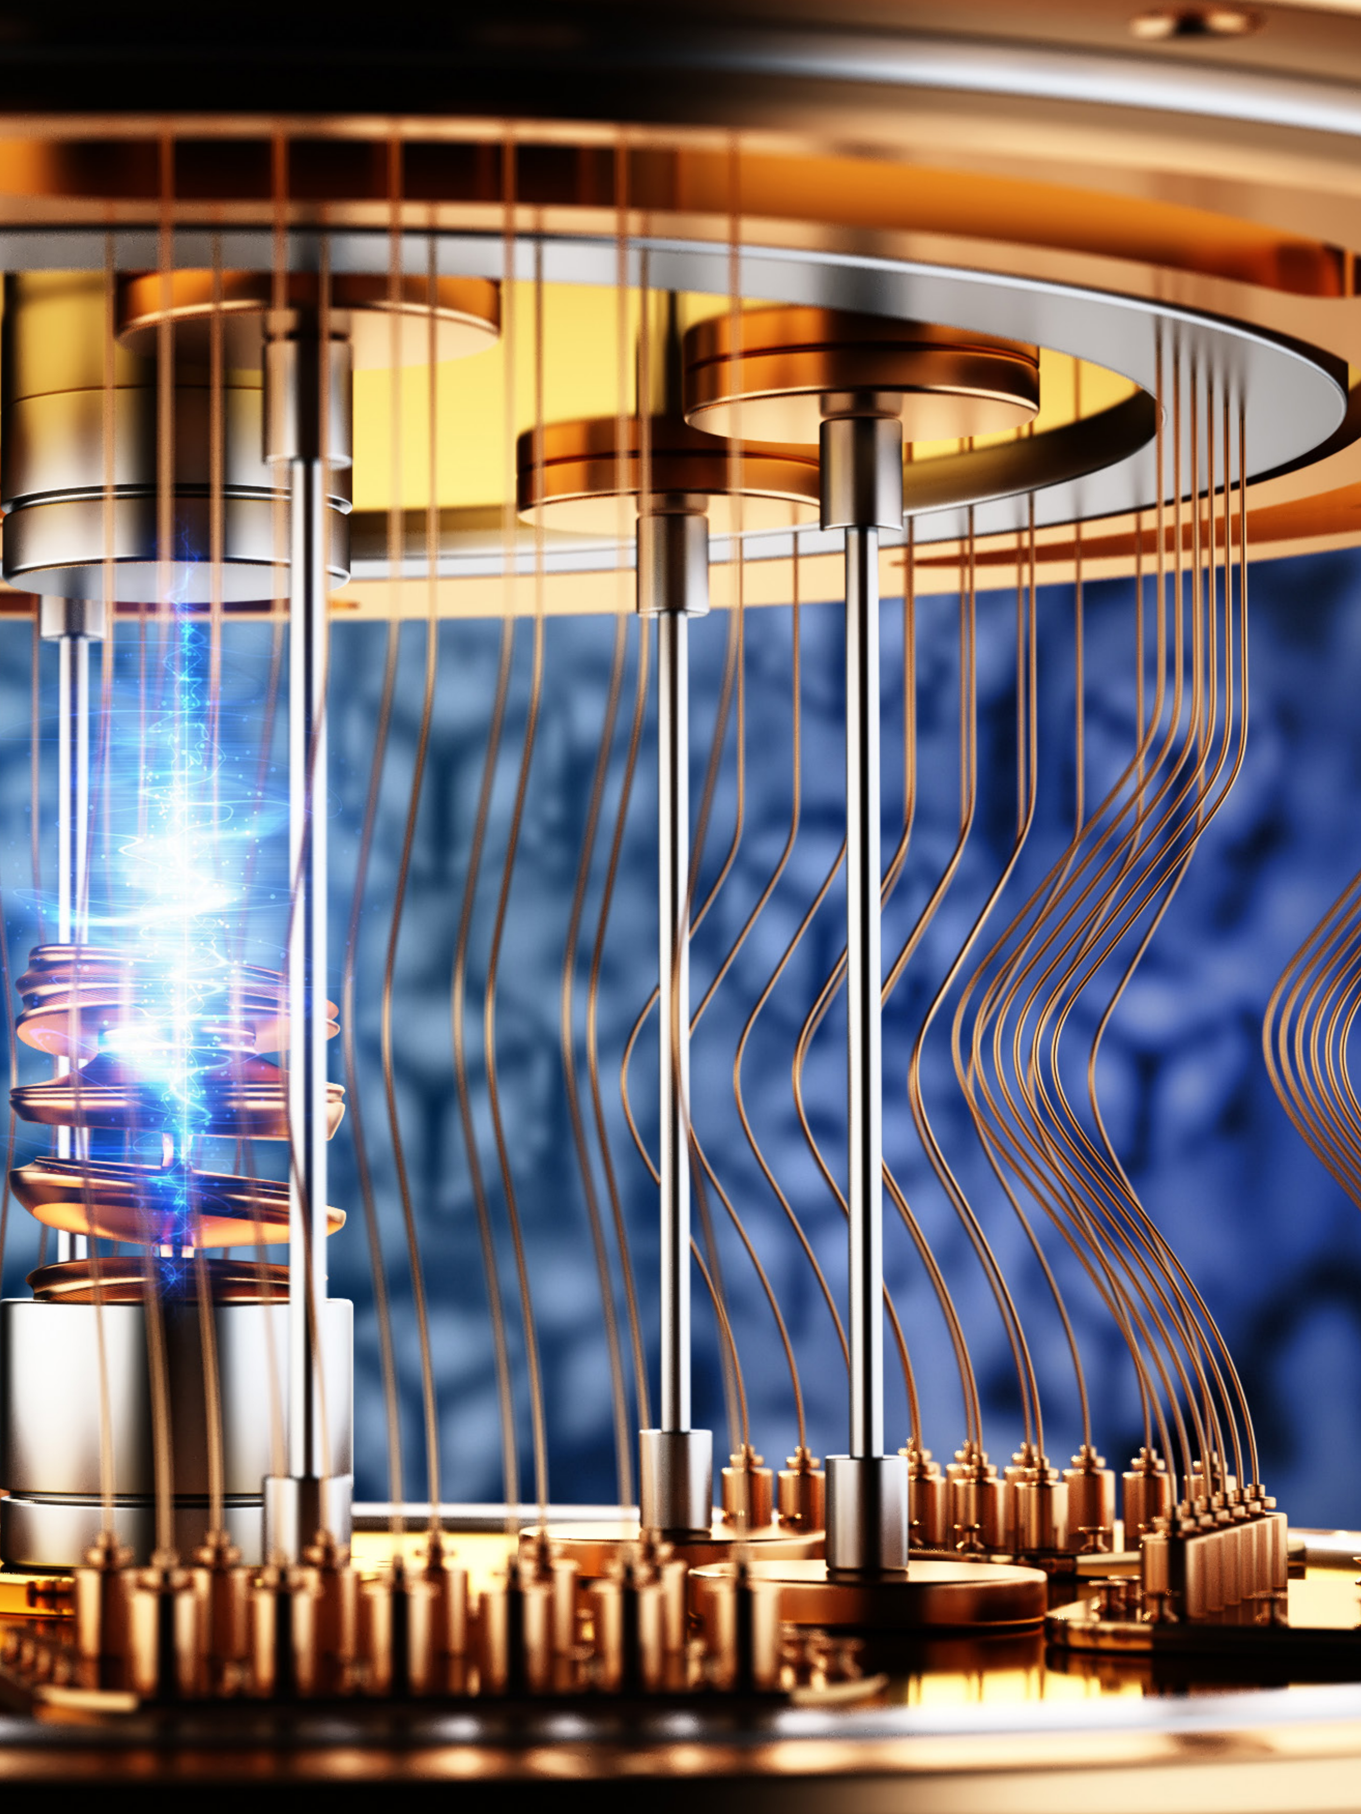

fault-tolerant hardware becomes available. So, if the realistic near-term contribution is a high-accuracy reference calculation feeding a larger classical workflow, the defensible asset is not the algorithm, which publication reproduces and which classical methods may match, but the proprietary validation data and workflow integration built around it. That distinction, not a speed-up claim alone, is what converts an early engagement into an owned position, on the same compounding logic set out in Section 4.

This conviction is conditional, and it is worth stating what would break it:

- If classical methods close the accuracy/cost gap before application-relevant fault-tolerant systems become available.
- If credible resource estimates place the useful workloads in scaled rather than early fault-tolerant regimes.
- If better calculations do not, in practice, change experiments or portfolio decisions.
- If application companies cannot retain reusable IP and validation data.
- If hardware, cloud or pharmaceutical incumbents internalise the capability rather than buy it from independents, a path some are already on, assembling integrated pipelines that combine classical, AI and quantum methods.

The strongest external challenge is not to the workflow argument but to the premise beneath it, whether the exponential quantum advantage in ground-state chemistry holds at all, which remains genuinely contested in the peer-reviewed literature. We treat that as a live risk to underwrite against, not a settled question.

These risks are real. They are also the reason the case rests on the workflow, not the algorithm. A company whose edge is a speed-up loses it the moment classical methods catch up. A company that owns the validated workflow, the proprietary data and the customer decision holds a defensible position even then, because what it owns is not the calculation but the trusted means of turning that calculation into a decision. What survives these risks is ownership of the workflow, the data and the decision, and that is what an investor should back.

6 Conclusion:

Where value will accumulate

We began by asking where value will accumulate in the quantum stack. The answer is that it will not accumulate by layer. It accumulates at specific control points: capabilities that are hard to reproduce, that customers build into their decisions, and that a company can own. Stack position does not predict value capture. Control does, and which control points are open depends on where the ecosystem is in its development journey.

Today, the application layer is the position that is scarce, ownable, capital-light and still open. Life science is the case that shows the framework in full: not because quantum will work there first, but because a better molecular decision is highly valuable, is something customers already pay to improve, and is delivered through a workflow that can be owned. In practice the company to back looks less like a licensor of algorithms and more like one that sells a validated scientific outcome today and compounds the workflow, data and customer relationships underneath it, so that its position deepens as the hardware matures. The investable opportunity is the specific bottleneck, not the sector, and the defensible asset is the validated workflow and proprietary data built around the calculation, not the algorithm itself.

The implication is the part that cannot wait. As the industry approaches application-relevant fault-tolerant regimes, the window to take an application-layer control point starts to narrow, as workflows emerge and embed, and the data compounds. To date, and rightfully so, capital has concentrated upstream into hardware, leaving the layer that converts constrained quantum resources into customer decisions underfunded relative to the potential value it will create. The value at stake for industries adopting quantum far exceeds the revenue that will accrue to the technology providers, and we read that asymmetry as directional support that the prize sits with those who convert the capability into applied outcomes, not with those who supply the hardware.

The next phase of quantum investment must finance application formation alongside continued hardware progress, backing companies that own the problem, the workflow, the validation and the data before application-relevant fault tolerance is established, because the position is built in the interval, not after it.

A strong ecosystem makes that easier, and the Nordics, with deep quantum and life-science strengths side by side, are well placed to build these companies.

Technology creates capability.
Dependency creates value.
Ownership determines
who captures it.

References

1. Scannell, J. W., Blanckley, A., Boldon, H., & Warrington, B. (2012). Diagnosing the decline inpharmaceutical R&D efficiency. Nature Reviews Drug Discovery 11(3), 191-200.
2. Wong, C. H., Siah, K. W., & Lo, A. W. (2019). Estimation of clinical trial success rates and related-parameters. Biostatistics 20(2), 273-286.
3. The Quantum Insider funding database, extracted 8 September 2026. Prior-year figures are restated as rounds are added retrospectively.
4. McKinsey & Company, Quantum Technology Monitor 2026, April 2026, pp. 6 and 47.
5. Calculations based on Synopsys FY2025 reported revenue categories
6. Calculations based on ASE reported financials for Q4 and FY 2025.
7. Company annual reports, regulatory filings, investor materials and official press releases; Novo Holdings Quantum Team analysis and calculations. Metrics reflect the latest available company-reported data accessed July 2026.
8. Based on published company roadmaps. Timing and specifications are targets, not independent forecasts.
9. Primary company roadmaps and technical announcements accessed September 2026. Roadmap metrics differ across architectures and include physical qubits, logical qubits, algorithmic qubits and operation counts; they are not directly comparable.
10. SoftBank Corp. & Quantinuum (2026). Quantum Computing Frontiers: A Use-Case Timeline for Quantum Chemistry and Topological Data Analysis. White Paper, July 2026, especially pp. 4, 18–24 and 43–45.
11. Lee, S., Lee, J., Zhai, H., et al. (2023). Evaluating the evidence for exponential quantum advantage in ground-state quantum chemistry. Nature Communications 14, 1952.
12. Izsák, R., Ivanov, A. V., Blunt, N. S., Holzmann, N., & Neese, F. (2023). Measuring electron correlation: the impact of symmetry and orbital transformations. JCTC 19(10), 2703-2720.
13. Reiher, M., Wiebe, N., Svore, K. M., Wecker, D., & Troyer, M. (2017). Elucidating reaction mechanisms on quantum computers. PNAS 114(29), 7555-7560.
14. Blunt, N. S., Camps, J., Crawford, O., et al. (2022). Perspective on the current state-of-the-art of quantum computing for drug discovery applications. JCTC 18(12), 7001-7023.
15. Sun, Q., & Chan, G. K.-L. (2016). Quantum embedding theories. Accounts of Chemical Research 49(12), 2705-2712.
16. Stein, C. J., & Reiher, M. (2019). autoCAS: a program for fully automated multiconfigurational calculations. J. Comput. Chem. 40(25), 2216-2226.
17. Rossmannek, M., Pavošević, F., Rubio, A., & Tavernelli, I. (2023). Quantum embedding method for the simulation of strongly correlated systems on quantum computers. JPCL 14(14), 3491-3497.
18. von Burg, V., Low, G. H., Häner, T., et al. (2021). Quantum computing enhanced computational catalysis. Physical Review Research 3(3), 033055.
19. Kovyrshin, A., Manawadu, D., et al. (2026). Approximate quantum circuit compilation for proton-transfer kinetics on quantum processors. Physical Chemistry Chemical Physics.
20. Zhou, Y., Casares, P. A., Dhawan, D., et al. (2025). Quantum algorithms for photoreactivity in cancer-targeted photosensitizers. arXiv:2512.15889.
21. Otten, M., Watts, T. W., Johnson, S. D., et al. (2024). Quantum resources required for binding affinity calculations of amyloid beta. arXiv:2406.18744.
22. Huang, P.-W., Boyd, G., Anselmetti, G.-L. R., et al. (2026). Fullqubit alchemist: quantum algorithm for alchemical free energy calculations. npj Quantum Information.

23. Holzmann, N. (2026). From promise to practice: closing the application gap in quantum computing. Philosophical Transactions of the Royal Society A (accepted). preprint: arXiv:2608.04936
24. Chipot, C. (2023). Free energy methods for the description of molecular processes. Annual Review of Biophysics 52(1), 113-138.
25. Senn, H. M., & Thiel, W. (2009). QM/MM methods for biomolecular systems. Angewandte Chemie Int. Ed. 48(7), 1198-1229.
26. Muratov, E. N., Bajorath, J., et al. (2020). QSAR without borders. Chemical Society Reviews 49(11), 3525-3564.
27. Baig, M. H., Ahmad, K., et al. (2016). Computer aided drug design: success and limitations. Current Pharmaceutical Design 22(5), 572-581.
28. Abbas, A., Ambainis, A., et al. (2024). Challenges and opportunities in quantum optimization. Nature Reviews Physics 6, 718-735.
29. Scannell, J. W., Blanckley, A., Boldon, H., & Warrington, B. (2012). Diagnosing the decline in pharmaceutical R&D efficiency. Nature Reviews Drug Discovery 11(3), 191-200.
30. Wong, C. H., Siah, K. W., & Lo, A. W. (2019). Estimation of clinical trial success rates and related parameters. Biostatistics 20(2), 273-286.
31. Fratus, K. R., Enenkel, N., Zanker, S., Reiner, J.-M., Marthaler, M., & Schmitteckert, P. (2026). Can a quantum computer simulate nuclear magnetic resonance spectra better than a classical one? arXiv:2508.06448.
32. Nordic Quantum Life Science Steering Committee, Nordic Quantum Life Science White Paper, 2026.

About Novo Holdings

Novo Holdings is a holding and investment company that is responsible for managing the assets and the wealth of the Novo Nordisk Foundation, one of the world's largest philanthropic enterprise foundations.

Novo Holdings is the controlling shareholder of Novo Nordisk A/S and Novonesis A/S, and manages an investment portfolio with a long-term return perspective.

In addition to managing a broad portfolio of equities, bonds, real estate, infrastructure and private equity assets, Novo Holdings is a world-leading life sciences investor. Through its Seed, Venture, Principal Investments, Planetary Health Investments and Asia teams, Novo Holdings invests in life science companies at all stages of development.

As of year-end 2025, Novo Holdings had Total Assets under Management of DKK 694 billion (EUR 93 billion).

novo **holdings**