

The QUBT Quandry: How Did This Company Reach a \$2.5B Market Cap?

Beniamino Green
September 10, 2025

Disclaimer: At the time of publishing (09/10/25) I currently hold a short position in QUBT, but I may change this position at any time without notice. This report is for informational purposes only and represents my good-faith interpretation of publicly available information. I believe the information presented here to be accurate at the time of writing, but it may contain errors or omissions. It is provided “as is” without any warranties of any kind, express or implied, including but not limited to accuracy, completeness, or fitness for any purpose. Nothing in this document should be construed as financial advice.

Quantum Computing Inc. ([NASDAQ:QUBT](#)) is a pure-play quantum computing company sporting an impressive \$2.5B market cap that has been driven up by recent investor enthusiasm for quantum stocks. While the firm brands itself as solving real-world optimization challenges using their proprietary quantum computing devices, they have generated almost no revenue despite partnering with several blue-chip customers. This is likely because the company’s quantum offerings do not speak to real business needs, and are easily replicated on classical computers: analysis of one of the company’s technical papers, which describes a machine-learning application of their quantum computer shows that in at least one case, their ‘quantum machine learning approach’ is nothing more than a simple regression that combines predictions from other machine-learning algorithms trained on a classical computer. Other promotional materials published on the company website and social media pages appear designed to mislead the company’s investors while carefully avoiding outright falsehoods. The combination of these factors leads me to believe the company is greatly exaggerating both the present capabilities and the future potential of its core business, and is significantly overvalued.

How QCI’s Quantum Computing Works

QCI occasionally publishes papers that aim to show how their quantum computer can address real business problems. A close examination of these papers shows that in many cases the quantum computer is only involved in performing a tiny step of the computation which can easily be replicated with a classical computer in less time.

This is best exemplified by a recent white paper put out by the company, titled [Financial Fraud Detection With Entropy Computing](#). The paper introduces an algorithm developed by QCI called CVQBoost, which makes use of QCI’s Dirac-3 Quantum computer. The paper favorably compares CVQBoost to XGBoost (a powerful classical alternative) finding that while XGBoost outperforms CVQBoost in the most difficult settings, CVQBoost comes close in performance for easier tasks, while being faster to run. The findings are taken as evidence that QCI’s quantum approach can be used as part of a machine-learning pipeline.

The paper does not highlight that all six of the authors of either own or used to own QCI stock, according to an [S-8 form filed with the SEC last year](#). Nicholas Chancellor, the author who owned the fewest shares, appears to have owned at 39,238 shares in March of 2025, worth over half a million dollars at today’s stock price.

What might have escaped most readers is that almost none of CVQBoost runs on a quantum computer, and the part that does could be run on a classical computer in a tiny fraction of the time. CVQBoost is a boosting algorithm, which means that it combines several other algorithms (‘weak learners’) together in order to create a stronger prediction. These weak learners are all trained on a classical computer, and then combined using a variant of linear regression, which is run on the quantum computer. The only quantum step is combining predictions using ridge regression — a standard statistical tool. This is trivial for classical computers, and in fact, if some of the problem constraints are eliminated, there is an exact closed-form solution that can be calculated in milliseconds.

(a) CVQBoost objective function, taken from QCI white paper

$$\min_{\mathbf{w}} \sum_{s=1}^S \left| \sum_{i=1}^N w_i h_i(\mathbf{x}_s) - y_s \right|^2 + \lambda \sum_{i=1}^N (w_i)^2$$

(b) Ridge regression objective function, from *Elements of Statistical Learning*, p. 63

$$\operatorname{argmin}_{\beta} \left\{ \sum_{i=1}^N (y_i - \beta_0 - \sum_{j=1}^p x_{ij} \beta_j)^2 + \lambda \sum_{j=1}^p \beta_j^2 \right\}.$$

Figure 1: The objective functions for CVQBoost (left) and Ridge regression (Right). The two expressions are mathematically identical, apart from the fact that ridge regression estimates an extra intercept parameter.

In short, QCI's "quantum" contribution to this algorithm is nothing short of homeopathic. The vast majority of the machine-learning workload is performed classically, and even the quantum aggregation step could be executed in milliseconds with a few lines of code on a classical computer. It seems likely that this quantum step exists mainly so the company can claim the algorithm is "quantum" in a move just as ridiculous as a plane manufacturer advertising their plane as "electric" because an electric starter ignites its jet engines.

Limited Adoption of Core Quantum Services

The fact that QCI's quantum computers perform worse than classical computers likely makes it very difficult to pitch them to potential clients. This may help explain why, as other analysts have noticed, their quantum compute as a service offering still generates next to no revenue according to their most recent [SEC filings](#), which showed that the company only generated 61K in revenue in Q2 2025. As of March 2025, [QCI charged \\$1,000 an hour for use of their quantum compute service](#), which suggests that they have only been able to charge users for less than half a week's worth of compute in a quarter.

If anything, I think that the revenue reported by QCI might overstate the promise of the business. One way to track the number of users of QCI's quantum compute service is to check the number of downloads for the software packages users use to interact with QCI's quantum computers. Because these software packages are distributed on public repository, the download statistics can be freely queried, and the results are shocking.




Figure 2: Downloads for 'qci-client', the python package that allows users to interface with QCI's flagship Dirac-3 Quantum computer. On most days, the package sees fewer than 5 downloads. Screenshot taken 08/11/25.

If we view the data on [pypistats.org](#), we can see that on some days, the only downloads for the qci-client package are by 'mirrors,' which are servers that automatically back-up all packages. Given that some of the remaining downloads are likely from QCI employees, this data strongly suggests that QCI's core quantum computer offering has either no or next to no users. Given that customers who contact QCI for a ten-minute trial of the quantum computer also use this package, it suggests that QCI has significant difficulty even generating sales leads.

While QCI papers are capable of fooling investors who do not understand machine learning, they cannot convince domain experts to fork over \$1000 an hour to perform easily-replicated calculations at a snail's

pace. I speculate that the \$61K in revenue generated from their quantum compute service reflects companies that purchase a small allocation of compute credits to explore the technology and decline to purchase more once they realize it is significantly less capable and more expensive than classical approaches.

Tempe, AZ chip foundry may not exist

These technical shortcomings are compounded by questions about QCI's physical operations—specifically its much-touted chip foundry in Tempe, Arizona. Outside of QCI's press releases, it's very difficult to find any evidence that this chip foundry exists. The activist short-selling firm Iceberg Research [sent an investigator to the address of the chip foundry](#), and found that while QCI appeared to be a tenant, the building did not appear suitable for manufacturing, lacking basic amenities such as a loading dock or space for large industrial equipment.

This makes sense when you realize that QUBT is literally across the street from Arizona State University's [MacroTechnology Lab](#), which is a shared R&D and prototyping hub, and has the type of clean-rooms necessary to create advanced photonic devices. In fact, ASU's core research facility [website even has a page](#) that advertises their capability and expertise in photonics development. It appears very likely that QCI is not developing a manufacturing center in their rented offices, but are using the ASU lab to develop and fabricate initial versions of their TFLN wafers.

This is corroborated by the fact that QCI does not appear to have significant headcount in Arizona based on their employees listed on LinkedIn, and that the company does not seem to recruit employees in the Tempe Arizona area: I can only find record of a single job posting made by the company in Tempe, was for a director of manufacturing operations position.

The company often posts updates on their LinkedIn page about the progress of the foundry, in ways that seem designed to imply that the foundry is further along than it is while carefully avoiding outright lies to investors. QCI posted the following set of images of thin film deposition equipment, which purportedly show the beginning of 'the completion of our Foundry in Tempe.'

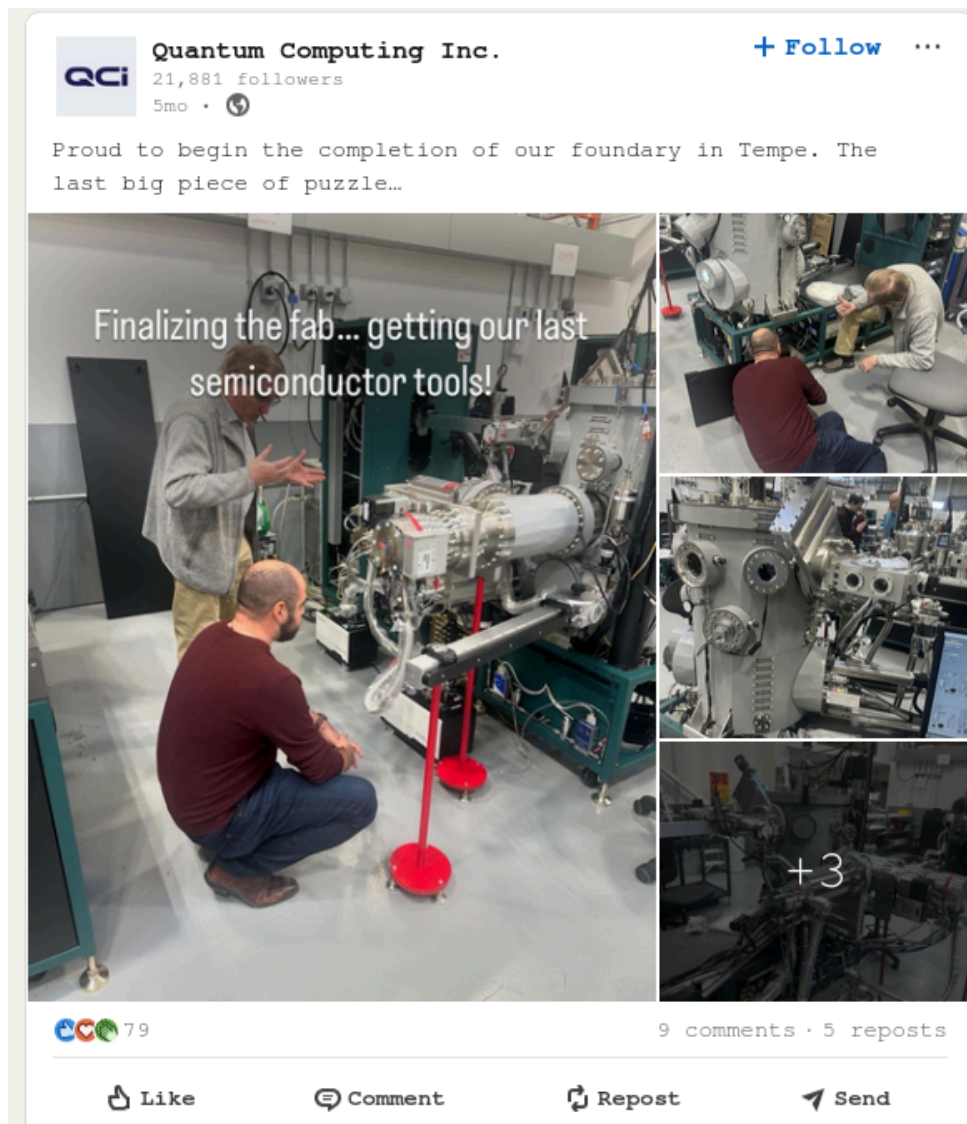


Figure 3: Instagram post of progress of QCI's foundry. These photos were actually taken at AJA international in Massachusetts

A diligent examination of the background and some online sleuthing reveals that these are actually photos from the premises of AJA international¹, the makers of the equipment in the picture. QCI executives appear to have made a trip to AJA to look at or purchase some equipment, and then posted online to make it look as if the equipment is installed in their Arizona facility. The fact that the company appears to fake even basic milestones in the development of its foundry such as installing a machine, suggests that they may be either behind schedule or incapable of finishing the foundry as described to investors.

Conclusion

QCI's market capitalization appears to have been driven largely by investor enthusiasm for quantum computing, rather than by demonstrable commercial traction or technical breakthroughs. While early-stage technology companies often generate limited revenue, the evidence suggests that QCI's quantum systems are inferior to classical alternatives, and adoption of its core quantum compute service is near zero.

Absent a quantum technology that meaningfully outperforms classical alternatives, QCI is unlikely to ever generate significant revenue. While it is possible that the Dirac-3 systems and Tempe foundry represent long-term investments that could pay off in the future, current evidence suggests that the market may be overestimating this probability. Independent research highlights the poor adoption of QCI's core quantum

¹The same piece of equipment can be seen in several of AJA's LinkedIn posts, including [this one](#).

compute service, likely reflecting the product's technical shortcomings. Moreover, the Tempe, Arizona foundry may be significantly further behind than investors realize.

Investor excitement may temporarily inflate QCI's valuation, but long-term sustainability ultimately depends on commercial traction. Given the substantial gap between the company's marketed capabilities and observable performance, the market will likely correct this disparity, bringing the valuation more in line with actual business fundamentals.