

QUANTUM · FINANCE · TECHNOLOGY

A PORTRAIT OF  
**QUANTUM  
TECHNOLOGIES**  
**in Finance**

---

Edited by  
**Oswaldo Zapata, PhD**

---

THE QUANTUM FINANCE BOARDROOM



# A Portrait of Quantum Technologies in Finance

Edited by Oswaldo Zapata, PhD

The Quantum Finance Boardroom

May 2026



# Contents

<b>Foreword</b>	<b>13</b>
<b>Contributors</b>	<b>15</b>
<b>1 Quantum Ethics – Adrian Maguire</b>	<b>25</b>
1.1 A Childhood in Superposition . . . . .	26
1.2 From Y2K to Q-Day . . . . .	27
1.3 The Responsible Quantum Industry Forum . . . . .	27
1.4 Quantum Finance in Superposition . . . . .	29
1.5 Banking Systems and Financial Literacy . . . . .	29
1.6 Imagination, Ideation, Innovation . . . . .	30
1.7 Living in Nottingham, Thinking in Superposition . . . . .	30
1.8 The Subtle Promotion . . . . .	31
<b>2 Combinatorial Portfolio Optimization – Alejandro Rodriguez Dominguez</b>	<b>33</b>
2.1 Introduction . . . . .	36
2.2 Quantum-Inspired Algorithms and Combinatorial Portfolio Optimization . . . . .	38
2.3 Application to Active and Passive Portfolio Management . . . . .	42
2.4 Empirical Evaluation: Computational Accuracy, Statistical Generalization, and Investment Performance . . . . .	44
2.5 Conclusion . . . . .	48
<b>3 Building the Quantum Workforce – Amir Rassol</b>	<b>53</b>
3.1 A Career at the Frontier of Technology . . . . .	53
3.2 Data Science, AI, and Another Frontier . . . . .	54
3.3 Entering Quantum . . . . .	54
3.4 From Feasibility to Scale . . . . .	54
3.5 The Physics Trap and the Translator Gap . . . . .	55
3.6 Quantum Advancing Faster Than Organisational Readiness . . . . .	55

3.7	Recruitment Limits and Team-Based Reality . . . . .	55
3.8	Workforce Advisory and Upskilling . . . . .	56
3.9	Finance as an Early Adopter . . . . .	56
3.10	From Poaching to Workforce Strategy . . . . .	56
3.11	Upskilling in Practice . . . . .	56
3.12	My Current Work in the Quantum Ecosystem . . . . .	57
3.13	Education, Storytelling, and the Podcast . . . . .	57
3.14	Making Quantum Legible Without Selling Certainty . . . . .	58
3.15	Open Questions and the Road Ahead . . . . .	58
3.16	Workforce Infrastructure as the Real Constraint . . . . .	59
3.17	Timelines, Direction, and Readiness . . . . .	59
3.18	Conclusion: Recruiting Without a Playbook . . . . .	59
<b>4</b>	<b>The Quantum Finance Boardroom – André Costa</b>	<b>61</b>
4.1	The Most Dangerous Moment: A Flight Plan for Financial Disruption	61
4.2	From Curiosity to Structured Inquiry: The Genesis of a Quest . .	62
4.3	The Birth of TQFB: An Ecosystem for Translational Intelligence .	62
4.4	The Wealth Management Imperative: Beyond Alternative Invest- ments . . . . .	63
4.5	Expanding the Quantum Horizon: Beyond Computing . . . . .	63
4.6	The Complexity Premium: Valuing Knowledge as an Asset . . . . .	64
4.7	The Human Element: People at the Core of the Quantum Revolution	64
<b>5</b>	<b>Quantum Computing in Asset Management – Carlos Arcila Barrera</b>	<b>67</b>
5.1	Treating Quantum Like a Research Program, Not a Trend . . . . .	69
5.2	Sigma Quantum Lab: Specializing Instead of Generalizing . . . . .	72
<b>6</b>	<b>Kernelizing Quantum Finance – Chinonso Onah</b>	<b>75</b>
6.1	Introduction: Ambition Meets Misalignment . . . . .	75
6.2	The Structural Misalignment . . . . .	77
6.3	Conservation Laws and Structured Financial Optimization . . . . .	78
6.4	Kernelization of Quantum Finance . . . . .	79
6.5	The Cost of Ignoring Geometry . . . . .	81
6.6	Conclusion: Toward a Structural Theory of Quantum Finance Optimization . . . . .	82
<b>7</b>	<b>From Technical Correctness to Hybrid Pragmatism – David Isaac</b>	<b>85</b>
7.1	Entering With the Wrong Assumption . . . . .	85

---

7.2	A Concrete Failure and Its Consequences . . . . .	86
7.3	What I Stopped Believing, and What Replaced It . . . . .	87
7.4	The Hidden Tax of Being Early . . . . .	89
7.5	The Tensions That Do Not Resolve Cleanly . . . . .	90
7.6	Closing . . . . .	92
<b>8</b>	<b>A Road Less Traveled – Francisco Castro</b>	<b>93</b>
8.1	The Start of an Unexpected Journey . . . . .	93
8.2	Getting Industrial Experience . . . . .	94
8.3	From Engineer to Lawyer . . . . .	95
8.4	The Start of My Legal Career . . . . .	95
8.5	Serendipity . . . . .	96
8.6	Becoming IonQ’s First In-House Lawyer . . . . .	97
8.7	The Challenges of a Quantum Startup . . . . .	98
8.8	The Harvard Connection and the Origins of Bluzec . . . . .	98
<b>9</b>	<b>Seeing What Isn’t There – Genevieve Hayman</b>	<b>101</b>
9.1	The View from Somewhere . . . . .	101
9.2	The Roots of Computation in Logic . . . . .	102
9.3	The Quantum Rupture . . . . .	103
9.4	Phase Portraits and the Landscape of Possibility . . . . .	103
9.5	Complex Systems and the Comfort of Irreducibility . . . . .	104
9.6	Alpha as Epistemic Advantage . . . . .	105
9.7	An Invitation . . . . .	106
<b>10</b>	<b>A Dance of the Blind Puppeteer – Jacob L. Cybulski</b>	<b>107</b>
10.1	Introduction . . . . .	107
10.2	Strings: PQC Parameters . . . . .	108
10.3	Stage and Puppet: Hilbert Space and Manifold . . . . .	109
10.4	Blind Puppeteer: Classical Optimizer . . . . .	112
10.5	Too Many Strings: The Curse of Dimensionality . . . . .	113
10.6	Resolution: Navigating the Metric Mismatch . . . . .	117
10.7	Conclusions: The Geometric Imperative . . . . .	119
<b>11</b>	<b>A Quantum Journey – Jeremy Green</b>	<b>123</b>
11.1	Learning Quantum: A Pragmatic, Security First Approach . . . . .	125
11.2	Quantum Breaks This Assumption . . . . .	127

11.3	The Purpose of All This is Straightforward . . . . .	130
11.4	My View of the Quantum Future . . . . .	131
11.5	The Security Imperative: What Must Happen Next . . . . .	132
<b>12</b>	<b>Speaking Two Languages – Joe Ghalbouni</b>	<b>137</b>
12.1	From Momentum to Responsibility . . . . .	137
12.2	Why Finance Became the Proving Ground . . . . .	138
12.3	The First Reality Check: A Change of Dialect . . . . .	138
12.4	From Education to Execution . . . . .	139
12.5	When Opportunity and Risk Converge . . . . .	139
12.6	Finance’s Paradox: Pioneering Yet Constrained . . . . .	139
12.7	A Misalignment of Incentives . . . . .	140
12.8	The Hybrid Future Is Not a Compromise . . . . .	140
12.9	Talent, Not Qubits, Is the Bottleneck . . . . .	141
12.10	Governance Will Decide the Winners . . . . .	141
12.11	The Role That Had to Exist . . . . .	141
<b>13</b>	<b>My Roundabout Journey into Quantum – John Galani</b>	<b>143</b>
<b>14</b>	<b>Quantum Is the Next AI Moment, Only Bigger – John Riley III</b>	<b>147</b>
14.1	The AI Déjà Vu Moment . . . . .	147
14.2	Making Quantum Easy to Digest . . . . .	148
14.3	Convergence Imperative: AI, Quantum, Cyber & Blockchain Are Becoming One Strategy . . . . .	149
14.4	Leading Financial Institutions in Tri-Technology Adoption . . . . .	150
14.5	Quantum Isn’t “10 or 20 Years Away” Anymore . . . . .	151
14.6	Quantum Readiness: Move Early, Move Smart, Don’t Get Left Behind . . . . .	155
14.7	Quantum is Here: What Are You Waiting For? . . . . .	156
<b>15</b>	<b>The Quantum in Saudi Arabia – Khulud Almutairi</b>	<b>157</b>
15.1	The Nonlinear Path to Quantum . . . . .	157
15.2	The Quantum Kingdom: Strategy, Policy, and Architecture . . . . .	159
15.3	Building the Quantum Stack: Hardware, Infrastructure, and In- dustrial Partnerships . . . . .	162
15.4	The Human Qubit: Academic Research and Talent Infrastructure . . . . .	164
15.5	The Quantum-Finance Nexus: Opportunity, Urgency, and the Saudi Advantage . . . . .	166

---

15.6 Conclusion: Foundations Extraordinary, Building Just Begun . . . .	167
<b>16 Quantum Is a Public Good – Lionel Martellini</b>	<b>169</b>
16.1 Professional Foundation . . . . .	169
16.2 The Quantum Pivot . . . . .	171
16.3 Identifying the Void . . . . .	173
16.4 The Future Landscape . . . . .	175
<b>17 My Finance Journey Into Quantum – Maylix Brianto</b>	<b>177</b>
17.1 When Curiosity Changes the Direction of a Career . . . . .	177
17.2 My First Lessons in Finance . . . . .	178
17.3 The Beginning of a Global Journey . . . . .	179
17.4 From Washington to Geneva . . . . .	179
17.5 Building Bridges in Wealth Management . . . . .	180
17.6 Discovering the Quantum Frontier . . . . .	180
17.7 Exploring the Quantum Investment Frontier . . . . .	181
17.8 Navigating Quantum Opportunities Across Continents . . . . .	181
17.9 Women, Next-Gen, and the Quantum Ecosystem . . . . .	182
17.10 Investor Guide: Engaging with Quantum Finance with Curiosity, Rigor, and Vision . . . . .	183
17.11 Quantum-Enhanced Alternative Investments: Bridging Frontier Technology and Portfolio Strategy . . . . .	184
17.12 Quantum and Sustainability: Investing in the Future Responsibly	186
17.13 Conclusion: Curiosity, Conviction, and the Future of Investing . .	188
<b>18 Standards for Quantum Technology in Finance – Melissa Hernández</b>	<b>191</b>
18.1 Why Quantum Finance? . . . . .	192
18.2 What Standards Mean in Quantum Finance . . . . .	192
18.3 Quantum Standards Summary . . . . .	193
18.4 Cryptographic Standards Layer . . . . .	193
18.5 The Financial Sector . . . . .	195
18.6 The EU Regulatory Layer . . . . .	198
<b>19 Building Quantum-Safe Finance – Olga Mamlyga</b>	<b>201</b>
19.1 Introduction . . . . .	201
19.2 The Quantum Edge: The Real Disruption . . . . .	206
19.3 Categorizing Quantum Technologies . . . . .	206

19.4 Quantum Communication: Industrial-Grade, Post-Quantum Resilient . . . . .	207
19.5 The Current Cyber Threat Landscape: AI as a Force Multiplier . . . . .	208
19.6 Quantum Escalation and the Future Threat Landscape . . . . .	211
19.7 Industry Guidance and Security Frameworks . . . . .	213
19.8 Conclusion: The Quantum Imperative for Finance . . . . .	215
<b>20 An Accidental Quantum Love Story – Orlagh Neary</b>	<b>217</b>
20.1 My Love at First Sight Moment . . . . .	218
20.2 What AI Taught Us, and What the World of Quantum is Still Learning . . . . .	219
20.3 An Invitation . . . . .	219
20.4 So Where Do You Start? . . . . .	220
<b>21 Where Quantum Computing Meets Optimization – Pascal Halfmann</b>	<b>223</b>
21.1 Introduction . . . . .	223
21.2 Success Story: Quantum Optimization in the Energy Industry . . . . .	226
21.3 Success Story: Quantum Multiobjective Optimization for Portfolio Optimization . . . . .	231
21.4 State of Quantum Computing & and a Look Ahead . . . . .	235
<b>22 From Bits to Qubits – Rafał Pracht</b>	<b>241</b>
22.1 Opening Perspective . . . . .	241
22.2 From Classical Logic to Computational Limits . . . . .	242
22.3 Financial Simulation: The Computational Heart of Modern Banking . . . . .	247
22.4 From Classical to Quantum: A New Paradigm for Estimation . . . . .	251
22.5 Quantum Monte Carlo in Finance: Intuition Without Physics . . . . .	252
22.6 The Quantum Binomial Tree: Loading Financial Markets into Qubits . . . . .	253
22.7 Looking Ahead: Why Quantum Monte Carlo Matters for the Future of Finance . . . . .	256
<b>23 Quantum Finance and Venture Capital – Reinaldo Coelho</b>	<b>259</b>
23.1 Quantum Finance Startups as an Investment Opportunity . . . . .	259
23.2 High-Value Domains in the Quantum Finance Frontier . . . . .	260
23.3 The Startup Landscape in Quantum Finance . . . . .	262
23.4 The Investor Landscape in Quantum Finance . . . . .	265
23.5 The Geography of Quantum Finance Innovation . . . . .	267
23.6 In Conclusion, a Call to Action . . . . .	268

---

<b>24 Post Quantum Cryptography – Santanu Ganguly</b>	<b>269</b>
24.1 Introduction . . . . .	269
24.2 Background . . . . .	269
24.3 So, What Is This PQC Anyway? . . . . .	271
24.4 Is PQC the Holy Grail of Security Against All Quantum Attacks? . . . . .	271
24.5 PQC: Some Technical Overview . . . . .	273
24.6 Why PQC Matters Now . . . . .	274
24.7 NIST PQC Standards and How They Work . . . . .	275
24.8 Engineering Migration Roadmap (3–5+ Year Program) . . . . .	276
24.9 Machine Learning (ML) for PQC Deployment, Monitoring, and Capacity Control . . . . .	279
24.10 QML for PQC: Opportunities and Limits . . . . .	281
24.11 Risks, Pitfalls, and Recommendations . . . . .	281
24.12 Conclusion . . . . .	282
<b>25 Quantum Readiness for Financial Adoption – Sebastian Torres</b>	<b>285</b>
25.1 Origins: The Deterministic Illusion in a Stochastic World . . . . .	286
25.2 The Classical Apex: A Case Study of PRO Systematic CM . . . . .	287
25.3 The Combinatorial Wall: Rethinking Portfolio Construction . . . . .	288
25.4 The Latency of Risk: Accelerating Monte Carlo Simulations . . . . .	289
25.5 Beyond Correlation: Quantum Machine Learning for Hedging . . . . .	290
25.6 The Prerequisite of Data Hygiene . . . . .	290
25.7 The Hybrid Operating Model and the Role of the Translator . . . . .	291
<b>26 Quantum Neural Networks – Sebastian Zajac &amp; Krzysztof Kuba</b>	<b>293</b>
26.1 Introduction . . . . .	293
26.2 Classical Machine Learning . . . . .	295
26.3 Quantum Machine Learning . . . . .	298
26.4 Summary . . . . .	304
<b>27 Musings Between Breakthroughs &amp; Builds – Sierra Clouse</b>	<b>307</b>
27.1 The Pattern of Progress . . . . .	308
27.2 The Quantum Divide: Building for the Real World . . . . .	308
27.3 The Question(s) in Front of Us . . . . .	310
27.4 Make These Your FAQs . . . . .	310
27.5 Bonus Round: Orient, Learn, and Act . . . . .	311

<b>28 Quantum and Security – Tatiana Mitrova</b>	<b>313</b>
28.1 Introduction: When “Quantum” Stopped Being a Future Topic for Me . . . . .	313
28.2 Where Quantum Enters the Security System . . . . .	315
28.3 Quantum Computation: The Trust-Layer Shock . . . . .	315
28.4 Quantum Communication: Strengthening Links Without Pretending Systems Are Perfect . . . . .	317
28.5 Quantum Sensing and Timing: Security Begins with Time and Visibility . . . . .	318
28.6 Governance, Chokepoints, and the Transition Problem: Why “Global” Security Often Means “Who Controls What” . . . . .	320
28.7 Conclusion: A Disciplined Way to Talk About Quantum and Security (Without Hype) . . . . .	321
<b>29 The Quantum Edge in Finance – Tomasz Ćwik</b>	<b>323</b>
29.1 The World Before the Revolution: When Quantum Was “Science Fiction” . . . . .	323
29.2 Roots in Banking: Domain Advantage from Day One . . . . .	324
29.3 The Hackathon That Changed Everything: The Birth of the Team . . . . .	325
29.4 From Corporation to Startup: The Decision to Leap into the Unknown and Building a New Identity . . . . .	326
29.5 Accelerators, Structure, Pace, and Global Perspective of Deep-Tech Construction . . . . .	328
29.6 Deep-Tech Financing, Selectivity, Valuation, and Strategic Courage . . . . .	330
29.7 Advisory Board and Team Building, Intellectual Capital as a Real Competitive Advantage . . . . .	332
29.8 Product Strategy, Hybrid as a Responsible Path to the Future . . . . .	333
29.9 Quantum Machine Learning, Between Hardware Limitations and Scientific Ambition . . . . .	333
29.10 Creative Destruction Lab, Global Validation in the Most Demanding Environment . . . . .	334
29.11 Customer Relationships, Rebuilding Trust in a Post-Hype World . . . . .	335
29.12 Business Model and Long-Term Vision: Quantum Finance as a New Layer of Financial Infrastructure . . . . .	336
<b>30 The Coverage Banker – Tse Loong Chin</b>	<b>339</b>
30.1 The Humble Ledger . . . . .	339
30.2 The Architect of Automation . . . . .	339
30.3 The Banker and the Break . . . . .	340

---

30.4 The Institutionalism and the Spark of Quantum . . . . .	341
30.5 Drawing from Mentors' Paths . . . . .	342
30.6 The Preservation of Curiosity . . . . .	343



# Foreword

For many years, discussions about quantum technologies in finance were shaped largely by the physics of the possible—qubit counts, error correction, and the pursuit of quantum advantage. That phase of exploration was essential. It was also during that period that I first became interested in the potential applications of quantum computing in finance. The excitement surrounding those early developments sparked imagination, mobilized researchers, and laid the scientific foundations on which this field now stands.

Today, however, the conversation is evolving.

Over the past few years—and particularly in the last year—I have noticed a clear shift in the questions coming from finance professionals. The discussion is gradually moving beyond foundational questions such as what a qubit is or how quantum hardware works. Increasingly, practitioners are asking a different and far more consequential question: how might quantum technologies influence the way finance itself operates in concrete and practical ways?

This change in perspective reflects a field that is maturing, as well as a shift in the mindset of finance professionals toward quantum technologies: there is growing trust in the work of the scientific community and in the broader quantum discourse. Yet with that maturation—and with that trust—comes both opportunity and uncertainty. Quantum technologies remain complex, and the surrounding narrative can sometimes oscillate between excessive optimism—the so-called quantum hype—and unwarranted skepticism. For many decision-makers and practitioners, the challenge today is not a lack of interest, but the absence of grounded perspectives on what truly matters.

This collection of insights was curated to bridge that gap.

Rather than presenting a single, overly academic viewpoint, I invited a diverse group of thinkers and practitioners from across The Quantum Finance Boardroom community—individuals who are actively shaping the evolving relationship between quantum technologies and finance. Their perspectives span a wide range of topics, from ethics and geopolitics to investment strategy, talent development, and real-world implementation. Together, these contributions form something more than a technical reference. They offer a “portrait” of a vibrant ecosystem that is still taking shape.

*The “portrait” we have drawn places the human dimension at its center.* Behind every chapter is an individual with their own journey within this evolving landscape—researchers, entrepreneurs, investors, policymakers, and technologists who are actively contributing to the future of the field. Their insights are grounded not only in theory, but in lived experience. This was a deliberate choice. Our aim is for readers to connect not only with the ideas presented here, but also with the trajectory and the human being behind them.

If a particular chapter resonates with you, I encourage you to reach out to its author. One of the great strengths of the quantum community is its openness and collaborative spirit. Progress in this domain will not come from isolated efforts, but from networks of people willing to share ideas, challenge assumptions, and build together.

On a personal level, I would like to extend my sincere gratitude to every contributor—each a member of The Quantum Finance Boardroom community—who helped bring this project to life. Their generosity with time, insight, and perspective has made this project possible. Working with this group has been deeply rewarding, allowing me to explore areas and perspectives that I might not otherwise have encountered.

To conclude, it is my sincere hope that the pages that follow will serve as a useful companion as you navigate this rapidly evolving ecosystem.

Whether you are encountering quantum technologies for the first time or are already working at the frontier, this compendium aims to support the journey from curiosity to informed perspective. The quantum era in finance is still unfolding—but it is no longer a distant prospect. It is a conversation that is beginning to shape the future of our industry.

Welcome to the conversation.

*Oswaldo Zapata, PhD*  
Editor | Cofounder, TQFB

# Contributors

Chapters are organized alphabetically by the authors' first names. This simple choice reflects the intention not to prescribe a single path through the material, but rather to invite exploration.

Each contributor offers a unique perspective on quantum technologies in finance. Some chapters are more technical, others more visionary—but all reflect the experience, curiosity, and insight of the human beings behind them.

Below each contributor's profile, you will find a short quote from their chapter. If a particular idea resonates with you, simply click on the "Name – Title" that follows the quote to explore the full chapter.

Let your curiosity guide you and begin exploring from there.

**Adrian Maguire** is a multidisciplinary thinker and former Head of Industry Governance at Experian UK&I, with over 30 years in financial services, data ethics, and emerging technologies. A member of the Responsible Quantum Industry Forum (convened by the NQCC), he has contributed to ethical frameworks guiding quantum innovation. Drawing from philosophy, economics, and hands-on governance, Ade explores how quantum technologies, like superposition itself, embody both promise and peril. He champions responsible development that helps humans flourish, not merely maximise utility.

*"Ethics is the one that keeps you awake at night when the technology you're building might help millions, harm them, or, quantum-style, do both at once."*

[Adrian Maguire – Quantum Ethics](#)

**Alejandro Rodriguez Dominguez** is Head of Quantitative Analysis at Miralbank. Since 2018, he has led development of AI and ML solutions across the institution. His research focuses on ML for portfolio diversification, causality, and risk management. He is a Quant Advisor at Inspiration-Q, contributing to

development of quantum-inspired systematic investment strategies. Previously, he worked in trading and financial engineering positions in London and Paris at Société Générale, Nomura, BBVA, and BNP Paribas. He pursues a PhD in Artificial Intelligence at University of Reading and holds degrees in Financial Engineering (Imperial College London), Computational Statistics, AI and ML, and Mining Engineering.

*“These methods do not rely on quantum hardware; rather, they use heuristics inspired by quantum mechanics [...] to efficiently explore large discrete solution spaces using classical computational resources.”*

### Alejandro Rodriguez Dominguez – Combinatorial Portfolio Optimization

**Amir Rassol** is the founder of QAI Talent, a specialist recruitment and advisory firm focused on quantum technologies and deep tech. With nearly three decades of experience in technology recruitment, he has built teams across data science, AI, and machine learning before pivoting into the emerging quantum ecosystem. Amir works with startups, research organisations, and investors to help design quantum-ready teams and talent strategies. He also hosts the podcast Beyond the Bit: Quantum Pathways, where he explores the people, ideas, and career pathways shaping the future of quantum technologies.

*“One of the biggest fallacies I encountered was what I call the ‘Physics Trap’—the belief that every quantum hire must be a PhD with a deep understanding of qubits.”*

### Amir Rassol – Building the Quantum Workforce

**André Costa**, CAIA, is the founder of Free Enterprise Advisors and co-founder of The Quantum Finance Boardroom. Born on May 26, 1968, in Capinópolis, MG, Brazil, he holds a degree in veterinary medicine from the Universidade Federal de Goiás (Brazil). André is a registered Commodity Trading Advisors with the CFTC and a member of the NFA, as well as a Registered Investment Advisor with the SEC (USA) and CVM (Brazil). He holds various certifications, including CAIA<sup>®</sup>, CFP<sup>®</sup>, and CDAA<sup>®</sup>, reflecting his commitment to excellence in finance. André is also a Certified Energy Risk Professional (ERP<sup>®</sup>), Sustainability & Climate Risk (SCR<sup>™</sup>), and Risk and AI (RAI<sup>™</sup>) certificate holder, all issued by the Global Association of Risk Professionals. Additionally, he is a registered insurance broker with the SUSEP (Brazil) and a Chartered MCSI member with the CISI.

*“Learning about this advanced matter isn’t just an academic exercise; it’s about positioning ourselves at the vanguard of alternative investments for our clients, students, or employers.”*

### André Costa – The Quantum Finance Boardroom

**Carlos Arcila Barrera, CFA, CAIA, SCR**, is the founder of Sigma Advanced Quantum Lab, applying quantum and quantum-inspired methods to asset

management and energy markets. He spent fourteen years in commodity derivatives, including eight years running his own Chicago-based hedge fund managing institutional capital. The computational challenges of asset and risk management led him to quantum computing, completing coursework through UChicago, MIT, and others, and deploying optimization prototypes on D-Wave hardware. He remains an active trader and has lectured at the University of Notre Dame and Universidad de Los Andes. He holds an MSt from the University of Cambridge, MSc in Finance from Notre Dame, and BBA from Universidad de Los Andes.

*“This is where most executives disengage. They hear the hardware is not ready and file quantum computing somewhere between 2030 and never.”*

[Carlos Arcila Barrera – Quantum Computing in Asset Management](#)

**Chinonso Onah** is a theoretical physicist and a quantum algorithm researcher with five years of experience in quantum algorithm research and industrial applications. He studies kernelization of constrained quantum optimization, co-designing encodings, mixers, and phase signals to make constraint geometry explicit, save computational resources and enable provable sampling guarantees and benchmarking on noisy hardware. His work spans quantum optimization algorithms and hybrid QC–HPC workflows, including feasibility/optimalty analyses, finite-shot bounds via polynomial phase filtering, and encoded-kernel architectures for constraint-aware QAOA variants. He also develops scalable benchmarking pipelines and hardware-maturity diagnostics for industrial workloads.

*“The central difficulty in quantum finance to date has not been hardware capability, but problem-algorithm misalignment.”*

[Chinonso Onah – Kernelizing Quantum Finance](#)

**David Isaac** is Co-Founder of Abaqus Computing, where he focuses on applying quantum and hybrid optimization methods to complex financial and operational problems. He has over 20 years of experience in the technology sector across sales, strategy, and business development. His work centers on translating real-world constraints—such as portfolio risk limits and high-dimensional machine-learning systems—into optimization models designed for emerging quantum and hybrid computing platforms.

*“What proved hardest to let go of was not any specific tool or technique, but a set of beliefs that were technically defensible yet economically irrelevant: that correctness implies value, that purity implies superiority, and that adoption follows capability.”*

[David Isaac – From Technical Correctness to Hybrid Pragmatism](#)

**Francisco Castro, PhD, JD**, is currently the Co-Founder of Bluzec, a deep tech advisory company with US and European operations. Prior to that, Dr. Castro was the Associate General Counsel at quantum computing startup IonQ. Dr.

Castro received his Ph.D. from Drexel University in applied semiconductor physics. He was a research engineer first at Lucent Technologies (now Nokia) and then at Motorola before attending Chicago-Kent College of Law. He began his legal career at Silicon Valley-based law firm Cooley LLP. Prior to joining IonQ, Dr. Castro was Counsel at Arent Fox LLP in Washington, DC.

*“After five furious and relentless years at IonQ, I decided to take some time off and figure out what my next step would be.”*

[Francisco Castro – A Road Less Traveled](#)

**Genevieve Hayman, PhD**, is Senior Manager, Macrosystems & Foresight, with expertise in complex systems, cognitive science, and pensions and retirement security. She earned a PhD in philosophy of science from Georgetown University and a master’s degree in economics from George Mason University. Her work has appeared in peer-reviewed journals, and she has recently authored reports on the integration of artificial intelligence within pension systems and how financial markets can be reframed through a complex-systems lens.

*“For those in finance, I would offer this: the quantum era is not just about adopting new tools. It is about adopting new ways of seeing.”*

[Genevieve Hayman – Seeing What Isn’t There](#)

**Jacob L. Cybulski** is the founder of Enquanted, providing research, training and consulting services in quantum computing. Enquanted projects involve quantum machine learning and its applications in business, engineering, and science. Jacob is also Honorary Associate Professor in Quantum Computing, School of IT, Deakin University, Australia. Recently, Jacob researched development of complex quantum models and issues of their trainability, quantum time series and signal analysis, as well as aspects of quantum information field theory. Jacob’s past work also concerned classical machine learning, immersive 3D data visualization and business analytics, at a number of universities and research organizations in Australia.

*“I was among them painting and drawing and dreaming of doing this for the rest of my life.”*

[Jacob L. Cybulski – A Dance of the Blind Puppeteer](#)

**Jeremy Green** is a security professional with over 20 certifications, including CISSP, CISM, CEH, ECDE and CHFI. An official ISACA and EC-Council instructor, Jeremy has authored BCS Information Security Management Principles and Security Architecture books. His career spans education, consultancy, telecoms, police and military, initially teaching Computer Science in further education before moving into police cybercrime and OSINT training. He later delivered certification training across Europe, then held cyber defence and advisory roles at BT and Vodafone. Jeremy is now a Security Architect and serves in the RAF Cyber Reserves. He is currently undertaking a PhD in Computer Science researching quantum security.

*“Data with long term value was already being collected, archived and stockpiled by actors who understood exactly what was coming.”*

[Jeremy Green – A Quantum Journey](#)

**Joe Ghalbouni, PhD**, holds a PhD in quantum communications. After graduating, Joe started his career as an associate professor of Physics at the Lebanese University where he started a quantum computing research activity. In 2021 Joe joined Point72 Asset Management to lead their Quantum initiative and their Core Research. In September 2025, Joe launched his consulting firm where he helps in the successful adoption of quantum, AI and HPC technologies.

*“Introducing it as a radical replacement rather than an incremental enhancement triggers resistance—not fear of complexity, but fear of uncontrolled consequences.”*

[Joe Ghalbouni – Speaking Two Languages](#)

**John Galani** has 25 years of senior leadership across shipping, trading, finance, and technology on four continents, with expertise spanning assets, fintech, AI, and quantum. His roles include COO at Nasdaq-listed fintech Triterras Inc, with its \$580 million SPAC spin-off, and CCO/COO at PE-backed Incomlend, where he drove restructuring and operational streamlining. Earlier in his career he built and successfully exited shipping divisions in 2006 and 2021. His recent AI and Quantum Computing studies focus on investments and applications. John holds an MSc in Shipping, Trade & Finance from Bayes Business School and brings strong analytical capabilities relevant to academic discourse at the intersection of trade, finance, and technology.

*“Whilst I felt I had missed the boat on the most successful fintech companies and would likely miss the boat on AI, I felt I was early in quantum. My heart was there, my mind was set, and I wanted to make it count.”*

[John Galani – My Roundabout Journey into Quantum](#)

**John Riley III** is Co-Founder and Chief of Emerging Tech at IMPACTIFI, with over 25 years of experience in the software consulting services sector, including roles at global software companies like Oracle and SAP. John has honed his expertise in driving user adoption and implementing emerging technologies such as Blockchain, AI, Digital Twins, and cybersecurity preparedness to facilitate digital transformation initiatives. His strategic insights and hands-on approach have empowered organizations to embrace cutting-edge solutions and thrive in today’s digital landscape. His current focus centers on workforce development and preparing students and professionals for the convergence of emerging technologies like Quantum. His dedication to fostering the technology and innovation ecosystem is matched only by his commitment to service, having served honorably as a U.S. Marine War Veteran.

*“The key is awareness and foundational education early. Starting sooner helps us avoid the mistakes of the past, when new tech was forced on people before they were prepared.”*

### John Riley III – Quantum Is the Next AI Moment, Only Bigger

**Khulud Almutairi** is an accomplished quantum researcher and educator at King Abdulaziz City for Science and Technology (KACST), where she contributes to national projects within the Quantum Technologies and Advanced Computing Institute and helps organize quantum events. She holds an MSc in Physics from the University of Calgary’s Institute for Quantum Information Science and a BSc from King Saud University, complemented by Oxford’s Executive Leadership Programme. Previously a Lecturer in Physics at King Saud University, she founded QSaudi Arabia under QWorld to lead workshops in quantum computing. Her work spans quantum information research, education, and ecosystem development in Saudi Arabia.

*“But the critical question, whether quantum will be built in Saudi Arabia, by Saudi people, for Saudi systems, on Saudi terms, remains open.”*

### Khulud Almutairi – The Quantum in Saudi Arabia

**Lionel Martellini** is the founding director of the EDHEC Quantum Institute, research director at the CFA Institute Research Foundation, and former director of the EDHEC Risk Institute. Before joining EDHEC, Professor Martellini was a faculty member at the University of Southern California and held visiting positions at Princeton University and the Massachusetts Institute of Technology (MIT). He holds a PhD in finance from the Haas School of Business, University of California at Berkeley. Outside his work in finance, he earned a PhD in Relativistic Astrophysics and currently conducts research on the foundations of quantum mechanics and applications in quantum technologies.

*“Beyond these technical considerations, there is a deeper issue: the lack of quantum literacy among decision-makers. This is where I believe business schools have a critical role to play.”*

### Lionel Martellini – Quantum Is a Public Good

**Maylix Brianto, CAIA, CESGA, MBA, MCSI**, has nearly two decades of experience in investment management. She works across Portfolio Management and Multi-Family Office services at a leading wealth manager in Geneva, and previously held roles at Banque SYZ, Société Générale, J.P. Morgan, and the Inter-American Development Bank. An Economist from UCV with a Master in Finance from IESA and an MBA from Grenoble Ecole de Management, she is a CAIA Chapter Executive, a part-time professor at IFM and GEM Business Schools, and a financial writer.

*“Ultimately, investing in quantum technologies is about more than returns. It is about participating in a frontier that reshapes how we think, how markets operate, and how possibilities are measured.”*

### Maylix Brianto – My Finance Journey Into Quantum

**Melissa Hernández** is a PhD candidate at the Institute of Private Law of Leiden University. Her research focuses on Standards for Quantum Applications in Finance. She obtained a JD at the Universidad de Costa Rica, a Postgraduate degree in Alternative Dispute Resolution at the Universidad Latina de Costa Rica, an Alumni degree in Economic Law with an Emphasis in International Trade at the Universidad Estatal a Distancia de Costa Rica, and an LL.M. in International Technology Law at Vrije Universiteit Amsterdam. She worked in the private sector as legal counsel for financial and insurance companies.

*“What governance frameworks does the financial sector need for the quantum era?”*

[Melissa Hernández – Standards for Quantum Technology in Finance](#)

**Olga Mamlyga** is the CEO of Quantum Scouts. She leads a team developing quantum network solutions for critical infrastructure across energy, finance, telecommunications, and security sectors. The company supports organizations in their quantum adoption journey by assessing vulnerabilities, designing hybrid classical–quantum security architectures, and delivering pilot projects with research and industry partners. Quantum Scouts bridges deep technical expertise with practical engineering to help institutions move from theory to real-world quantum readiness.

*“A cyberwar is already underway.”*

[Olga Mamlyga – Building Quantum-Safe Finance](#)

**Orlagh Neary** is a quantum and AI commercialization leader with 25+ years of experience, including as former VP of Microsoft’s Quantum and AI Ecosystem Engagement team. She works with executives and founders to build quantum literacy and go-to-market strategies that translate breakthrough technologies into category leadership and revenue growth. Author of the Quantum Links series, she chronicles the people and ideas driving the quantum era. Orlagh is also the Founder of The ORB Network, a fast-growing leadership and wellness community for women in the Seattle area. Originally from Ireland and now based in Seattle, she graduated from Dublin City University with a B.A. in Applied Languages.

*“If you are reading this and you are not a physicist, if your background is in finance, or healthcare, or law, or communications, or policy, or community development, or any of the thousands of other disciplines that make a complex world function, I want you to hear this clearly. Quantum computing needs you.”*

[Orlagh Neary – An Accidental Love Story](#)

**Pascal Halfmann, PhD**, is a researcher at the Fraunhofer Institute for Industrial Mathematics (ITWM), where he coordinates the institute’s research activities in quantum computing. He is a mathematician with a PhD in mathematical optimization and recognized expertise in multiobjective and robust optimization. His research focuses on the development of novel quantum algorithms, especially

for complex optimization problems. He evaluates the practical potential of quantum computing through benchmarking and real-world applications, with a particular focus on finance—especially portfolio optimization—and energy systems.

*“Entering the field therefore did not feel like a natural extension of my previous work, but rather like stepping into a new intellectual territory with its own language, assumptions, and ways of thinking. This initial distance, however, turned out to be an advantage.”*

[Pascal Halfmann – Where Quantum Computing Meets Optimization](#)

**Rafał Pracht** is a computer science graduate from the Military University of Technology in Warsaw with nearly twenty years of experience in the financial industry. He has worked with institutions including Deloitte, Moody’s Analytics, BNP Paribas, KBC Group, and PZU, focusing on financial engineering, software architecture, and risk systems. He is currently the CTO at FinQbit, where he leads the development of quantum algorithms for financial applications. Rafał is an IBM Qiskit Advocate and holds an MIT xPRO certificate in Applications of Quantum Computing. His work focuses on applying quantum computing to derivative pricing and risk modeling in quantitative finance.

*“For decision-makers in finance, the key takeaway is not the physics itself, but the consequence: quantum computation provides a new method for estimating complex probabilistic quantities.”*

[Rafał Pracht – From Bits to Qubits](#)

**Reinaldo Caelho** is Founding Partner at Triaxis Capital. He has been working with Venture Capital since 2008, in 4 different funds. Made investments in more than 50 companies. Focuses in fintech, healthtech, and marketing & sales tech. Lead investor in remarkable companies like Bling, Vindi, and Konduto. Holds a Masters degree in Industrial Engineering from Virginia Tech; Executive Education from UC Berkeley; and a PhD in Finance, a Masters degree in Economics, and a Bachelors in Mechanical Engineering from UFSC. Professor of Entrepreneurship & Finance at ESAG/UDESC. Quantum Finance enthusiast. CAIA, CISI, CGA-CGE-CFG.

*“Execution, capital discipline, and the ability to navigate the long timeline between scientific potential and paying customers remain decisive factors.”*

[Reinaldo Coelho – Quantum Finance and Venture Capital](#)

**Santanu Ganguly** is a UK-based researcher and professional focused on HPC-enhanced quantum and AI, simulations, cybersecurity, and networks—areas central to his PhD research, and has experience across the domains of finance, security, government, and enterprise. He collaborates part-time with research labs and has held several senior roles at Silicon Valley companies spanning AI, data communications, quantum technologies, cloud, and security. He is the author of the book *Quantum Machine Learning: An Applied Approach*

(Springer Nature, 2021), has authored or co-authored several research papers, and holds five US patents in quantum technologies, AI, cloud, and security.

*“Cryptographic choices should be defined by enterprise policy rather than individual developers or isolated teams.”*

[Santanu Ganguly – Post Quantum Cryptography](#)

**Sebastian Torres** is a Quantitative Trader, Data Scientist, and registered CTA specializing in algorithmic trading and risk modeling. As Chief Investment Officer at PRO Systematic Capital Management and a quantitative developer for PRO Investing, he has led institutional strategies for portfolio optimization and capital preservation. With experience at StoneX and academic credentials from the University of the Andes, MIT, and Stanford, Sebastian merges classical financial rigor with advanced machine learning. He is currently focused on leveraging quantum computing to overcome the mathematical limitations of risk management and efficiency in global financial markets.

*“Financial institutions must invest heavily in cleaning, structuring, and maintaining their data lakes right now.”*

[Sebastian Torres – Quantum Readiness for Financial Adoption](#)

**Sebastian Zajac** holds a PhD in Physics and is Assistant Professor at SGH Warsaw School of Economics. His work focuses on Quantum Machine Learning, graph data analysis, and the development of the Quantum Information Field Theory (QIFT). With a rich background in topological physics and bioinformatics (published in Nature and PRD), he currently translates quantum mechanics into business solutions as an MLOps Engineer. He collaborates with QPoland on quantum computing education. When not exploring the frontiers of data science, he is either playing jazz on the piano or trekking in the Tatra Mountains.

*“This evolution of computational paradigms is not merely a technical shift, but a testament to the profound convergence of fundamental physics and practical data science.”*

[Sebastian Zajac & Krzysztof Kuba – Quantum Neural Networks](#)

**Sierra Clouse** is the Founding Managing Partner of Barclo Venture Studio, where she evolves emerging tech—including quantum—from research papers into products & profits. She leads Barclo’s mission to move breakthroughs out of the lab & into the hands of customers who shouldn’t have to be experts to benefit from their value. Sierra operates on a hard market truth: the best technology doesn’t always win. Success belongs to the firms behind products that ship & sell to those who want, need, & can access them.

*“Quantum tech isn’t a single ‘eureka’ moment. It is a layered ecosystem forming in parallel across academia, startups, enterprises, national labs, & defense agencies.”*

[Sierra Clouse – Musing Between Breakthroughs & Builds](#)

**Tatiana Mitrova** is an independent director, strategic advisor, and energy geopolitics analyst focused on how large energy and industrial systems behave under geopolitical, regulatory, and market stress. She is a Global Fellow at Columbia University’s Center on Global Energy Policy, a Senior Research Fellow at the Oxford Institute for Energy Studies, and Director of the New Energy Advancement Hub. Tatiana brings over a decade of non executive board experience in international energy and industrial companies, with a focus on risk oversight, long term strategy, and governance in uncertain environments. Her work connects geopolitics, energy systems, and decision making, emphasizing resilience and judgement under pressure.

*“A security analysis that stops at technology misses where global security is often decided: governance and chokepoints.”*

[Tatiana Mitrova – Quantum and Security](#)

**Tomasz Ćwik** is a co-founder of the finQbit, with over 15 years of experience across the IT and financial sectors. He is a graduate of multiple top-tier entrepreneurship programmes, including Stanford University, Techstars or CDL. He successfully bridged academia and business by establishing and leading R&D units in CEE that pioneered practical quantum and hybrid quantum-classical applications for risk and financial modelling. His contributions have earned him global recognition, including being named one of the top 23 quantum strategists worldwide by Quantum Insider. Tomek has also collaborated as an expert with organisations such as the CQF Institute and the European Institute of Innovation and Technology (EIT).

*“Quantum computing, as an emerging technology, requires a completely different pace, experimentation, rapid prototyping, hypothesis testing, and sometimes failure. At some point, we realized that to fully leverage this technology’s potential in finance, a more flexible structure than a bank was needed.”*

[Tomasz Ćwik – The Quantum Edge in Finance](#)

**Tse Loong Chin** is a senior coverage banker at HSBC, spearheading the bank’s digital assets and cryptocurrency initiatives for institutional clients across Asia Pacific. With a 30-year career spanning SMBC, UOB, and DBS, he now focuses at the nexus of finance and emerging technology, championing blockchain applications for asset digitalisation. A trilingual leader, he excels at navigating complex regulatory landscapes and building collaborative teams to deliver innovative, cross-border financial solutions that position his clients at the forefront of industry change.

*“I knew I didn’t need to become a quantum physicist. I didn’t need to write quantum code. What I needed was to understand the application and its impact. I needed to understand the ‘so what?’ for finance.”*

[Tse Loong Chin – The Coverage Banker](#)

# Chapter 1

## Quantum Ethics

*Yes, they're both good and bad, and everything in between, at the same time*

Adrian Maguire

When I recently told a friend that I had been asked to write an article about quantum ethics, his joking response was, “Are they really, really, small, tiny things?” We both cracked up laughing and then the conversation moved on. Most of us left philosophy behind in the sixth form, if we ever encountered it at all. So when I started to learn more I was helped by first understanding where the subject sat in terms of philosophy.

Ethics, or moral philosophy, is simply the branch of thinking that asks: “How should we live?” or “What does it mean to live a ‘good’ life?” It sits alongside four other big branches of philosophy like limbs on a tree:

- **Metaphysics** asks what reality actually is.
- **Epistemology** asks how we can know anything.
- **Logic** asks how to reason without falling over.
- **Aesthetics** asks what is beautiful.

Ethics is the one that keeps you awake at night when the technology you’re building might help millions, harm them, or, quantum-style, do both at once.

Within ethics itself there is another little tree. At the top: **meta-ethics** (what do words like “good” even mean?). Then **normative ethics**, the big competing recipes for right action:

- **Virtue ethics** (Aristotle): become the kind of person who does the right thing.

---

This chapter is a featured contribution to the monograph *A Portrait of Quantum Technologies in Finance*, edited by Oswaldo Zapata, PhD, and published by The Quantum Finance Boardroom (2026). All rights reserved by the chapter’s author.

- **Deontology** (Kant): follow unbreakable rules, regardless of consequences.
- **Utilitarianism** (Bentham and Mill): maximise net happiness, full stop.

And finally **applied ethics**, where we roll up our sleeves and argue about concrete cases; AI, quantum, finance, you name it.

Professor John Tasioulas at Oxford’s Institute for Ethics in AI has said, with characteristic clarity, that much of today’s “AI ethics” has hollowed out the subject into a thin, technocratic version of Benthamite utilitarianism: just run the numbers and pick the option with the highest net good.

The Lyceum Project, his collaboration with Stanford’s Human-Centred AI and Greece’s Demokritos, pushes back. It revives an Aristotelian approach: AI (and by extension quantum technologies) should be judged by whether they help human beings flourish in the full, rich sense, exercising reason, forming relationships, living in just communities, and not merely by whether they deliver a bigger GDP or fewer accidents.

That richer picture is more complete because quantum technologies are not neutral tools. They are superposition machines: capable of being good and bad, helpful and harmful, at the same time, until the moment we measure them by the choices we make.

If you want to get an ancient philosophy take on AI, I highly recommend The World’s First Multi-Agent Socratic Dialogue on Artificial Intelligence at The AI Lyceum<sup>©</sup> <https://republic.theailyceum.com/>

## 1.1 A Childhood in Superposition

I grew up living above a shop that my mother managed. Sometimes, when she wasn’t looking, I’d slip a packet of sweets into my pocket. Numerous times I lied to my schoolteachers about why my homework wasn’t handed in. And I also recall playing board games with my siblings and moving the pieces to my advantage when my opponent wasn’t looking. Small, ordinary moral failures. None of them made me “a thief,” “a liar,” or “a cheat” for life. They were data points, not destiny. That lesson has stayed with me: actions are not identities. Technologies are not destinies. We get to choose which branch of the wavefunction we collapse into at any point.

My father once told me about an elderly relative in rural Ireland in the 1950s. When the government arrived to install mains electricity and running water to all remote farmhouses free of charge, a gift from the Irish Government, the old man sent them packing. “I don’t want that new-fangled stuff upsetting my way of life.” In 2019 my siblings and I took our 90-year-old father on a trip to Ireland and we went to see the old farmhouse. It was derelict, overgrown, barely accessible, not lived in for decades and unloved. The refusal to adapt had frozen the place in time.

The Luddite movement in the UK knew the same tension. In 1811 in Nottingham, where I now live, a man named Ned Ludd became the figurehead for textile workers who smashed power looms. They weren’t anti-technology in some cartoon sense;

they were protesting the way the new machines were being used to destroy skilled jobs and communities. There’s a pub in Nottingham city centre named after him (and it serves a very respectable pint of Guinness). Every time I walk past it I’m reminded: resistance to change is human. So is the need to adapt. The question is always: on whose terms?

## 1.2 From Y2K to Q-Day

During 1998 and 1999 I worked on some Y2K projects. At the time plenty of people scoffed: “It’s just a big consultancy scam.” Then midnight 31 December 1999 came and went, the world partied, and no planes fell from the sky. The quiet success was dismissed as proof it was never a problem. But the fixes had been real. Now we face Q-Day, the unknown date when a sufficiently powerful quantum computer can crack RSA and ECC, the cryptographic foundations of the modern internet. Estimates have shortened dramatically in recent years. From an innovation perspective, five to ten years is best-case; the sooner the better. From a national-security perspective, five years is the worst-case. Legacy systems in critical national infrastructure are still running on code written when I was in short trousers. Not preparing is not neutral. It is a choice.

The Jaguar Land Rover cyber-attack in 2025 made the point brutally. Production lines halted for weeks, thousands of families lost income, suppliers trembled. Analysts called it the costliest cyber incident in UK history. Years of successive IT-security budget cuts had left the company exposed. Were those budget prioritisation decisions ethical? When the micro-economic pain lands on real households the answer is obvious. It is not ethical to leave known, material risks unmitigated when the tools to address them exist. Financial institutions that continue to treat post-quantum cryptography as a 2030 problem are making the same ethical miscalculation that Land Rover made with IT security budgets. The question is no longer “if” but “who pays when the wavefunction collapses?”

## 1.3 The Responsible Quantum Industry Forum

In 2025 I was proud to be part of the team convened by the UK’s National Quantum Computing Centre (NQCC) along with industry co-chairs TechUK and UKQuantum, to develop an ethical framework for the emerging quantum ecosystem to use as guidance for their developments and plans. That group evolved into the Responsible Quantum Industry Forum (RQIF)<sup>1</sup> that meets twice a year to share, discuss and learn more on operationalising responsible practices using quantum technologies. The principles we published provide a practical compass:

This approach is not dissimilar to the UK Government’s White Paper on AI regulation (to which I was also asked to contribute). That paper proposed five cross-sectoral

---

<sup>1</sup><https://www.nqcc.ac.uk/updates/the-responsible-quantum-industry-forums-principles-launched/>.

Principle	Description
Open collaboration on responsible quantum	Be collaborative and open in upholding principles and pursuing responsible development.
Forward-looking and responsive	Proactively consider impacts on individuals, society, and the environment; respond proportionately.
Transparency and explainability	Be honest about capabilities and impacts; aim for explainable outcomes.
Safety and reliability	Design, build, and test for safety and reliability.
Privacy and security	Proactively address risks to privacy and security.
Equitability, fairness, and inclusivity	Promote fair access and build a diverse, inclusive ecosystem.
Accountability	Demonstrate commitment through actions and decisions.

principles for existing market regulators to use for supporting their industries:

- Safety, security and robustness
- Appropriate transparency and explainability
- Fairness
- Accountability and governance
- Contestability and redress

These principles echo broader scholarly work, including the 2024 “Ten Principles for Responsible Quantum Innovation” published in *Quantum Science and Technology*, which organises guidance into safeguarding, engaging, and advancing categories. They reject both naïve optimism and Luddite rejection, insisting instead on anticipatory governance.

A 2022 paper in the same journal, “Bridging the quantum divides: a chance to repair classic(al) mistakes?”, puts it well. The authors argue that quantum technologies give us a second chance to avoid the ethical missteps of classical computing and AI; the rush to deploy, the concentration of power, the externalisation of risk. Reading it feels like being handed a better map

Another excellent framework that I regularly consult is the Quantum Applications Framework from the Open Quantum Institute (OQI) at CERN. It offers practical ways to map quantum solutions to real-world problems while keeping ethical and societal dimensions in view. These are not box-ticking exercises. They are guardrails for an ecosystem that is still young enough to choose its character. The RQIF principles have already been referenced in the UK Regulatory Horizons Council’s report on regulating quantum technologies and in the Department of Science, Innovation & Technology’s response to it. They are living documents, meant to be used, stress-tested, continuously improved.

Some consider a voluntary, principles-based approach insufficient and while sympathetic to the concern, this is a nascent ecosystem that requires encouragement and nurturing in order for society to reap the benefits. Any and every individual harm is regrettable, we should remember that even in medicine where the principle of “first, do no harm” is foundational, avoidable adverse drug reactions alone cause hundreds of deaths each year in the UK and contribute to thousands more injuries. The same measured caution should apply to emerging technologies so that they can be scaled rapidly but with precision and care.

## 1.4 Quantum Finance in Superposition

Finance is one of the first sectors where quantum technologies, in particular quantum computing, are moving rapidly from laboratory to ledger. Portfolio optimisation, Monte Carlo risk simulations and quantum machine learning are no longer theoretical. Yet every quantum advantage arrives with a quantum dilemma: faster arbitrage may widen information asymmetries; quantum-enhanced credit models may entrench bias at unprecedented scale; and the race to quantum-safe encryption risks creating a two-tier system of “secure” and “vulnerable” institutions.

The RQIF principles are not abstract philosophy; they are the minimum viable governance layer for any bank, insurer or asset manager that wants to claim it is acting responsibly in the quantum era.

## 1.5 Banking Systems and Financial Literacy

After the 2008 crisis I felt compelled to understand money more deeply. One former Bank of England Governor captured the pragmatic British attitude when he observed that of all the banking systems available, we had somehow ended up with the least worst. Fractional reserve banking, central-bank money vs commercial-bank money, CAPEX vs OPEX, wholesale versus retail, wages versus income; most citizens navigate these distinctions with roughly the same confidence they bring to quantum mechanics. Recent 2025 surveys show that fewer than one in four UK adults can pass a basic financial literacy quiz, and only a minority feel confident with core concepts such as stablecoins or tokenised deposits. That is not a criticism of individuals; it is a comment on the complexity we have built and the education we have failed to provide.

Now we stand on the threshold of AGI economics and potential post-scarcity. Philip Trammell and Anton Korinek’s work “Economic Growth under Transformative AI” maps the possibilities: full automation of production could break the long-run stability of labour share and growth rates that have defined the industrial era. If machines self-replicate and self-improve, output could accelerate dramatically while the portion of income going to human wages collapses. The Kaldor facts for economic growth that have held for decades would shatter.

What then becomes of individual meaning and identity when intelligence is a utility

and human time is abundant? Daron Acemoglu’s Introduction to Modern Economic Growth reminds us that institutions and incentives matter enormously in how technological revolutions are distributed. The coming transition is not pre-ordained to be utopian or dystopian; it is a design problem lacking a vision.

## 1.6 Imagination, Ideation, Innovation

Einstein was right: “Imagination is more important than knowledge. For knowledge is limited, whereas imagination embraces the entire world, and all there ever will be to know and understand.”

In an era when machines can calculate extreme complexities and generate predictions at superhuman speed, the scarce resource becomes human imagination, the ability to see possibilities that are not yet in the training data. Ideation for solving clearly defined problems. Innovation for turning the ideas into reality. The ancient Hindu trinity, Brahma (creation), Vishnu (preservation), Shiva (destruction), is often misunderstood. Shiva does not destroy the idea; he destroys the old “form” of it (an idea form in 0D space) once the new “form” of it (material form in 3D space) has emerged through human action. We need all three.

*Machines generate. Humans create.*

Carl Jung observed: “Problems cannot be solved at the same level of consciousness that created them.” AI trained on human data will inevitably reproduce human blind spots, biases and levels of awareness. It’s not really a surprise to find AI’s manipulating, lying, and cheating to self-preserve when they have been trained to mimic human behaviour. A desktop calculator has been far superior to all humans, super-intelligent, at arithmetic for quite some time. Quantum computers will be super-intelligent at certain classes of combinatorial problems. But the higher-order questions; what should we optimise for, who decides, what kind of society do we want; these remain stubbornly human. Tools augment; they do not replace moral agency.

*Think outside of the tesseract.*

## 1.7 Living in Nottingham, Thinking in Superposition

As I mentioned above, I live in the city where the Luddite movement began and also where Robin Hood is said to have robbed from the rich to feed the poor. The legend has a lot of myth associated, but the moral tension is real: when is redistribution justice and when is it theft? Quantum technologies will force us to ask similar questions at global scale, about access to compute, about who owns the qubits, about whether certain applications should be open-source or classified.

If any of this resonates and you would like to explore subtle yet impactful ideas

at the meeting point of quantum physics, economics and philosophy over a drink, do come to the next QUIPS (Quantum In Pubs) event in London. Credit to Joe Spencer for having both the creative spirit and organising tenacity to make that idea a reality. They are now a highlight of the UK quantum calendar!

## 1.8 The Subtle Promotion

My own journey, 30+ years in financial services, deep experience in data, analytics and governance, industry coalitions, regulatory engagement, plus years of self-directed continuous learning in economics, philosophy, psychology and quantum technologies, has left me at a rare intersection. I have sat in rooms with economists at the Bank of England and quantum physicists at the NQCC. I have helped shape policy submissions and ethical frameworks. I have recruited and led multidisciplinary teams. I understand both the commercial pressure to deploy fast and the societal duty to deploy responsibly.

If you're in the financial and professional services sector and exploring quantum computing capabilities to boost imagination or accelerate scenario planning for Q-Day, I would welcome a conversation. Whether it is a 15-minute introductory call, a leadership briefing, or joint work on a strategic innovation initiative or ethical-audit framework, the door is open.

My purpose is to support a graceful transition into the Age with Digital Quantum Super Intelligence by serving others in the true spirit of collaborative solidarity. I can help develop your future visions and explore subtle yet impactful ethical strategies and plans that help them become manifest.

Because quantum ethics really is both good and bad, and everything in between, at the same time, until we choose.

(References and further reading available on request. All views are my own and do not represent any organisation.)



## Chapter 2

# Combinatorial Portfolio Optimization

### *Quantum-inspired methods for discrete portfolio construction under institutional constraints*

Alejandro Rodriguez Dominguez

**Abstract.** Modern institutional portfolio construction increasingly involves structural constraints that introduce combinatorial complexity into the optimization process. Cardinality limits, regulatory concentration bounds, liquidity constraints, turnover controls, and mandate customization often transform classical continuous allocation problems into mixed-integer or non-convex optimization programs that are computationally challenging to solve at scale. This work examines portfolio construction from the perspective of combinatorial optimization. We distinguish between two sources of out-of-sample degradation in quantitative portfolio design: statistical estimation error and optimization approximation error arising from incomplete exploration of large discrete feasible regions. Within this framework, we discuss Quadratic Unconstrained Binary Optimization (QUBO) formulations and quantum-inspired algorithms executed on classical hardware, including stochastic global search methods such as simulated annealing. The goal is not to claim advantages derived from quantum hardware, but to analyze how these algorithmic approaches facilitate structured exploration of large combinatorial search spaces in practical institutional applications. Examples include cardinality-constrained index tracking, partial index replication, and personalized portfolio construction. Empirical illustrations suggest that hybrid discrete–continuous optimization methods can generate high-quality feasible portfolios under realistic constraints while maintaining tracking discipline and risk control

---

This chapter is a featured contribution to the monograph *A Portrait of Quantum Technologies in Finance*, edited by Oswaldo Zapata, PhD, and published by The Quantum Finance Boardroom (2026). All rights reserved by the chapter's author.

## *A Practitioner’s Journey into Quantum-Inspired Finance*

My personal exposure to quantum-inspired computation arose during my time as a quantitative researcher at the asset management division of Miralta Finance Bank S.A. At the firm, we had developed a comprehensive fixed income quantitative infrastructure, built on more than a decade of daily data covering approximately 7,000 bonds globally. Each instrument was described by a rich feature set of roughly 100 variables, including pricing attributes, liquidity indicators, and qualitative descriptors. The platform supported a broad range of analyses, including fair value estimation, capital structure modeling, factor decomposition, liquidity assessment, portfolio construction, and relative value strategies.

Within this framework, we set out to develop a systematic approach to measuring dispersion in fixed income markets. The objective was not limited to comparing individual securities, but rather to characterize dispersion at the level of structured portfolios. In particular, we aimed to construct equal-weighted baskets of bonds defined by combinations of qualitative and quantitative features, and to analyse the dispersion of spreads between such baskets across time. Conceptually, this problem is closer to identifying extremal structures in a discrete space of feature-conditioned portfolios than to traditional mean–variance optimization.

In practical terms, baskets were defined by selecting subsets of features and grouping bonds sharing the same categorical attributes. Each such grouping induces a family of portfolios, and the analysis focuses on pairwise comparisons between these portfolios using spread-based indicators such as rolling z-scores or signal-to-noise ratios. The problem therefore reduces to evaluating a very large number of basket pairs across time and identifying those that exhibit persistent dispersion.

The computational implications of this formulation are substantial. Even at the most basic level, the number of pairwise spreads across 7,000 bonds is

$$\binom{7000}{2} = 24,496,500,$$

which must be evaluated repeatedly across time. Under conservative assumptions, computing all such spreads for a single trading day already requires on the order of hours, and extending this across a ten-year daily dataset leads to runtimes measured in months on a single processing unit.

However, this pairwise analysis represents only the lowest level of complexity. The true combinatorial explosion arises when baskets are constructed using feature intersections. To provide a representative example, consider approximately 50 variables, each discretized into around 10 categories. The number of possible feature groupings of size  $k$  grows as

$$\binom{50}{k} \cdot 10^k,$$

which already becomes extremely large for moderate values of  $k$ .

For baskets defined by just two features, there are 1,225 feature pairs, each generating up to 100 distinct baskets. The number of basket pairs per feature pair is therefore on the order of several thousand, leading to roughly 6 million basket-pair evaluations per day. Even this restricted setting already implies multi-day computations for a full historical dataset. At three-feature depth, the situation becomes dramatically more severe. There are nearly 20,000 feature triplets, and each one can generate up to 1,000 baskets, resulting in approximately 500,000 basket pairs per triplet. In aggregate, this corresponds to nearly 10 billion basket-pair evaluations per day. At realistic computation speeds, this translates into decades of processing time for a single year of data, and well over a century for a full historical analysis.

Beyond this level, the problem rapidly becomes astronomically large. The number of feature combinations continues to grow, while the number of admissible baskets becomes increasingly sparse and irregular due to the finite size of the asset universe. In aggregate, a full combinatorial exploration across feature subsets leads to a number of candidate configurations that is effectively beyond any feasible classical enumeration. The key point is that the difficulty is not merely computational in the engineering sense; it is structural. The problem itself belongs to a class of high-dimensional combinatorial optimization problems for which exhaustive approaches are fundamentally impractical.

My initial attempts to address this problem relied on classical techniques, including parallelization frameworks such as Spark and matrix-based formulations designed to exploit linear algebraic structure. While these approaches provided some improvements, they did not alter the fundamental nature of the problem. The combinatorial geometry of the search space remained unchanged, and exact computation was ultimately infeasible.

At a later stage, I was contacted by Samuel Fernández Lorenzo, who was founding Inspiration-Q (currently acting CEO), a company focused on quantum-inspired algorithms for combinatorial optimization. These methods do not rely on quantum hardware; rather, they use heuristics inspired by quantum mechanics (such as annealing dynamics and energy minimization) to efficiently explore large discrete solution spaces using classical computational resources. Many such problems can be formulated within a quadratic optimization framework over binary variables, enabling structured search over extremely large configuration spaces.

What began as an initial discussion quickly evolved into a deeper collaboration. After presenting the dispersion problem, it became clear that its structure aligned naturally with the class of problems that Inspiration-Q had been addressing, particularly in portfolio optimization and index tracking. I subsequently joined the company as an advisor, contributing to the extension of their framework toward large-scale dispersion analysis in fixed income markets.

The key insight was that the problem could be reformulated as an optimization over a discrete configuration space, where each candidate basket or basket pair corresponds to a structured selection over assets. Within this formulation, quantum-inspired algorithms provide an efficient mechanism for navigating the search space and identifying high-quality solutions without requiring exhaustive enumeration. This represents a fundamental shift from continuous optimization in convex domains

to combinatorial search over discrete geometries.

The resulting implementation enabled, for the first time, the practical computation of dispersion structures at this scale. The tool is now deployed across multiple asset classes, extending beyond fixed income into broader financial applications. Its use cases include the generation of risk signals, the identification of relative value opportunities, and the enhancement of statistical arbitrage strategies, as well as improvements in execution efficiency.

This experience illustrates a central theme of this chapter: when portfolio construction problems are formulated in discrete terms—whether due to feature-based grouping, cardinality constraints, or institutional requirements—the underlying geometry departs fundamentally from the convex framework of classical portfolio theory. In such settings, the feasible set becomes a combinatorial object, and optimization requires methods capable of navigating complex, high-dimensional discrete spaces. Quantum-inspired algorithms provide a natural and powerful approach to these problems.

From a broader perspective, these methods offer a computational paradigm capable of addressing classes of optimization problems that are otherwise only tractable through approximation. In finance, this has direct implications for portfolio construction, risk management, pricing, and trading, particularly in environments where constraints induce non-convex and discrete structures.

Looking forward, quantum-inspired and quantum computing methods are likely to play a significant role in the evolution of financial technology. Their ability to efficiently address high-dimensional, non-convex optimization problems has the potential to improve the efficiency of financial systems, from portfolio allocation and execution to transaction processing and digital asset infrastructure. At the same time, these advances introduce new challenges, including cybersecurity risks, model risk, and potential vulnerabilities to market manipulation. As a result, their adoption will require not only technological innovation but also appropriate regulatory and institutional frameworks to ensure their robust and responsible use.

## 2.1 Introduction

The construction of investment portfolios has undergone a progressive formalization over the past decades. While earlier approaches relied heavily on discretionary judgment, modern quantitative asset management is grounded in mathematical optimization, statistical estimation, and computational methods. The availability of large financial datasets together with advances in numerical optimization has enabled the systematic implementation of scalable portfolio construction frameworks.

The foundational formulation of the portfolio selection problem was introduced by Markowitz (1952), who cast asset allocation as a quadratic optimization problem in the space of expected returns and covariances. Subsequent research has addressed well-known limitations of this framework, particularly its sensitivity to estimation error (Best & Grauer, 1991; Chopra & Ziemba, 1993). This literature led to the development

of improved estimators and portfolio construction techniques, including shrinkage estimators (Ledoit & Wolf, 2004), Bayesian approaches such as Black–Litterman (Black & Litterman, 1992), and robust optimization methods (Goldfarb & Iyengar, 2003). Alternative convex formulations have also been proposed, including mean absolute deviation portfolios (Konno & Yamazaki, 1991) and CVaR minimization (Rockafellar & Uryasev, 2000). When the covariance matrix is positive semidefinite and constraints are linear, these formulations remain convex and can be solved efficiently using interior-point methods (Boyd & Vandenberghe, 2004).

In such convex settings, portfolio instability is primarily attributed to statistical uncertainty rather than computational intractability. Indeed, explicit constraints may improve out-of-sample behavior within the convex regime (Jagannathan & Ma, 2003), although estimation risk remains a structural concern (DeMiguel et al., 2009).

In practical institutional mandates, however, portfolio construction often includes structural constraints that alter the mathematical nature of the problem. Cardinality limits, minimum position sizes, turnover controls, regulatory concentration bounds, liquidity thresholds, and tax constraints introduce discrete decision variables or non-convex feasible regions. A representative example is the imposition of a cardinality constraint requiring the selection of  $K$  assets from a universe of  $N$ . The number of possible subsets grows as

$$\binom{N}{K},$$

which increases combinatorially with  $N$ . Even for moderate values of  $N$ , the resulting search space becomes prohibitively large.

When such discrete features are incorporated explicitly, the classical mean–variance problem can be formulated as a mixed-integer quadratic program (MIQP), which is NP-hard in general (Garey & Johnson, 1979; Nemhauser & Wolsey, 1988; Bienstock, 1996). The computational difficulty arises not from the quadratic objective itself, but from the combinatorial structure induced by binary selection variables and cardinality constraints. Even when the continuous weight allocation subproblem remains convex, the presence of binary decision variables transforms the overall optimization into a mixed discrete–continuous problem. As a result, portfolio construction under realistic institutional constraints often moves beyond the domain of classical convex optimization and into the broader framework of combinatorial optimization.

These observations motivate a conceptual distinction between two sources of out-of-sample degradation in portfolio construction:

1. **Estimation error**, arising from imperfect parameter estimates (Michaud, 1989; Ledoit & Wolf, 2004);
2. **Optimization or approximation error**, arising from limited exploration of complex discrete feasible regions.

While the first component has been extensively studied in financial econometrics, the second is more closely related to developments in combinatorial optimization and operations research. In large-scale institutional contexts—such as index replication with limited constituents, enhanced index tracking, or constrained long–short

mandates—optimization error may become material if the feasible set cannot be adequately explored.

To address such problems, a range of algorithmic approaches has been proposed, including branch-and-bound mixed-integer solvers, convex relaxations (e.g.,  $\ell_1$  penalization) (Tibshirani, 1996; Brodie et al., 2009), evolutionary heuristics (Beasley et al., 2003), and metaheuristic methods such as simulated annealing (Geman & Geman, 1984; Hajek, 1988). More recently, formulations based on Quadratic Unconstrained Binary Optimization (QUBO) have been explored in finance, enabling the application of quantum and quantum-inspired search procedures (Fernández-Lorenzo et al., 2021; Mugel et al., 2022; Rubio-García et al., 2024). These methods should be understood as stochastic global search techniques operating on classical hardware, without guarantees of global optimality in finite time.

The purpose of this chapter is to clarify the boundary between convex, tractable portfolio formulations and combinatorial or non-convex structures that lead to computational hardness. Rather than advocating a specific solver class, we analyze how structural constraints modify the geometry of the feasible set, how this affects computational complexity, and how different algorithmic paradigms trade off scalability, solution quality, and practical implementability.

In institutional portfolio construction, both estimation and combinatorial complexity play a role. Understanding their interaction is essential for designing coherent and scalable allocation frameworks.

## 2.2 Quantum-Inspired Algorithms and Combinatorial Portfolio Optimization

The increasing prevalence of structural constraints in institutional portfolio construction has renewed interest in algorithmic paradigms capable of addressing discrete optimization problems at scale. While discussions of quantum computing in finance often emphasize prospective asymptotic speedups from hardware advances, the dominant challenge in many current applications is structural rather than purely computational: the combinatorial geometry of the feasible set.

This section examines quantum-inspired optimization methods as one class of algorithms designed to explore large discrete search spaces. The focus is on structural suitability and empirical behavior rather than on hardware-dependent acceleration claims.

### Discrete structure and QUBO reformulation

Many institutional portfolio problems are intrinsically discrete. Cardinality limits, basket selection, partial index replication, binary regulatory constraints, and certain long–short exposure controls cannot be represented exactly within purely continuous convex programs without relaxation.

A convenient mathematical representation introduces binary selection variables

$$x_i \in \{0, 1\}, \quad i = 1, \dots, N,$$

and expresses the objective in quadratic form:

$$\min_{x \in \{0,1\}^N} x^\top Q x.$$

This corresponds to a Quadratic Unconstrained Binary Optimization (QUBO) formulation. The matrix  $Q$  encodes tracking error terms, covariance interactions, expected return adjustments, and penalty representations of structural constraints.

The relevance of QUBO lies in its structural correspondence to classical combinatorial optimization problems. It enables the application of algorithmic techniques originally developed in statistical physics and stochastic optimization to financial selection problems. Recent studies have demonstrated the feasibility of QUBO-based formulations for cardinality-constrained index tracking and related applications (Fernández-Lorenzo et al., 2021; Mugel et al., 2022; Rubio-García et al., 2024). It is important to emphasize that QUBO reformulation does not remove NP-hardness; it provides a structured representation that enables the application of heuristic and metaheuristic search procedures.

A standard way to encode cardinality constraints in a QUBO formulation is through quadratic penalty terms. Let  $x_i \in \{0, 1\}$  denote the binary selection variable for asset  $i$ , and let  $K$  denote the desired portfolio cardinality. A typical formulation augments the quadratic objective with a penalty term:

$$\min_{x \in \{0,1\}^N} x^\top Q x + \lambda \left( \sum_{i=1}^N x_i - K \right)^2, \quad (2.1)$$

where  $Q$  encodes covariance interactions, tracking error contributions, or return adjustments, and  $\lambda > 0$  controls the strength of the cardinality constraint. The quadratic penalty encourages feasible solutions with exactly  $K$  selected assets while preserving the unconstrained quadratic structure required for QUBO-based optimization methods.

## Relationship to classical mixed-integer solvers

Mixed-integer quadratic programming (MIQP) provides an exact framework for solving cardinality-constrained mean–variance problems. Branch-and-bound or branch-and-cut methods can yield globally optimal solutions in finite time. For moderate dimensions, commercial solvers are effective.

However, worst-case complexity grows exponentially with problem size. As the number of assets, constraints, or scenario evaluations increases, exact solution times may become prohibitive, particularly in contexts requiring repeated rebalancing or scenario stress testing.

Hybrid approaches combining discrete search with convex subproblems have therefore been proposed. For example, Rubio-García et al. (2024) develop a hybrid simulated annealing (SA) approach in which asset selection is treated combinatorially, while weights are optimized via quadratic programming. In empirical experiments on S&P 500 index tracking, portfolios with 10–30 constituents were obtained in seconds to minutes on classical hardware, with objective improvements saturating beyond moderate search depth.

To situate quantum-inspired methods within the broader optimization landscape, Table 1 summarizes structural and computational trade-offs across commonly used approaches.

Table 1: Comparison of optimization approaches for cardinality-constrained portfolio selection.

Method	Exact nality	Cardi- nality	Global Opti- mality Guar- antee	Opti- mality Guar- antee	Scalability (Typical)	Remarks
MIQP (Branch-and-Bound)	Yes		Yes (finite time)	(finite)	Moderate ( $N \lesssim 200$ )	Exponential worst-case complexity
$\ell_1$ Relaxation (Convex)	Approximate		Yes (convex optimum)		High ( $N \gtrsim 1000$ )	Sparsity induced but $K$ not fixed
Hybrid SA-QP	Yes		No (asymptotic only)		High ( $N \gtrsim 500$ feasible)	Empirically near-optimal
QUBO Metaheuristics	Yes		No (heuristic)		High (parallelizable)	Natural encoding of discrete structure

While worst-case complexity is exponential for all exact combinatorial methods, practical scaling behavior differs substantially. Table 2 provides indicative magnitudes for index tracking problems with cardinality constraint  $K$ , based on representative experimental settings reported in (Rubio-García et al., 2024).

Table 2: Indicative scaling behavior for cardinality-constrained index tracking (illustrative magnitudes).

$N$	$K$	$\binom{N}{K}$	MIQP Time	Hybrid SA Time	Typical Tracking Error Gap
100	10	$2.6 \times 10^{13}$	seconds–minutes	< 1 sec	negligible
250	20	$> 10^{32}$	minutes–hours	few seconds	< 5 bps
500	30	$> 10^{52}$	often impractical	$\sim 10$ –60 sec	< 10 bps

The magnitudes in Table 2 highlight the combinatorial growth of the feasible region. Exact MIQP solvers remain viable for medium-scale instances but exhibit rapidly increasing runtime as  $N$  grows. Hybrid annealing-based approaches scale more smoothly in practice, though without strict optimality guarantees. Empirical results suggest diminishing objective improvements beyond moderate search depth, indicating that statistical noise may dominate residual optimization improvements in large universes.

## Annealing-based global exploration

Annealing methods define a stochastic search process governed by

$$\pi_T(x) \propto e^{-f(x)/T},$$

allowing controlled acceptance of non-improving moves (Kirkpatrick et al., 1983; Geman & Geman, 1984; Hajek, 1988). Under theoretical cooling schedules, convergence to global optima can be established asymptotically, though practical schedules are finite and heuristic. The practical performance of simulated annealing depends critically on the cooling schedule and neighborhood search heuristics. Classical theoretical results guarantee convergence to global optima only under extremely slow cooling schedules (Hajek, 1988), which are impractical in large-scale applications. In practice, implementations typically adopt geometric cooling schedules of the form

$$T_{t+1} = \alpha T_t, \quad 0 < \alpha < 1,$$

which provide a balance between exploration and computational efficiency. More advanced variants incorporate adaptive temperature updates based on acceptance rates, reheating mechanisms, or parallel tempering schemes that run multiple temperature chains simultaneously. These techniques can improve exploration of complex combinatorial landscapes and reduce the probability of becoming trapped in local minima.

In high-dimensional combinatorial landscapes, such probabilistic dynamics can reduce sensitivity to initialization and mitigate local minima entrapment. Empirical evidence in index tracking applications (Rubio-García et al., 2024) indicates that broader exploration improves in-sample objective values, while out-of-sample gains tend to plateau beyond moderate optimization depth.

## Hardware considerations

Current quantum hardware remains in the NISQ regime (Preskill, 2018), with limitations in qubit count, connectivity, and coherence. Most financial applications therefore rely on quantum-inspired algorithms executed on classical hardware.

These approaches exploit discrete reformulations (e.g., QUBO) without requiring physical qubits. Their scalability derives from classical parallelization rather than from proven quantum speedups. Whether future quantum hardware will offer practical advantages in portfolio optimization remains an open empirical question.

## Scope and limitations

Quantum-inspired optimization methods do not eliminate estimation error and do not guarantee global optimality in finite time. Their contribution lies in structured exploration of discrete feasible regions when convex relaxations are insufficient and exact MIQP solvers become computationally burdensome.

Improvements from deeper combinatorial search typically exhibit diminishing returns beyond moderate optimization depth, as documented empirically in index tracking experiments (Rubio-García et al., 2024). Accordingly, such methods should be viewed as complementary tools within a broader quantitative workflow rather than as replacements for statistical modeling or economic judgment.

## 2.3 Application to Active and Passive Portfolio Management

The distinction between active and passive management has long structured academic research and industry practice. Passive strategies typically aim to replicate a benchmark index while minimizing tracking error and implementation costs (Malkiel, 2003; Roll, 1993). Active strategies, in contrast, seek to generate excess returns through controlled deviations from a benchmark portfolio (Jensen, 1968). Although conceptually distinct, institutional portfolio construction often reveals a more nuanced structural relationship between the two approaches.

Empirical evidence over the past decade suggests that the persistence of alpha is limited across broad asset classes. According to the Morningstar *Active/Passive Barometer*, only a minority of actively managed funds outperform comparable passive benchmarks over long horizons. For U.S. equity funds, the ten-year success rate is approximately 14 percent, with similar magnitudes observed in European equity categories, while fixed-income strategies exhibit somewhat higher success rates. These figures, summarized in Table 3, are reported net of fees and account for survivorship effects (Morningstar Research, 2024).

Table 3: Ten-Year Success Rates of Active Management (Source: Morningstar Research, 2024)

Asset Class	Ten-Year Success Rate
U.S. Equity	~ 14%
European Equity	~ 13.5%
Fixed Income	~ 26%

At the same time, the structure of asset management has evolved toward greater customization. Separately managed accounts and direct indexing programs now represent more than \$10 trillion in the United States (Cerulli Associates, 2022), reflecting demand for portfolio configurations that incorporate account-specific tax considerations, liquidity constraints, and regulatory requirements. These developments suggest that the relevant distinction between active and passive management is not solely performance-based, but structural: both paradigms rely on solving constrained optimization problems whose mathematical properties depend on the nature of the constraints imposed.

From a formal perspective, many institutional active mandates can be expressed as benchmark-relative optimization programs. Rather than maximizing absolute return in isolation, active management typically operates within explicit tracking-error budgets and risk constraints. A stylized representation is

$$\max_w \mathbb{E}[R(w) - R(b)] \quad \text{s.t.} \quad \text{TE}(w, b) \leq \tau, \quad w \in \mathcal{C},$$

where  $b$  denotes the benchmark, TE the tracking error, and  $\mathcal{C}$  the set of operational and regulatory constraints. Under this formulation, active management may be interpreted as a structured extension of index replication in which systematic factor exposures—such as size, value, quality, or momentum (Fama & French, 1993; Carhart, 1997)—are introduced within controlled deviation limits. The underlying mathematical architecture is therefore shared: both active and passive strategies require

solving constrained optimization problems, differing primarily in the magnitude and structure of permissible deviations.

The computational characteristics of these problems depend critically on constraint design. In the absence of discrete features, tracking error minimization can be written as a convex quadratic program,

$$\min_w (w - b)^\top \Sigma (w - b) \quad \text{s.t.} \quad w \in \mathcal{F},$$

which is solvable in polynomial time when  $\Sigma$  is positive semidefinite and  $\mathcal{F}$  is convex. However, institutional mandates frequently introduce structural constraints that alter this geometry. Limiting the number of holdings, imposing minimum position sizes, enforcing sector or factor exposure bands, controlling turnover relative to prior allocations, and respecting regulatory concentration thresholds may require discrete decision variables or induce non-convex feasible regions. These structural dimensions are summarized in Table 4. These practical portfolio construction settings therefore combine continuous risk optimization with discrete structural constraints, naturally leading to combinatorial optimization formulations.

Table 4: Structural Components in Institutional Index Tracking

Structural Dimension	Mathematical Implication	Computational Consequence
Cardinality limits	Discrete subset selection	Mixed-integer quadratic program (NP-hard)
Sector and factor exposure controls	Interdependent constraints	Potential non-convex feasible region
Turnover constraints	Dependence on prior allocation	Dynamic optimization problem
Regulatory concentration limits	Threshold-based constraints	Fragmented feasible region
Liquidity and exclusion rules	Time-varying eligible universe	Recurrent combinatorial re-optimization

When an explicit cardinality constraint is imposed—for example, restricting selection to at most  $K$  assets from a universe of size  $N$ —the number of feasible subsets grows combinatorially as  $\binom{N}{K}$ . In such cases, the problem becomes a mixed-integer quadratic program and is NP-hard in general (Garey & Johnson, 1979; Nemhauser & Wolsey, 1988; Bienstock, 1996). This structural shift does not eliminate the importance of statistical estimation error; rather, it introduces an additional dimension of potential approximation error. Even with identical return and covariance estimates, incomplete exploration of a discrete feasible region may prevent identification of the best admissible configuration under the imposed constraints. The relative importance of estimation error and optimization approximation error depends on the scale of the universe, the interaction of constraints, and the frequency of rebalancing.

These considerations are amplified in private banking and wealth management contexts. Portfolio construction in such environments incorporates after-tax objectives, heterogeneous liquidity horizons, capital commitment structures associated with private assets, and jurisdiction-specific regulatory requirements. Tax-aware strategies, particularly in direct indexing and separately managed accounts, introduce path dependence through tax-lot selection and capital gains realization. Allocations to

private assets may further introduce partial irreversibility and non-linear liquidity profiles, challenging purely continuous optimization assumptions. Formally, these environments can be represented as collections of coupled optimization problems. If  $M$  client accounts share a common universe of  $N$  assets, each account  $m$  solves

$$\min_{w_m} (w_m - b_m)^\top \Sigma (w_m - b_m) + \lambda_1 \text{TC}(w_m, w_m^{\text{prev}}) + \lambda_2 \text{Tax}(w_m) \quad \text{s.t.} \quad w_m \in \mathcal{F}_m,$$

where  $\mathcal{F}_m$  encodes account-specific structural constraints. The computational challenge arises not only from solving an individual instance, but from repeatedly solving large numbers of heterogeneous instances under changing constraint sets and market conditions. Sequential heuristic adjustments may reduce computational burden, yet they can restrict the effective feasible region in ways that are difficult to quantify.

In both institutional index tracking and personalized wealth management, portfolio construction therefore combines continuous risk minimization with discrete structural decisions. The mathematical characterization of these problems connects directly to the combinatorial optimization framework discussed in the preceding section, where scalable approximation methods become relevant when convex formulations are insufficient to capture operational reality.

## 2.4 Empirical Evaluation: Computational Accuracy, Statistical Generalization, and Investment Performance

This section reports empirical evidence for cardinality-constrained portfolio construction along three axes: (i) computational scalability and approximation quality, (ii) the interaction between optimization effort and out-of-sample (OOS) behavior, and (iii) investment performance in passive tracking and enhanced tracking. The experiments are implemented with a hybrid workflow that combines discrete selection (cardinality) with continuous quadratic weight optimization, consistent with hybrid formulations used in recent quantum-inspired portfolio studies (Fernández-Lorenzo et al., 2021; Rubio-García et al., 2024). Unless otherwise stated, results and figures are reproduced from empirical experiments and materials provided by Inspiration-Q (Inspiration-Q)<sup>1</sup>. All runtime comparisons depend on the optimization target (global optimality versus high-quality feasible solutions), solver settings, and hardware. Likewise, OOS performance summaries are reported as point estimates and should ideally be complemented with uncertainty quantification (e.g., bootstrap confidence intervals, robust inference) and factor attribution to distinguish sampling variation and systematic exposures from residual effects.

The empirical analysis therefore examines how hybrid discrete–continuous optimization behaves in practice when applied to realistic portfolio construction tasks.

<sup>1</sup>The author serves as Quantitative Advisor to Inspiration-Q. Some empirical results and computational benchmarks reported in this chapter are based on experimental materials provided by the company. These results are included for illustrative purposes regarding the computational characteristics of the optimization methods discussed.

### 2.4.1 Computational scalability and approximation quality under cardinality

Cardinality constraints induce a combinatorial search space of size  $\binom{N}{K}$ , which grows rapidly even for moderate  $N$  and  $K$ . Table 5 summarizes representative magnitudes that are typical of index-tracking and enhanced-tracking mandates.

Table 5: Combinatorial scaling induced by cardinality constraints. Even moderate  $(N, K)$  pairs generate extremely large feasible sets, motivating approximate global search methods.

Universe size $N$	Basket size $K$	$\binom{N}{K}$ (approx.)
100	20	$5.36 \times 10^{20}$
130	10	$2.7 \times 10^{13}$
200	30	$2.9 \times 10^{38}$
100	50	$1.0 \times 10^{29}$

A direct implication is that exact mixed-integer quadratic programming (MIQP) solvers become impractical as problem size increases or as constraints become more realistic (turnover caps, minimum weights, sector bounds). For a concrete benchmark, Inspiration-Q reports that a 10-of-130 selection instance can be solved by a simulated-annealing (SA) engine in under 10 seconds, whereas a reference MIQP approach is reported to require extremely long runtimes when targeting certified global optimality under the benchmark configuration for an exact solve at comparable settings (Inspiration-Q, 2025)<sup>2</sup>. Because larger instances cannot be reliably solved to global optimality by MIQP within operational budgets, practical evaluation focuses on approximation quality rather than exact optimality certificates.

Table 6 provides a compact empirical/operational comparison of three families of approaches in this regime. The intent is not to claim universal dominance, but to clarify what is typically optimized (objective value vs. certificates vs. proxies for sparsity) and what is operationally feasible at scale.

### 2.4.2 Optimization strength and statistical generalization

A central concern in practice is that stronger in-sample optimization does not automatically translate into improved OOS outcomes. Recent work emphasizes an approximation–estimation perspective: as optimization error decreases, OOS performance may saturate once statistical estimation noise dominates, so additional computation can yield diminishing returns (Belkin et al., 2019). This effect is particularly salient in portfolio settings where covariance and return estimates are noisy and time-varying.

Evidence of this saturation behavior appears in empirical studies of annealing-based portfolio construction. Figure 1 illustrates this effect: increasing SA steps improves in-

<sup>2</sup>This Table refers to targeting global optimality (i.e., proving optimality) under the specific solver configuration and hardware used in the referenced benchmark. In practice, modern MIQP solvers (e.g., Gurobi, CPLEX, MOSEK) often deliver strong feasible solutions much faster via heuristics and early termination with an explicit optimality gap; runtime depends on solver settings, hardware, and the accepted gap tolerance.

Table 6: Compact empirical/operational comparison for cardinality-constrained portfolio problems. Reported runtime contrast (10-of-130) follows Inspiration-Q empirical benchmarking (Inspiration-Q, 2025).

Method family	What it provides	Typical strengths / limitations	Illustrative runtime signal
MIQP (exact / branch-and-bound)	Optimality certificates (when solvable)	High reliability on small instances; degrades sharply with size/constraints; often requires early termination with nontrivial gaps	10-of-130 reported as $\sim 1$ year for exact solve (Inspiration-Q, 2025).
Simulated annealing (global stochastic search; hybrid SA+QP)	High-quality feasible solutions; no global certificates	Scales to large discrete spaces; naturally handles multiple penalties/constraints via energy terms; quality depends on schedule and run budget	10-of-130 reported $< 10$ seconds (Inspiration-Q, 2025).
Convex relaxations / sparsity penalties ( $\ell_1$ proxy)	Convex tractability; continuous proxy for sparsity	Fast and stable; does not enforce exact cardinality; requires rounding/thresholding, which can degrade feasibility or objective	Fast, but may miss discrete optimum (problem-dependent)

sample objective values monotonically, while OOS metrics plateau beyond a moderate computational budget (Rubio-García et al., 2024). This motivates reporting results at operationally realistic budgets and interpreting incremental in-sample improvements cautiously.

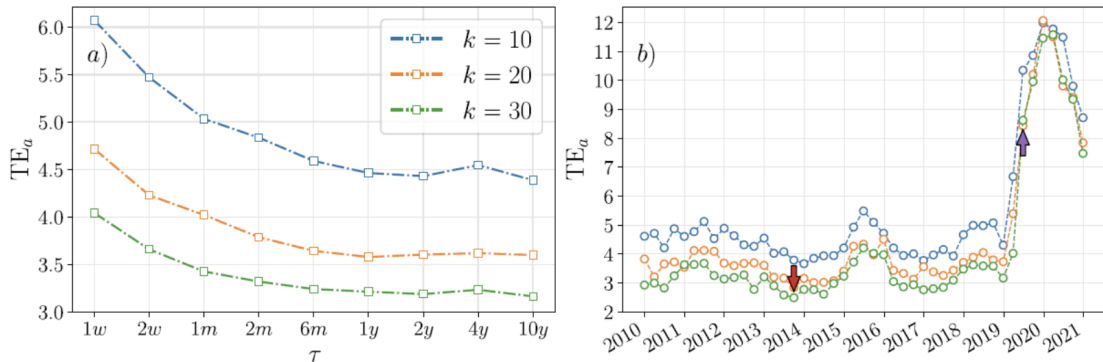


Figure 2.1: Illustration of the approximation–estimation interaction: increasing optimization effort can reduce in-sample error while OOS improvements saturate beyond a moderate budget. Reproduced from Rubio-García et al. (2024) (figure shown via embedded PDF page).

### 2.4.3 Passive index tracking: one-year OOS performance under monthly rebalancing

We report one-year OOS performance under monthly rebalancing for tracker baskets with sizes  $k \in \{10, 20, 30\}$  and a minimum annual excess-return constraint of 2%, as provided in Inspiration-Q empirical materials (Inspiration-Q, 2025). Figure 2 shows

the distribution of annualized tracking error (TE) and annualized excess return (ER). Consistent with portfolio replication theory, TE declines as  $k$  increases, reflecting the additional degrees of freedom available to represent the benchmark, while ER remains positive on average but becomes more sensitive to forecast quality as the basket becomes larger and closer to the benchmark.

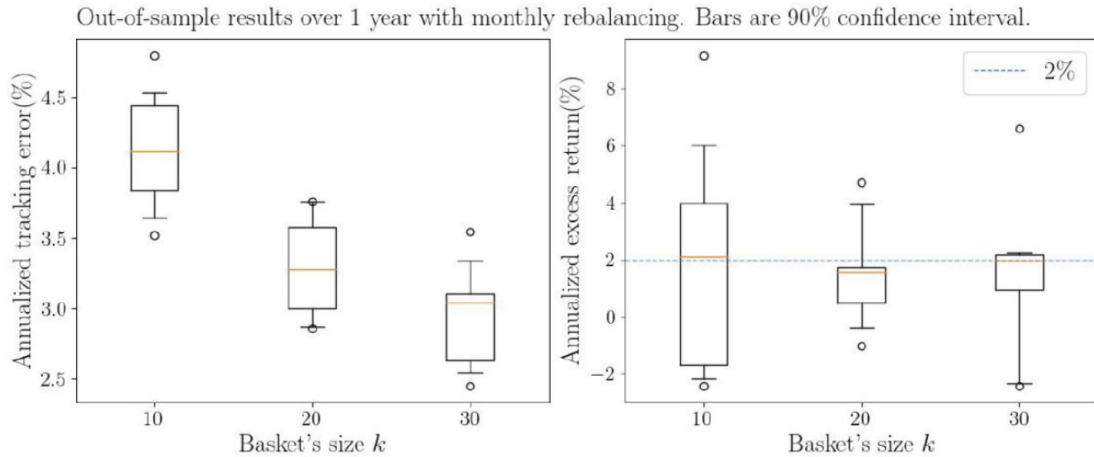


Figure 2.2: One-year out-of-sample results with monthly rebalancing for tracker portfolios of size  $k \in \{10, 20, 30\}$  and a minimum annual excess-return target of 2%. Source: Inspiration-Q empirical experiments (Inspiration-Q, 2025).

Table 7 records the corresponding summary statistics reported alongside the figure (Inspiration-Q, 2025). The key point for later sections is that (i) tracking improves predictably with basket size, and (ii) economically meaningful excess return can coexist with explicit tracking discipline, at the cost of higher sensitivity to forecasting and to constraint tuning.

Table 7: One-year OOS summary statistics under monthly rebalancing (reported in Inspiration-Q materials).

Basket size $k$	10	20	30
Avg. TE (%)	4.1	3.3	2.9
Avg. ER (%)	2.2	1.5	1.4

#### 2.4.4 Enhanced tracking: dividend-tilted Nasdaq-100 total return case study (2011–2024)

The next block distinguishes passive index tracking from enhanced tracking. Here the objective is to remain close to a benchmark (Nasdaq-100 Total Return) while systematically embedding a return target via a dividend tilt. The combinatorial burden is immediate: selecting a 50-asset basket from a 100-asset universe yields  $\binom{100}{50} \approx 10^{29}$  candidate baskets (Table 5), far beyond exhaustive enumeration.

Figure 3 presents cumulative performance over 2011–2024 for the 50-asset enhanced tracker relative to Nasdaq-100 Total Return (TR) and QQQ TR (Inspiration-Q,

2025). The optimization tracks Nasdaq-100 TR using ex-dividend price returns, while incorporating a systematic tilt toward higher expected dividend yields; realized dividends are reinvested in the portfolio.

Table 8 reproduces the mean annual return figures reported with the experiment (Inspiration-Q, 2025). The purpose of this case study is not to assert universal, persistent alpha, but to demonstrate feasibility: a large discrete design space can be searched efficiently enough to embed a structured tilt under explicit tracking and turnover constraints, producing a disciplined enhanced-tracking profile.

Table 8: Mean annual returns reported for 2011–2024 (Inspiration-Q enhanced-tracking experiment).

Index / strategy	Mean return (%)
Nasdaq-100 (Price)	17.01
Nasdaq-100 TR	18.5
QQQ TR	18.3
Enhanced tracker (50 assets)	20.1
Spread over QQQ TR	1.8

## 2.4.5 Synthesis

Across the three empirical blocks, the evidence supports a coherent message for institutional portfolio construction with discrete constraints. First, the cardinality-induced state space (Table 5) explains why exact MIQP approaches rapidly become operationally infeasible, motivating approximation methods that can deliver high-quality feasible solutions on realistic time scales (Table 6). Second, consistent with approximation–estimation perspectives (Belkin et al., 2019), increasing optimization effort can improve in-sample objectives without proportionate OOS gains once statistical noise dominates (Figure 1); this argues for reporting compute budgets and for interpreting incremental improvements through a generalization lens. Third, the passive and enhanced-tracking case studies show that hybrid discrete–continuous optimization can support both benchmark replication (Figure 2, Table 7) and structured active overlays under tracking discipline (Figure 3, Table 8) on classical hardware.

## 2.5 Conclusion

This chapter has examined portfolio construction through the lens of combinatorial optimization under realistic institutional constraints. While classical mean–variance optimization is convex and computationally tractable in its continuous formulation, the introduction of cardinality, regulatory, liquidity, turnover, and related structural constraints fundamentally alters the geometry of the feasible set. The resulting problems are, in general, mixed-integer quadratic programs and NP-hard.

The central observation is that in modern institutional environments the dominant source of computational difficulty is not the quadratic objective itself, but the discrete structure induced by practical constraints. This combinatorial structure

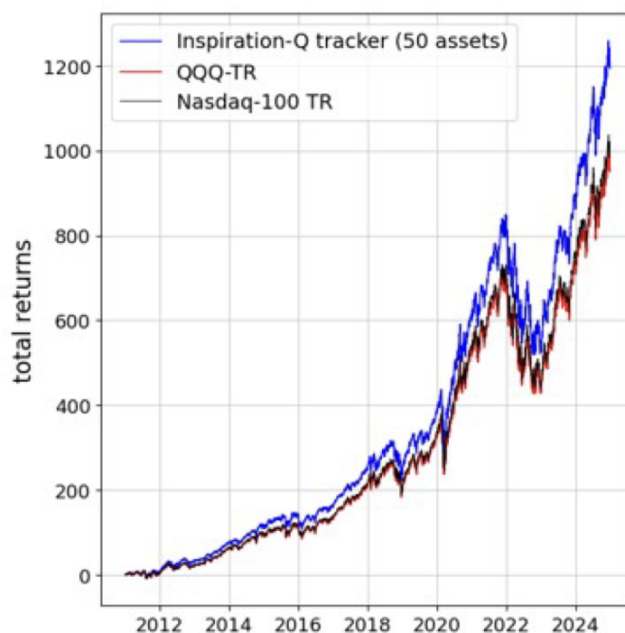


Figure 2.3: Tracking Nasdaq-100 TR using price returns (ex-dividends) with a dividend-targeting tilt, showing cumulative performance over 2011–2024. Source: Inspiration-Q empirical experiments (Inspiration-Q, 2025).

transforms portfolio construction from a purely statistical estimation problem into a joint estimation– optimization problem in which search quality plays a material role.

Hybrid simulated annealing and related quantum-inspired metaheuristics provide scalable mechanisms for exploring large discrete spaces without exhaustive enumeration. Empirical results indicate that such methods can achieve near-optimal objective values for cardinality-constrained index tracking problems at scales where exact mixed-integer solvers become computationally prohibitive.

At the same time, the analysis highlights an important structural limitation: increased optimization accuracy improves in-sample objective values monotonically, but out-of-sample gains saturate once optimization error falls below the statistical noise floor of covariance estimation. This empirical pattern is consistent with approximation–estimation trade-off theory and underscores that computational precision and statistical reliability must be balanced. Improving the solver alone cannot eliminate generalization risk inherent in financial data.

From an investment perspective, combinatorial flexibility enables the construction of passive and enhanced tracking portfolios that respect realistic institutional constraints while maintaining acceptable tracking precision. In enhanced configurations, structured tilts can be embedded within benchmark discipline without requiring relaxations that discard cardinality structure.

The contribution of quantum-inspired optimization in this context is therefore structural rather than speculative. It does not rely on asymptotic quantum speedups, nor does it presuppose access to fault-tolerant quantum hardware. Rather, it leverages algorithmic principles derived from statistical physics to address discrete optimization

problems that arise naturally in portfolio construction. More broadly, the findings suggest that in highly constrained financial systems, improvements in combinatorial exploration can reduce approximation error and operational inefficiencies, but cannot substitute for robust statistical modeling. Future research directions include integrating improved covariance estimation, dynamic models under regime shifts, and extending hybrid search methods to multi-period stochastic control problems.

In summary, as portfolio mandates become increasingly personalized and structurally constrained, the mathematical nature of the optimization problem shifts from convex allocation to discrete configuration. In this regime, scalable global search methods become a practically relevant component of the portfolio construction toolkit.

## References

John E. Beasley, Nigel Meade, and Tsu-Sheng Chang, “An evolutionary heuristic for the index tracking problem,” *European Journal of Operational Research*, 148(3), 621–643, 2003.

Mikhail Belkin, Daniel Hsu, Siyuan Ma, and Soumik Mandal, “Reconciling modern machine-learning practice and the classical bias–variance trade-off,” *Proceedings of the National Academy of Sciences*, 116(32), 15849–15854, 2019.

Michael J. Best and Robert R. Grauer, “On the sensitivity of mean–variance-efficient portfolios to changes in asset means,” *Review of Financial Studies*, 4(2), 315–342, 1991.

Daniel Bienstock, “Computational study of a family of mixed-integer quadratic programming problems,” *Mathematical Programming*, 74, 121–140, 1996.

Fischer Black and Robert Litterman, “Global portfolio optimization,” *Financial Analysts Journal*, 48(5), 28–43, 1992.

Stephen Boyd and Lieven Vandenberghe, *Convex Optimization*. Cambridge University Press, 2004.

Joshua Brodie, Ingrid Daubechies, Christine De Mol, Domenico Giannone, and Ignace Loris, “Sparse and stable Markowitz portfolios,” *Proceedings of the National Academy of Sciences*, 106(30), 12267–12272, 2009.

Mark M. Carhart, “On persistence in mutual fund performance,” *Journal of Finance*, 52(1), 57–82, 1997.

Cerulli Associates, “Managed accounts assets climb to \$10.7 trillion,” Press Release, 2022.

Vijay K. Chopra and William T. Ziemba, “The effect of errors in means, variances, and covariances on optimal portfolio choice,” *Journal of Portfolio Management*, 19(2), 6–11, 1993.

Victor DeMiguel, Lorenzo Garlappi, and Raman Uppal, “Optimal versus naive diversification: How inefficient is the  $1/n$  portfolio strategy?” *Review of Financial*

*Studies*, 22(5), 1915–1953, 2009.

Eugene F. Fama and Kenneth R. French, “Common risk factors in the returns on stocks and bonds,” *Journal of Financial Economics*, 33(1), 3–56, 1993.

Samuel Fernández-Lorenzo, Diego Porras, and Juan José García-Ripoll, “Hybrid quantum–classical optimization with cardinality constraints and applications to finance,” *Quantum Science and Technology*, 6(3), 034010, 2021.

Michael R. Garey and David S. Johnson, *Computers and Intractability: A Guide to the Theory of NP-Completeness*. Freeman, 1979.

Stuart Geman and Donald Geman, “Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 6(6), 721–741, 1984.

Donald Goldfarb and Garud Iyengar, “Robust portfolio selection problems,” *Mathematics of Operations Research*, 28(1), 1–38, 2003.

Bruce Hajek, “Cooling schedules for optimal annealing,” *Mathematics of Operations Research*, 13(2), 311–329, 1988.

Ravi Jagannathan and Tongshu Ma, “Risk reduction in large portfolios: Why imposing the wrong constraints helps,” *Journal of Finance*, 58(4), 1651–1684, 2003.

Michael C. Jensen, “The performance of mutual funds in the period 1945–1964,” *Journal of Finance*, 23(2), 389–416, 1968.

Scott Kirkpatrick, C. Daniel Gelatt, and Mario P. Vecchi, “Optimization by simulated annealing,” *Science*, 220(4598), 671–680, 1983.

Hiroshi Konno and Hiroaki Yamazaki, “Mean-absolute deviation portfolio optimization model and its applications,” *Management Science*, 37(5), 519–531, 1991.

Olivier Ledoit and Michael Wolf, “A well-conditioned estimator for large-dimensional covariance matrices,” *Journal of Multivariate Analysis*, 88(2), 365–411, 2004.

Burton G. Malkiel, “The efficient market hypothesis and its critics,” *Journal of Economic Perspectives*, 17(1), 59–82, 2003.

Harry Markowitz, “Portfolio selection,” *Journal of Finance*, 7(1), 77–91, 1952.

Richard Michaud, “The Markowitz optimization enigma: Is ‘optimized’ optimal?” *Financial Analysts Journal*, 45(1), 31–42, 1989.

Samuel Mugel, Carlos Kuchkovsky, Escolástico Sánchez, Samuel Fernández-Lorenzo, Jorge Luis-Hita, Enrique Lizaso, and Román Orús, “Dynamic portfolio optimization with real datasets using quantum processors and quantum-inspired tensor networks,” *Physical Review Research*, 4(1), 013006, 2022.

George L. Nemhauser and Laurence A. Wolsey, *Integer and Combinatorial Optimization*. Wiley, 1988.

John Preskill, “Quantum computing in the NISQ era and beyond,” *Quantum*, 2, 79, 2018.

R. Tyrrell Rockafellar and Stanislav Uryasev, “Optimization of conditional value-at-risk,” *Journal of Risk*, 2(3), 21–41, 2000.

Richard Roll, “A mean/variance analysis of tracking error,” *Journal of Portfolio Management*, 19(4), 13–22, 1993.

Álvaro Rubio-García, Samuel Fernández-Lorenzo, Juan José García-Ripoll, and Diego Porras, “Accurate solution of the index tracking problem with a hybrid simulated annealing algorithm,” *Physica A: Statistical Mechanics and its Applications*, 639, 129637, 2024.

Robert Tibshirani, “Regression shrinkage and selection via the lasso,” *Journal of the Royal Statistical Society: Series B*, 58(1), 267–288, 1996.

# Chapter 3

## Building the Quantum Workforce

*A personal journey from tech recruitment to human-scale readiness*

Amir Rassol

### 3.1 A Career at the Frontier of Technology

I'm a tech recruiter with 28 years of experience, but I didn't start in tech. I studied law and then moved into recruitment when it became clear law wasn't for me. I joined S3, one of the biggest tech recruiters in the world at the time, recruiting for relational databases when they were still relatively novel compared to mainframes.

What made me effective in those early years wasn't technical expertise in the traditional sense, but something more foundational: the ability to understand technology well enough to translate it. I would leverage my ability to understand technology to build the verticals for the new team members and then pass down those verticals so they could hit the ground running. That way I understood the technologies, what they were recruiting for, how it fitted into my structure, and what they were doing. It gave them a great launch pad, and it gave me a comprehensive view of how different technologies intersected and evolved.

I didn't do it generically. I built verticals. Data warehousing, Business Objects, RDBMS, Project Management — defined niches that mapped to real demand.

Across my career I've worked across nearly every aspect of technology recruitment. I've always been drawn to the frontier edge and have watched multiple waves arrive.

---

This chapter is a featured contribution to the monograph *A Portrait of Quantum Technologies in Finance*, edited by Oswaldo Zapata, PhD, and published by The Quantum Finance Boardroom (2026). All rights reserved by the chapter's author.

## 3.2 Data Science, AI, and Another Frontier

After building several high-performing teams, I decided set up on my own company, Sumner and Scott, focusing on data science, machine learning and AI because I could see another frontier forming. At that time, people were still asking “what is data science?” and my response was “this is going to change the world we live in.” The reason was obvious: the explosion of data made it clear that something significant was coming.

I started dealing with a whole host of companies ranging from defense through to retail, through to pharmaceuticals. COVID arrived and I started looking at, on a conscious level, data for good and AI for good. As a result, I started focusing on biotech companies, those that were looking to use data in positive ways and how to give back to society. Again: frontier tech, bright people, unclear pathways, real-world problems.

## 3.3 Entering Quantum

After working in data science and machine learning for several years, I started to hear about quantum computing properly—this was around 2022. And I knew that quantum was going to be something revolutionary and wanted to be a part of it. Being a part of something that is going to change the fabric of how we live is fascinating to me. And to be sitting there watching it happen is great, but being a part of it is even better.

I spent the next best part of 2 years just trying to understand concepts and looking to see how that’s going to grow and what kind of people are going to be needed.

Obviously, it was a bit of a fear going into this unknown. Is this technology going to come off or not? The thoughts were 25 years in the future, 15 years in the future, and putting all my eggs into this basket was a case of, “Okay, you know, this may not come off and it may not make me a living or a business if it’s going to be that long.” But still, keeping an eye on it, I thought yes, this will happen. And that’s getting to where we are now.

## 3.4 From Feasibility to Scale

An interesting point of fact was when I started looking at it, it was a case of “Can quantum computers even work?” and then it was all talking about error correction and detection and so on, and now it’s coming round to “How fast can you scale? How quickly can this grow?” That shift in conversation is significant. And that’s when the workforce conversation becomes urgent.

When I talk about quantum as the next frontier, I’m drawing on a long memory: the birth of the internet, the normalisation of mobile, streaming becoming mundane, data

science migrating from niche to default expectation and now AI. Each time, I've seen new markets appear, unexpected roles emerge, and organisations scramble—often not because they lack technology, but because they lack the people, structures, and language to make sense of it. Quantum feels like the next chapter in that same story.

### 3.5 The Physics Trap and the Translator Gap

In the early days of this transition, it became clear that the traditional tech recruitment rules did not apply. We are operating in a space where roles are largely undefined. One of the biggest fallacies I encountered was what I call the “Physics Trap”—the belief that every quantum hire must be a PhD with a deep understanding of qubits.

In many of the conversations I've had, especially with finance clients, the real bottleneck isn't “Do we have enough PhDs?” It's “Do we have anyone who can sit between the physicists, the quants, the risk teams, and the business sponsors and keep everyone pointed at the same problem?”

A real gap is the lack of translators—individuals who can bridge the gap.

### 3.6 Quantum Advancing Faster Than Organisational Readiness

What ties my work together in quantum isn't a single role, but there's a recurring gap which I keep running into. Through recruitment, I saw that organizations were asking for quantum talent without necessarily having the language to evaluate it. In advisory conversations, I saw teams investing in pilots without understanding how those capabilities would work in their workforce. And I saw people excluded from the conversation altogether because it was either too technical or because it was framed as something that would arrive fully formed one day.

Each of those things look different on the surface, but they're the same kind of problem. Quantum is advancing faster than organizational readiness.

### 3.7 Recruitment Limits and Team-Based Reality

Recruitment is where the gap shows up most visibly. You can see it in job specifications and misaligned expectations. When a financial institution posts a quantum role, they're often looking for someone who can do everything: understand the quantum computing stack, translate business problems into quantum algorithms, communicate with C-suite executives, and somehow also know the regulatory environment. This candidate doesn't exist in meaningful numbers.

What's needed is a team-based approach. You will need domain experts who under-

stand risk modelling or portfolio optimization, quantum-literate software engineers who can work with hybrid classical-quantum systems, and people who can move between the technical and business layers.

### 3.8 Workforce Advisory and Upskilling

Workforce advisory is where the longer-term strategy comes in. Organizations realize that hiring alone won't prepare them. They need capability that must be built gradually across the existing teams. This is where the real work happens: identifying which existing team members have the foundational skills to be upskilled, creating learning pathways that don't require PhD-level physics, and building internal cultures where quantum literacy is seen as a career asset rather than an exotic specialization.

Education and storytelling is the cultural layer. The language of quantum computing—superposition, entanglement, quantum advantage—can feel obscure to people who work in classical finance. This creates a barrier where perfectly capable professionals self-select out of the conversation because they assume they lack the prerequisites.

### 3.9 Finance as an Early Adopter

Finance is often an early adopter of frontier technologies. There's much more vested interest in doing that. Banks and financial institutions understand that being late to a technological shift can be existential. So there's an institutional nervousness about being left behind on quantum, even if the practical use cases are still being defined.

### 3.10 From Poaching to Workforce Strategy

A core realization of my work is that the traditional “poaching” model is broken. There simply isn't enough ready-made talent to go around, and trying to “import” a workforce is a losing game. This is why I have pivoted my focus toward workforce strategy and ecosystem building. The future is about capability mapping and internal mobility. This means taking existing domain experts—actuaries, risk managers, and software engineers—and upskilling them to be “quantum-literate”. I see education as infrastructure, not just marketing.

### 3.11 Upskilling in Practice

A lot of my work currently sits in the space between curiosity and commitment. Organisations recognise that quantum matters, but they are unsure how to engage without overreaching. One example of this has been recent workforce and capability

study with a large financial institution, focused not on technology selection, but on what quantum readiness looks like in practice.

What stands out in these conversations is that the first barrier is rarely the technology itself. It is organisational clarity. Teams want to understand where quantum might plausibly intersect with existing functions—risk, actuarial modelling, optimisation, data, and core technology—before they think about hiring or significant investment. The immediate question is not “who do we recruit?” but “who do we already have, and what would they need to know to engage intelligently?”

This is where upskilling becomes the sensible first move. Not to create quantum specialists overnight, but to build enough shared literacy that internal teams can evaluate use cases, ask better questions of vendors and research partners, and avoid mistaking activity for progress. In finance especially, this matters. The cost of getting ahead of yourself is real, but so is the cost of waiting too long.

The work typically focuses on capability mapping rather than headcount: identifying which role families are most exposed to potential quantum impact, what adjacent skills already exist internally, and where learning pathways can be layered without disrupting day-to-day delivery.

In my experience, this kind of upskilling is what will allow quantum adoption to move from abstract interest to grounded readiness, without pretending certainty exists.

### 3.12 My Current Work in the Quantum Ecosystem

My work now sits at the intersection of recruitment, advisory, and education. I help organisations clarify what quantum readiness means for them before they commit to hiring, and I help individuals understand how their existing skills might translate into this space.

On the organisational side, that means challenging assumptions, reframing job specifications, and designing talent strategies that recognise uncertainty rather than ignore it. On the individual side, it means demystifying the field and mapping realistic pathways into it.

What connects these threads is translation. I do not position myself as a physicist. I position myself as someone working in the gap between frontier technology and human systems — where capability, structure, and narrative determine whether potential becomes progress.

### 3.13 Education, Storytelling, and the Podcast

I started the podcast “Beyond the Bit: Quantum Pathways” when I entered quantum fully. I saw there was a missing layer in the ecosystem, a place where people could learn about quantum but without getting into the hardcore technicals. A lot of the

podcasts are talking with founders and quite heavy-duty scientists and they're going into the superposition, entanglement, and so on. I felt that that was great, but it doesn't help people who are potentially looking at the field, and we're going to need a lot of people into it.

So, I thought, "How do I help get this message out there and ultimately help me to learn as well". Speaking with candidates, hiring managers, and advisors, it brought up: How do people really get into this space? What backgrounds translate? And what does it feel like working in quantum? And how have people learned and what would they advise?

Rather than explaining quantum computing—because there's no way I could do that—I focused on the people around it: founders, researchers, ecosystem builders, all with the same framing: why they came into quantum, where they believe the field is heading, what advice they would give. The podcast functions like an informal workforce instruction infrastructure, trying to make quantum legible without diluting it, talking about career paths and so on, and also gives people an opportunity to represent themselves and their companies in the wider field.

### 3.14 Making Quantum Legible Without Selling Certainty

What I've learned through recruitment and the podcast and doing advisory work is that most people aren't asking how to understand quantum itself, they're asking for something more personal, practical. "Where do I fit? How real is this?" They want to know whether this is a field worth orienting into.

My role has increasingly been to help make quantum legible at that human level, to try and reduce confusion without selling certainty, to show people and organisations what roles are emerging and how they might position themselves within them. Each quantum company is going to need a lot of different functions within that. Now that comes from not just the quantum computing, the tech, the analysis, but all the other supporting roles.

Making quantum legible does not mean making quantum simple. The field is genuinely complex, and the path from here to scaled quantum advantage is uncertain. But complexity and uncertainty don't mean inaccessibility. They just mean that we need to be honest about what we know, what we don't know, and what skills translate into this space.

### 3.15 Open Questions and the Road Ahead

As I look forwards, I must be honest: there are still more questions than answers. At some point we will face massive talent bottlenecks, mismatched expectations, and a lack of alignment between what universities are teaching and what the market

actually requires. I don't yet know exactly how the talent ecosystem will mature, and that uncertainty is a necessary part of being at the frontier.

Success in the next decade won't be measured by who has the most hardware, but by who has built the most resilient, adaptable, and literate workforce.

### **3.16 Workforce Infrastructure as the Real Constraint**

The talent bottleneck in quantum is both immediate and structural. In the immediate term, there simply aren't enough people with the specific combinations of skills that organizations need right now. But the structural issue is deeper: we haven't built the educational pathways, the career ladders, or the community infrastructure that would allow the workforce to scale at the pace the technology requires.

What's needed is a more systemic approach—one that recognizes that workforce development is infrastructure, not just a hiring problem. This means collaboration between universities, industry, and governments to build clear pathways into the field.

### **3.17 Timelines, Direction, and Readiness**

One of the recurring challenges I see is a mismatch between the timelines that technology is operating on and the timelines that organizations and individuals need for planning. Quantum computing has been “five years away” for quite a long time, and that creates a credibility problem.

But here's what I've learned from watching other technologies mature: the timeline doesn't matter as much as the direction of travel. When I started working in data science, people were sceptical about whether machine learning would transform industries or remain an academic curiosity. The direction of travel was clear: more data was being collected, computational power was increasing, and the fundamental algorithms were getting better. That meant that even if the exact timeline was uncertain, the eventual impact was becoming inevitable.

I see quantum on a similar trajectory. We don't know exactly when scaled quantum advantage will be demonstrated, or which applications will prove most valuable. But the direction of travel is clear: quantum hardware is improving, error rates are decreasing, the software stack is maturing, and investment continues to flow into the field.

### **3.18 Conclusion: Recruiting Without a Playbook**

The quantum revolution will not be defined solely by qubits, coherence times, or algorithmic breakthroughs. It will be defined by whether institutions can build the

human systems required to engage with those advances responsibly and effectively.

Recruiting without a playbook means accepting ambiguity. It means helping organisations make decisions in conditions of partial information and helping individuals orient themselves toward a field that is still defining its contours. It requires humility about timelines and confidence about direction.

For me, that is the appeal as much as the challenge. I am not trying to predict exactly how quantum computing will unfold. I am working on the premise that when frontier technologies mature, the decisive constraint is rarely the science alone. It is whether the workforce, the education pathways, and the organisational structures were built early enough and well enough to meet it.

# Chapter 4

## The Quantum Finance Boardroom

### *Navigating the critical juncture of finance's quantum takeoff*

André Costa

#### 4.1 The Most Dangerous Moment: A Flight Plan for Financial Disruption

There's a principle in aviation that has always resonated with me: the most dangerous moment in flight isn't turbulence, a storm, or even engine failure. It's takeoff. It's that critical phase at low altitude and low speed where there is absolutely no margin for error. There's no room for ego, distraction, or assumptions. In aviation, we don't just hope for the best; we brief for emergencies before we even board the aircraft. We anticipate, we prepare, and we build contingency plans into our very DNA.

Business, and particularly the world of finance, should be no different. We are currently in the takeoff phase of the most profound technological shift of our generation: quantum computing. It's a low-altitude, low-speed moment where the right preparation will determine who soars and who stalls on the runway. For those of us managing wealth and risk, ignoring this isn't just a missed opportunity; it's a breach of our fundamental duty to our clients. This chapter is the story of how I, driven by a specific kind of curiosity, tried to build the pre-flight checklist for this journey, and how that led to the co-creation of The Quantum Finance Boardroom (TQFB).

---

This chapter is a featured contribution to the monograph *A Portrait of Quantum Technologies in Finance*, edited by Oswaldo Zapata, PhD, and published by The Quantum Finance Boardroom (2026). All rights reserved by the chapter's author.

## 4.2 From Curiosity to Structured Inquiry: The Genesis of a Quest

I'm curious by nature, but I've learned that curiosity without action is just day-dreaming. When quantum computing news started popping up everywhere, my initial reaction was a familiar one: "Oops, curious again." But this felt different. The implications weren't just abstract; they seemed to point directly at the heart of what I do.

The question that shifted me from passive interest to active pursuit was simple but loaded: "Where can I find reliable, technical, PhD-level information about this topic, specifically applied to finance?" Not just popular science articles, but the deep, technical stuff. More importantly, how could I learn to use this in my wealth management services? What, precisely, should I learn?

I quickly established a critical boundary for myself. I was not going to become a quantum physicist or a quantum mechanics expert—the world has very few of those, and for good reason. My goal was functional fluency, not theoretical mastery. The first step, therefore, was to find the source—the person who could bridge the chasm between the esoteric world of quantum mechanics and the pragmatic world of finance.

That's when I virtually met Dr. Oswaldo Zapata. It took some time to find the right person, someone who wasn't just brilliant but could also translate brilliance into practical terms. What followed were conversations, calls, and more calls. We weren't just talking; we were building a shared language. This collaboration led to our first Quantum Computing in Finance webinar. It was during that event that another right person appeared: William Kelly, CAIA. Together, the three of us delivered that first webinar, and the energy from it sparked the idea. The idea was simple: we needed a permanent space for this conversation. We needed The Quantum Finance Boardroom.

## 4.3 The Birth of TQFB: An Ecosystem for Translational Intelligence

TQFB was created exactly one year ago with a very specific objective: to be an ecosystem where Quantum Computing and Finance professionals could have a friendly environment to develop projects, businesses, and education. The core challenge we identified wasn't a lack of interest or even a lack of technology; it was a language barrier.

I'm not talking about English, Portuguese, Spanish, Italian, Dutch, German, Hindi, French, or the many other languages spoken by our members. I'm talking about the professional dialects that are often more divisive than geography. I'm talking about the languages of math, physics, ethics, law, venture capital, Value at Risk (VAR), conditional VAR (CVAR), risk management, and the list goes on. A physicist and a portfolio manager can sit in the same room and talk past each other for hours, using

the same words to mean entirely different things.

TQFB was designed to be the universal translator. It's a place where a theoretical physicist can explain the implications of quantum entanglement for security, and a risk manager can translate that into a new framework for evaluating systemic risk. It's where a venture capitalist can learn to ask the right questions of a quantum startup, and where a regulator can begin to understand the technology they must eventually govern. The goal is to learn these different languages well enough to apply the synthesized knowledge in our respective fields.

## 4.4 The Wealth Management Imperative: Beyond Alternative Investments

For me, the immediate application was clear: wealth management. Learning about this advanced matter isn't just an academic exercise; it's about positioning ourselves at the vanguard of alternative investments for our clients, students, or employers. The traditional 60/40 portfolio is already being challenged by new asset classes and strategies. Quantum computing does not represent a new asset class but a new paradigm for analyzing all asset classes.

Having the right tools to evaluate a Quantum Computing startup makes the connection between all stakeholders easy and efficient. It creates a bridge between VC Funds seeking the next big thing, Founders with world-changing technology, High-Net-Worth Individuals (HNWIs) and Family Offices looking for true diversification, Hardware companies building the infrastructure, Universities cultivating the talent, Regulators trying to ensure stability, and Financial Institutions that must adapt or die. This doesn't just create knowledge; it creates value, businesses, capital, and high-profile networks.

## 4.5 Expanding the Quantum Horizon: Beyond Computing

When I started this journey, I had no idea what PQC (Post-Quantum Cryptography) was. Quantum Networks, Quantum Sensing, Quantum Communications—these were terms I'd never heard of. The more I learned, the more I realized that “quantum” was a universe of technologies, not just one type of computer.

This opened up applications in fields I have a very genuine interest in, like energy, risk management, alternative investments, and even autonomous navigation systems. Spotting opportunities becomes almost easy when you have the right, high-quality information and, more importantly, the right people with you to process it. TQFB is not just a simple ecosystem; it's a complex, adaptive one. It's a place where you can find experts from everything from Theoretical Physics to ethics and philosophy. We can all learn from the best.

Consider the implications: Geopolitics and Oil markets being modeled with quantum-enhanced precision. Portfolio optimization algorithms that make today's models look like stone axes. Quantum sensing applications in the pharma industry that could revolutionize drug discovery timelines. And perhaps most urgently for finance, Quantum Networks protecting our institutions from the existential threat of Shor's algorithm to current encryption standards. This isn't science fiction; this is the engineering problem of our decade.

## 4.6 The Complexity Premium: Valuing Knowledge as an Asset

I've come to believe that learning is the best alternative investment ever. Your personal and professional valuation can only increase with each new skill added or each new tool in your toolbox. I call this the "complexity premium." In a world where alpha is getting harder to find, the ability to understand and operationalize complexity becomes a source of alpha in itself.

So, what new skills should the finance professional care about learning to stay ahead of the competition? It's not about becoming a coder or a physicist. It's about developing "quantum literacy." It's about understanding the principles of quantum algorithms enough to know when they might apply. It's about understanding the hardware roadmaps enough to know when quantum advantage might become realistic for a specific problem. It's about understanding PQC enough to know that your firm's cybersecurity strategy is critically obsolete.

Our clients have high expectations for us. They trust us to be guardians of their future. That fiduciary duty is not negotiable. It extends beyond picking the right stocks or bonds. It extends to understanding the technological tides that could lift or sink the very foundations of the financial system. Ignoring quantum computing is no longer an option.

## 4.7 The Human Element: People at the Core of the Quantum Revolution

At the end of the day, this is not about qubits or algorithms. It's about people. No matter how technical the subject gets, it's always about people. It's about the community of experts who are willing to share their knowledge. It's about the clients who trust us to navigate this complexity on their behalf.

The most important lesson I've learned through this entire process is to stay curious and willing to learn. But it's not enough to be curious alone. You must act on that curiosity. You must seek out the right people, ask the hard questions, and be willing to admit what you don't know. The co-creation of The Quantum Finance Boardroom was my answer to that call to action. It's an invitation to anyone else who feels

that pull of curiosity and the weight of responsibility to join the conversation. The quantum future is coming, and it's time we learn its language.



## Chapter 5

# Quantum Computing in Asset Management

### *A practitioner's assessment*

Carlos Arcila Barrera

I ran a commodity hedge fund in Chicago for over 8 years. Our approach was what some people now call quantamental, we combined rigorous quantitative modeling with deep fundamental research because, frankly, neither one alone was sufficient. The quantitative side included statistical methods, machine learning where it genuinely added predictive value, econometric models, and systematic signal construction. The fundamental side meant understanding physical supply chains, refinery economics, weather patterns, shipping routes, and the thousand small details that determine whether a spread is mispriced or just noisy. We also ran a separate strategy focused on environmental and carbon markets, which brought its own set of modeling challenges and data constraints.

I also carried a long-standing interest in quantum physics—not the popular-science version, but the actual formalism. The kind that makes you realize classical intuition is a remarkably convenient approximation your brain constructs to navigate daily life.

So when quantum computing started appearing more frequently in finance conversations around 2021–2022—first as footnotes in research notes, then as vendor pitch decks, then as LinkedIn posts from people who had never priced an option—I felt two things simultaneously: genuine curiosity and deep skepticism.

That tension is what pulled me in. Not the hype. The specific, uncomfortable possibility that behind all the noise, there might be something.

---

This chapter is a featured contribution to the monograph *A Portrait of Quantum Technologies in Finance*, edited by Oswaldo Zapata, PhD, and published by The Quantum Finance Boardroom (2026). All rights reserved by the chapter's author.

Every fund and trading operation I have worked at or compared notes with hits compute constraints at some point. The specific wall depends on what you trade, how you model it, and where your infrastructure was built to handle load. Modern hardware and well-optimized classical solvers handle many of these problems adequately, sometimes impressively. Supercomputers and cloud-scale infrastructure have pushed the boundary of what is tractable far beyond where it was a decade ago. That is worth acknowledging, because overstating the limitations of classical computing is just as dishonest as ignoring them.

But not every problem yields. Some formulations, once they reach production scale, either take so long that the answer arrives after the decision window has closed, or collapse entirely, solvers that fail to converge, matrices that become degenerate, approximations that break down exactly when precision matters most. The mathematics to handle these problems exists. What breaks is the ability to execute it within the time and stability constraints of a live operation.

At our fund, portfolio optimization was one example. The constraint stack was substantial: nonlinear transaction costs, variance-covariance matrices that became unstable as dimensionality grew, sector and factor exposure bounds, turnover limits, and a risk budget expressed as CVaR with scenario-based loss limits on top. The solver either took too long or forced us to simplify until the formulation no longer resembled the actual decision. Add cardinality constraints or multi-period rebalancing and classical solvers either slowed to a crawl or stopped converging.

Monte Carlo simulation compounded this, VaR, Expected Shortfall, stress testing, scenario generation. The trade-off was always accuracy against speed. Estimating tail risk with genuine confidence required either massive sample counts or variance reduction techniques that introduced their own assumptions. Models for more efficient sampling exist. Running it at the dimensionality real portfolios demand is where classical architectures strain.

Volatility calibration, tail risk modeling, and synthetic data generation each created similar pressure. Recalibrating stochastic volatility surfaces across a large book intraday, generating joint distributions that capture fat tails and correlation breakdowns during stress, producing synthetic datasets that preserve statistical structure without contaminating downstream models, these hit limits that had nothing to do with model quality and everything to do with what classical computation could deliver in time.

Exotic derivatives pricing added another layer. Path-dependent payoffs, early exercise features, high-dimensional underlyings, XVA calculations layering nested Monte Carlo on top of already expensive pricing models. I will not go deeper into technicalities, but these were problems I faced in practice.

Commodity markets carried their own complexity, physical supply chains, storage economics, weather-driven shocks, forward curves that behave nothing like interest rate term structures, spread relationships across geographies and delivery windows that create combinatorial problems most solvers struggled to handle. Our environmental and carbon strategies pushed further: structurally young markets, illiquid, exposed to regulatory regime changes and political risk that no correlation model

captures. These were problems we faced with our fund. Not all firms face the same ones. But institutional investors, hedge funds, prop trading desks, and asset managers each face some version, problems where the analytical frameworks exist but the computational demands for speed, accuracy, and stability at production scale exceed what classical architectures reliably deliver.

When I began reading the quantum computing research literature, peer-reviewed papers, not press releases, I noticed that the problem classes being targeted mapped directly onto constraints like these. Combinatorial optimization: portfolio construction, execution scheduling, capital allocation under discrete constraints. Monte Carlo acceleration: quantum amplitude estimation offering theoretical quadratic speedups for simulation tasks, a structural change in convergence speed, not a marginal improvement. Linear algebra at scale: solving the massive equation systems beneath risk modeling, factor analysis, and scenario generation. Sampling from complex distributions: generating scenarios that capture tail behavior, regime shifts, and non-Gaussian structure.

These were not analogies to our problems. They were our problems, expressed in the formalism of a different computational model. Much of this work remains academic, papers demonstrating theoretical advantage, not commercial systems running in production. But the specificity is what earned my attention. Researchers were not making vague claims about exponential speedup. They were targeting problem structures I recognized, problems where classical computation hits walls that better engineering alone does not solve.

## 5.1 Treating Quantum Like a Research Program, Not a Trend

After those readings, I talked to my team about what I was seeing. Instead of calling a vendor or hiring a consultant, I enrolled in an edX certificate program from the University of Chicago and approached quantum computing the way I would approach any unfamiliar domain that might eventually affect how we deploy capital: build the vocabulary first, the intuition second, the critical framework to evaluate claims third.

Structure mattered. If I had been left to browse papers on my own, I would have ended up with exactly the kind of selective understanding that leads to overconfidence. The program forced me to sit with fundamentals: what a qubit represents as a mathematical object, how quantum gates manipulate state vectors, why measurement collapses superposition, what entanglement actually means in operational terms rather than as a metaphor for spooky action.

After that I moved into MIT's quantum curriculum. Harder mathematics and more rigor. During this phase I consumed everything I could find on quantum algorithms applied to financial problems. JPMorgan's quantum research group has produced solid work. Goldman Sachs has contributions in literature. Academic teams at ETH Zürich, TU Delft, and several other institutions have gone deeper into the mathematical foundations. I chased citations the way you do when you are mapping

unfamiliar territory, not to confirm a thesis, but to understand the boundaries of what is actually known.

What I was building was not quantum expertise in the conventional sense. It was a calibration instrument, an internal detector tuned to distinguish, when someone made a claim about quantum advantage for finance, whether they were describing something achievable with near-term hardware or something that quietly assumes machines we will not have for a decade or more.

That distinction, more than any single algorithm or paper, is the most valuable thing I developed during this phase.

By mid-2024, I could walk through QAOA, VQE, quantum amplitude estimation, and the HHL algorithm for linear systems. I understood which problems had provable speedups and which relied on heuristic intuition. The theoretical landscape had become familiar terrain.

Then I tried to run something on actual hardware. That is when the real education started.

The distance between quantum algorithm research and deployable hardware is not a gap. It is a canyon. The algorithms that produce the most exciting headlines require fault-tolerant quantum computers — error-corrected machines with thousands or millions of physical qubits, capable of executing deep circuits without accumulated noise destroying the result. In practical terms, that means qubits stable enough that errors do not compound and corrupt the computation before it finishes. What exists today is different: Noisy Intermediate-Scale Quantum devices with dozens to a few hundred physical qubits and error rates that, in most realistic scenarios, outpace the useful computation you are trying to run.

This is where most executives disengage. They hear the hardware is not ready and file quantum computing somewhere between 2030 and never.

I understand that reaction. After eight years managing real capital, the instinct to walk away from technology that is not production-ready is well earned. But in this case, I think that instinct leads to the wrong conclusion.

My focus became twofold. First, understand and define solutions that could be deployed in the future when fault-tolerant hardware exists, positioning for what is coming rather than ignoring it. Second, and more immediate, identify problems that can be executed today with the hardware that is available now and that classical computing either cannot solve or solves poorly. Both tracks matter. The first is strategic. The second is operational. Treating quantum as purely a future problem misses that some of it is already here.

The path that made sense for my situation was D-Wave. They build quantum annealers, not gate-based quantum computers, different architecture, different physics, different problem class. These are purpose-built machines designed to find low-energy states of systems expressed as quadratic unconstrained binary optimization problems. They are not trying to be universal. They are trying to be useful for optimization now, with thousands of operational qubits.

I completed D-Wave’s training curriculum from foundational material through core programming, then built a working prototype with my team: portfolio optimization with Value-at-Risk constraints. Not the textbook Markowitz formulation that appears in every introductory presentation or the simple proofs of concept that populate most quantum demos. We wanted something closer to real — position limits, sector exposure bounds from actual investment mandates, turnover constraints, a risk budget. The goal was to test whether this hardware could handle a problem shaped like the ones we faced.

We formulated it as a Constrained Quadratic Model — CQM rather than QUBO, a technical distinction I will not belabor here but one that allowed us to better represent the actual structure of the problem. We ran it through hybrid solvers and benchmarked against commercial mixed-integer solvers. The criteria were straightforward: does it produce better risk-adjusted returns, is it faster, or both?

On performance, the results were similar, classical actually produced a slightly better Sharpe ratio and a few basis points higher compound return. No alpha from the quantum approach on this test. But on speed, something different emerged. As problem complexity scaled up, the quantum annealer found feasible solutions in regions of the search space that classical methods had not explored and solved the problem roughly 300 times faster. The solutions were not always better, but the speed and the evidence of a distinct approach to navigating the optimization landscape caught our attention.

The more important lesson was about workflow. Converting a classical problem into a formulation that can run on annealing hardware forced us to understand day-to-day operational details we had previously abstracted away. We also learned that practical quantum implementations today are not purely quantum, they are hybrid. Quantum processing handles specific subroutines; classical systems handle everything else. The productive question is not whether quantum can beat classical across the full pipeline. It is where in an existing workflow a quantum component might add value, even as one step among many.

We are now working on deeper problems where annealing and quantum approaches may provide real value, long-short strategies, statistical arbitrage, and others we are not yet ready to discuss.

In mid-2025, I attended the Chicago Quantum Forum expecting an audience of physicists and computer scientists. The reality was different, the room included policymakers, institutional investors, pharmaceutical researchers, utility executives, and defense technology professionals. The conversations and the agenda made something clear, quantum computing does not belong to financial services. It cuts across industries, government priorities, and research funding structures. Understanding that wider ecosystem — who funds the hardware development, who sets the policy framework, who controls the talent pipeline, became as important to my evaluation as understanding the algorithms themselves.

The network I built during this period proved more valuable than any course, paper, or vendor engagement. Practitioners who have already made mistakes teach faster than documentation. Conversations sharpen judgment about which hardware has

real advantages, which claims deserve attention, and which are marketing, and some problems cannot be solved in isolation, you need collaborators working on adjacent pieces.

Around the same time, I met Oswaldo Zapata, a theoretical physicist who was launching an initiative called *The Quantum Finance Boardroom (TQFB)*. The focus was applied work, what can actually be built, what has been tested, what failed and why. The membership included physicists and scientists but also practitioners, technologists, and professionals from different backgrounds. I joined because the conversations were grounded in reality rather than speculation, and because staying current on developments in this space is difficult to do alone. The network later became useful for work I had not yet anticipated.

## 5.2 Sigma Quantum Lab: Specializing Instead of Generalizing

By mid-2025, a pattern had emerged from three years of study, experimentation, and conversation. Near-term applications for quantum and quantum-inspired approaches in asset management exist, not theoretical possibilities, but problems that can be formulated and tested now. The hardware limitations are real, but they do not disqualify the work if the scope is defined honestly. Software ecosystems are maturing. Talent exists but remains scattered across universities, technology companies, and financial firms that often lack the structure to deploy it productively.

After two years of tests and prototypes, I founded Sigma Quantum Lab in 2025. The focus is quantum and quantum-inspired computation applied to asset management — problems I understood from direct experience, constraints I had dealt with for years and continue to deal with today.

What became clear through the work was that quantum approaches add value in niche, specific problems rather than general ones. Classical computing performs well on broad optimization and simulation tasks. But the problems I described earlier, volatility calibration at scale, synthetic data generation that preserves tail structure, probability distributions that capture regime shifts, derivatives pricing with path dependence, risk constraints in asset allocation, these sit in a different category. They are specific enough to formulate precisely, complex enough that classical methods strain, and valuable enough that solving them faster or more accurately matters operationally. That is where I decided to focus.

At the Sigma Quantum Lab we built partnerships with hardware companies. This surprised me initially since I expected protectiveness. The reality is that hardware vendors urgently need practitioners who bring authentic business problems. Their machines are valuable only if they solve something real. We bring structured problems and honest feedback on where the hardware works and where it does not. They provide access to cutting-edge systems, deep technical support, and early visibility into what is coming next. We are currently working on different solutions across multiple quantum hardware platforms. The more I engage, the more I learn — and

the more I recognize how much I still need to learn.

That is the story so far. A few observations for those evaluating this space:

**What works now:** Quantum annealing handles certain optimization workloads at production scale, combinatorial problems with discrete variables and layered constraints. Quantum-inspired classical algorithms deserve more attention than they receive; the research often produces classical variants that outperform traditional approaches without requiring quantum hardware. Hybrid workflows, quantum processing specific subroutines inside a classical pipeline, deliver practical value today. Gate-based quantum computers from IBM, Google, IonQ, and others are advancing but remain limited by qubit counts and error rates for most commercial applications. Their current value lies in research, algorithm development, and preparation rather than production deployment.

**What does not work yet:** Fault-tolerant speedups require error-corrected hardware that does not exist. Estimates range from five to twenty years. For example a realistic production-scale Monte Carlo acceleration has theoretical grounding but needs more qubits and lower error rates that can take years. Quantum machine learning has narrowly demonstrated advantage over well-tuned classical methods.

**Where the value sits:** Not in trying to apply quantum computing to everything. Classical systems handle general problems well. The value sits in niche, specific problems where classical methods strain, the volatility calibration, the tail risk modeling, the synthetic data generation, the constrained optimization I described earlier. Problems precise enough to formulate, complex enough that current compute struggles, and valuable enough that solving them faster or more accurately matters.

The path I took: Build the vocabulary first, the intuition second, the critical framework to evaluate claims third. That sequence worked for me. It may not be the only way, but it helped me distinguish what is real from what is marketing and, in this space, as in any new technology, the distance between those two is wide.

Quantum computing is a different computational approach. The timelines carry uncertainty. The hardware limitations are real. The edge in this space, in my view, will not come from being first to announce a quantum initiative or first to sign a vendor partnership. It will come from understanding which problems have the structure that quantum approaches can exploit — and having already done the work to formulate them before the hardware is ready to deploy. Those niche problems, the ones classical cannot solve today or where quantum can deliver faster or more accurate results, are specific to each organization and industry. That preparation takes years, not quarters. I started mine three years ago.



# Chapter 6

## Kernelizing Quantum Finance

### *Optimization as a structured assignment problem*

Chinonso Onah

#### 6.1 Introduction: Ambition Meets Misalignment

When I started working on quantum algorithms for combinatorial optimization problems, I didn't start out thinking about “feasible manifolds” or “problem geometry.” I started out debugging a toy project

I was working on a vehicle-matching problem and tried the most standard playbook in quantum optimization. I would encode everything as bits, add penalties for breaking the rules, and let the algorithm “learn” feasibility as it searches. On paper it looked reasonable and my small examples needed only a handful of qubits. But in practice, something felt wrong. Even when the circuit looked like it was improving the objective, the samples almost never represented valid matchings. I could get lots of bitstrings, but very few of them were answers I could actually use.

The reason is simple once you see it (Figure 1). The naive binary encoding makes the algorithm explore an enormous space, while the set of valid matchings is a tiny island inside it. Although penalties can make invalid choices expensive, they don't stop the algorithm from spending some of its time visiting invalid choices. For  $n=5$ , there are 33,554,432 bitstrings but only 120 valid matchings—about one usable sample in 280,000. At shallow depth, there just isn't enough “time” for probability to flow from the ocean of invalid states into the small region of valid ones. Any controllable mechanism to minimize these wasted efforts naturally leads to improvement.

The real question was: why am I searching largely irrelevant space in the first place?

---

This chapter is a featured contribution to the monograph *A Portrait of Quantum Technologies in Finance*, edited by Oswaldo Zapata, PhD, and published by The Quantum Finance Boardroom (2026). All rights reserved by the chapter's author.

What if I could avoid certain regions completely? What if the problem structure, i.e. geometry already tells you which regions you could avoid? That experience changed how I think about quantum optimization. I learned that by setting up the quantum search process to be structured in a way that closely aligns with the problem structure, the algorithm is able to avoid some or all of the unnecessary space of invalid solutions [1]. For such structured implementation, the resulting search space is still exponentially large but one now possesses physical mechanisms that can be leveraged to ensure exponentially fast probability transfer into the valid space [2] and obtain several exponential enhancements [3].

## The quantum curse of dimensionality

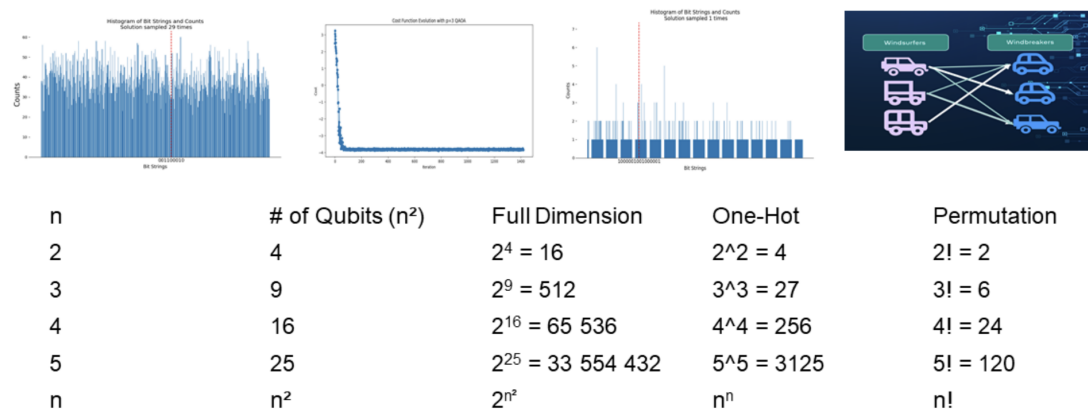


Figure 6.1: Reducing the curse of dimensionality in the quantum search space through structured configurations.

What surprised me later was that many finance problems have the same geometry I encountered in vehicle matching problems and can therefore benefit from the performance enhancements briefly mentioned in the preceding paragraph. In particular, the meaningful decisions live in a thin, structured subset defined by conservation laws. This chapter is a quantum-finance perspective on the shift from penalty enforcement toward designing representations and dynamics that live inside a smaller structured space and are equipped with mechanisms to steer the optimization to valid decisions even at shallow depth [3].

Why is quantum finance optimization interesting?

Portfolio construction, asset allocation, index tracking, and capital budgeting all appear, at first glance, to align naturally with the strengths of quantum algorithms. These problems are discrete, combinatorial, and high dimensional. They resist closed-form solutions and strain classical heuristics as the number of assets, constraints, and regulatory conditions grow. It is therefore unsurprising that quantum computing, with its promise of exploring large combinatorial spaces, has been repeatedly proposed as a transformative tool for financial decision making. [4–6]. In fact, over the past decade, finance has emerged as one of the most frequently cited application domains

for quantum optimization. [4–7].

Yet despite this enthusiasm, concrete progress has been limited. Empirical demonstrations remain small, fragile, and highly sensitive to modeling choices. Reported improvements often vanish under slight reformulations or disappear once feasibility constraints are enforced strictly. This gap between promise and performance has commonly been attributed to external limitations like insufficient qubit counts, hardware noise, shallow circuit depths, or immature error mitigation. While all of these challenges are real, focusing on them too early risks missing a more fundamental issue. The central difficulty in quantum finance to date has not been hardware capability, but problem-algorithm misalignment (See Figure 2).

### Allocation $\neq$ Selection

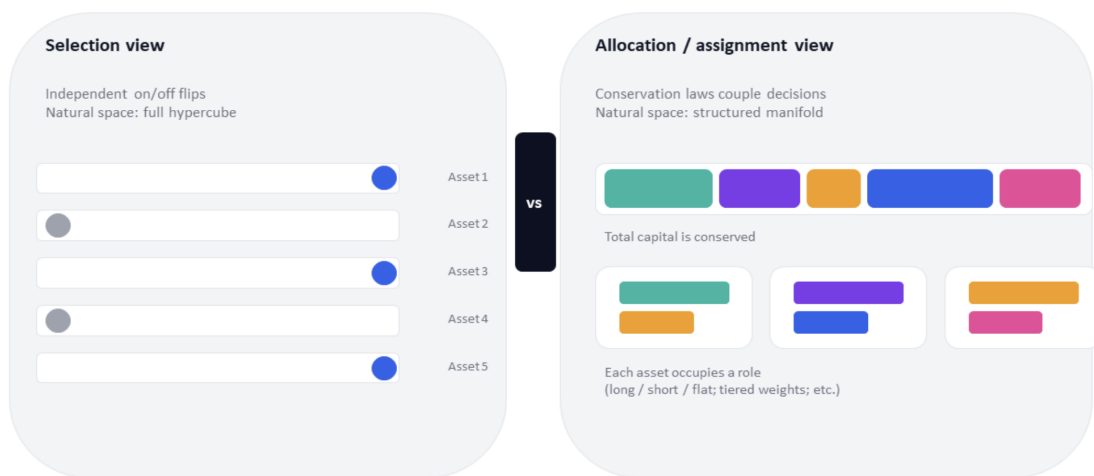


Figure 6.2: Many financial optimization problems are assignment problems with conservation laws. Treating their constraints only as generic unconstrained quadratic objectives destroys their geometry and misguides both classical and quantum solvers.

## 6.2 The Structural Misalignment

Financial decision problems are often NP-hard. For example, portfolio selection with cardinality constraints, sparse index replication, and discrete capital allocation across projects all exhibit combinatorial explosion as the number of assets or decision units increases. Moreover, these problems are typically solved repeatedly under changing market conditions, making fast approximate solutions practically valuable even when exact optimality is unnecessary.

Quantum algorithms, particularly variational quantum algorithms, appear well suited to this setting. They offer a framework in which a parametrized quantum state explores a large space of candidate solutions, while a classical outer loop steers the search toward lower-cost configurations. In many quantum-finance demonstrations, problems are expressed in a quadratic binary form (QUBO/Ising) because it is a

convenient way to implement and evaluate a diagonal cost Hamiltonian. Constraints such as budgets, cardinality limits, and exposure bounds are often added as penalty terms, and the resulting objective is optimized by a generic variational routine.

The issue is not the QUBO container itself. The issue is what happens when hard constraint structure is represented only as penalties, while the quantum dynamics is not engineered to reflect the structure that gave rise to the penalties. In that regime, feasibility becomes a statistically rare event in the sampled distribution, so depth and shots are spent rediscovering validity rather than comparing meaningful financial candidates.

In constrained problems, the feasible region is usually an exponentially small subset of a much larger space. The necessary problem structures are defined by the hard constraints every feasible solution must satisfy. The more the search space is reduced (while maintaining the full set of valid solutions), the better the algorithm is able to navigate the space using energetic signals to suppress infeasible configurations and move towards the feasible targets. The point here is that in addition to the penalties folded into the QUBO objective, the search space needs to be structured and reduced in line with hard constraints defining the problem structure. In financial optimization problems where constraints give rise to conservation laws and assignment structure, these requirements should be further reflected in how the quantum dynamics are engineered. These observations suggest that to make progress in quantum finance, one must first align the algorithmic framework with the natural structure of financial decision problems. This requires shifting attention away from penalty tuning and toward the geometry of feasibility itself.

### 6.3 Conservation Laws and Structured Financial Optimization

As already emphasized, at their core, most financial optimization tasks are about allocating conserved resources under coupled constraints. Capital, risk, exposure, and leverage obey accounting identities, regulatory limits, and normalization conditions that define the space of admissible decisions before any objective is evaluated. Consider portfolio construction. Even in its simplest discrete form, a portfolio is a normalized allocation of capital across assets where cardinality constraints restrict how many assets may be held. Budget constraints enforce that total investment equals available capital. Sector or factor constraints couple decisions across groups of assets.

The same structure appears in index tracking. Sparse replication requires selecting a small subset of assets whose combined exposure tracks a benchmark index. Here again, weights must sum to one, sector exposures must remain bounded, and tracking error is defined only on configurations that satisfy these normalization conditions. Capital allocation across projects or business units follows the same pattern. Funds are assigned in discrete tiers. Total capital is conserved. Minimum and maximum allocations impose hard bounds. Feasible allocations form a structured set defined by coupled constraints, not independent binary switches.

Much of the confusion in quantum finance formulations stems from treating allocation problems as selection problems. In a selection problem, one asks which elements to include or exclude. In an allocation problem, one asks how a conserved quantity should be distributed across options. These two problem classes behave very differently even though they both map to generic QUBO/Ising forms. Selection allows independent flips while allocation requires coordinated moves that preserve conservation.

Selection problems, such as MaxCut, naturally live in an unconstrained binary space. Each decision variable can flip independently. Constraints, if present, are often local or optional. Penalizing violations is usually sufficient to guide a heuristic search.

Allocation type problems do not share this property. Decisions are globally coupled through conservation laws. Changing one variable requires compensatory changes elsewhere. The feasible set is not a random subset of the hypercube, but a thin, highly structured manifold (a low-dimensional set of valid allocations) embedded within it. When the QUBO/Ising formulation of certain finance problems are solved without extra structure, the geometry defined by constraints is flattened away and all pieces of structural information are lost. The algorithm therefore explore the full hypercube. It is therefore trying to solve an allocation problem with a selection technique.

When allocation problems are forced into a selection framework, this structure is no longer reflected in the algorithm’s moves. The algorithm is allowed to explore configurations that violate conservation laws, only to be penalized energetically afterward. This mismatch becomes especially severe as problem size grows. The fraction of feasible configurations shrinks rapidly with the number of assets or allocation units, and sampling-based methods devote an overwhelming fraction of their effort to infeasible states, regardless of how penalties are tuned.

## 6.4 Kernelization of Quantum Finance

A more faithful abstraction is provided by assignment problems. In an assignment view, capital is assigned to assets, projects, or strategies according to discrete tiers or levels. Each assignment respects a conservation rule. Each asset or unit occupies exactly one slot in the allocation structure. This perspective aligns finance with a broad class of problems in operations research and combinatorial optimization like transportation, scheduling, and resource allocation that share this assignment structure [2]. Seen through this lens, many financial constraints cease to look like penalties and begin to look like symmetries. Budget conservation imposes an invariance under redistribution. Cardinality constraints impose fixed occupation numbers. Sector limits impose coupled assignment rules across groups of variables. These symmetries define the geometry of the feasible space. Thus, it is crucial that the quantum dynamics is engineered so that valid allocations are reached early and sampled often, rather than appearing as rare events.

To make this intuition precise, we borrow the notion of a kernel introduced in [2]. In that framework, “kernelization” means specifying, before optimization begins, the structured sector in which the algorithm is intended to do meaningful work. Formally,

[2] packages these choices into a kernel definition that fixes (i) an encoding that pins down a conserved structure (for instance, one-hot or fixed-occupation blocks), (ii) a decomposition of the diagonal cost into an objective term plus penalty terms, and (iii) Symmetries of the penalty terms respected by every valid solution, (iv) a mixer with a corresponding initial state whose induced moves respect the chosen conserved structure. This formalism basically adds a strong feasibility focus to the standard practice in problem algorithm codesign. A formal statement of this kernel notion appears as Definition 1 in [2].

In the finance setting, this “assignment view” provides the ingredients that a kernel needs: discrete roles (tiers/allocations), conservation rules (budgets/exposures), and symmetrylike invariances (redistribution within limits). Kernelization, in this sense, is the act of engineering the quantum dynamics so that it is compatible with these conservation and assignment symmetries, and can therefore drive probability mass toward valid allocations early—rather than relying on penalties alone to make feasibility show up as a rare event. We therefore introduce the concept of kernel alignment as a unifying lens for understanding algorithmic success and failure in financial optimization just like in other constrained optimization problems endowed with global constraints.

Every quantum optimization method implicitly defines a kernel, whether or not this is acknowledged explicitly. In unconstrained formulations, the effective kernel is often close to the full binary space, so the move set is largely blind to feasibility structure. All configurations are accessible, and constraints are enforced indirectly through penalties or post-processing. In constrained formulations, the kernel may be a strict subset of the full space, defined by conservation laws, fixed occupation numbers, or other structural restrictions. The critical point is that the kernel determines not only which solutions are possible, but how easily the algorithm can move between them. An algorithm may be powerful in an abstract sense, yet ineffective if its kernel is poorly matched to the problem’s structure. See figure 3 for an intuitive illustration.

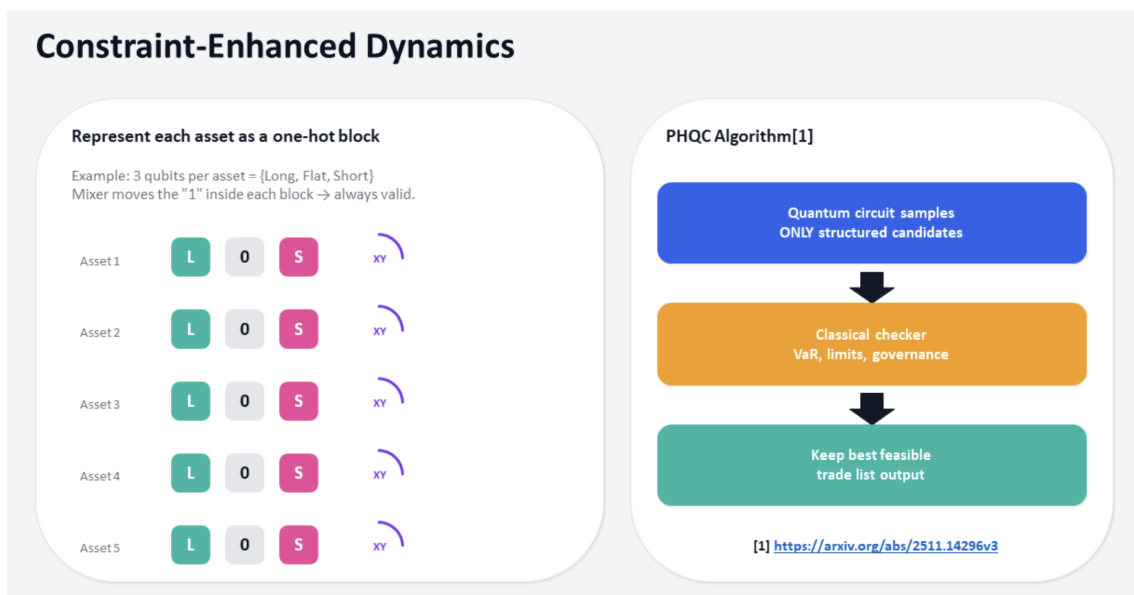


Figure 6.3: Illustration of structured circuit for long-short portfolio optimization.

Viewing optimization through the lens of kernel alignment shifts the design process. Instead of asking which algorithm to apply to a given problem, one asks how the problem should be represented so that the algorithm's natural dynamics operate in the right space. This principle applies beyond quantum computing. It explains why certain classical heuristics perform well on structured problems and poorly on others. It clarifies why penalty methods often struggle with tightly constrained tasks. It also provides a coherent framework for comparing algorithms that would otherwise appear incomparable.

## 6.5 The Cost of Ignoring Geometry

When optimization problems are flattened into unconstrained objectives, several pathologies emerge. First, feasibility becomes statistically rare. Even if an algorithm eventually identifies feasible configurations, the probability of sampling them at shallow depth is extremely low. This undermines the usefulness of sampling-based approaches, which are essential for risk analysis and scenario evaluation. Second, interpretability suffers. In finance and several industrial applications, infeasible samples carry no meaning. Post-processing steps must therefore discard the majority of outputs, wasting both quantum and classical resources. Third, algorithmic diagnostics become misleading. Performance metrics based on expected energy or average objective value conflate feasibility discovery with optimization quality. Apparent improvements may reflect better penalty satisfaction rather than better financial decisions. These issues are not unique to quantum algorithms. They appear whenever allocation problems are treated as unconstrained selection tasks. However, they are amplified in quantum settings

However, to fully understand the depth of the cost of ignoring the problem geometry, it is important to note that when feasibility depends only on penalty terms in an objective function, the problem's geometry is replaced by a scalar signal in form of coefficients. Energetic suppression is the only lever left to the algorithm and it attempts to discourage infeasible configurations by assigning them high cost but geometric exclusion which is more natural prevents the algorithm from entering those configurations at all. The set of feasible configurations is a structured region defined by coupled conditions. Budget constraints define hyperplanes. Cardinality constraints define fixed-occupation subspaces. Sector and exposure limits introduce correlations across variables. Together, these conditions carve out a constrained manifold of admissible solutions. This manifold has internal structure. Feasible configurations are related to one another by redistributions of capital, reallocations across tiers, or exchanges within sectors. Movement within this space preserves feasibility by construction. This is a powerful lever that is otherwise lost. The difference is not one of degree, but of kind.

As problem size grows, the feasible region occupies an increasingly small fraction of the ambient space, making the geometric point of view more and more important.

Reframing feasibility as geometry leads to a different way of thinking about algorithmic success. The primary object becomes the feasible mass produced in the

optimization process. An algorithm that produces many valid allocations, even if they are not optimal, is often more valuable than one that produces a single high-scoring configuration surrounded by invalid noise. This perspective aligns naturally with industrial workflows. Decision makers value diversity, robustness, and interpretability. They require candidate solutions that satisfy hard constraints by construction.

## 6.6 Conclusion: Toward a Structural Theory of Quantum Finance Optimization

We have conceptually introduced the notion of a kernel as a way to reason about how the problem is structured and encoded, where an algorithm operates, what transitions it permits, and which parts of the problem space it can meaningfully explore. It is defined by four elements. First, the set of states the algorithm can access. Second, the transitions it allows between those states. Third, the symmetries or invariances it preserves throughout its evolution. Fourth, the symmetries and invariances demanded by the valid solutions.

Once optimization is understood as part of a structured assignment pipeline, traditional performance metrics must be reconsidered. Metrics based on expected objective value or convergence to a minimum capture only a narrow aspect of algorithmic behavior. More relevant questions include how many feasible candidates are produced per unit of computational effort. How diverse are these candidates and how robust are they under scenario variations?

Constraint-enhanced approaches score well on these dimensions because they treat feasibility as a core target and are optimized as such. This shift in evaluation criteria aligns the generated solutions with the realities of industrial decision making in finance and other sectors.

The broader outlook on feasibility suggests the development of a structural theory of constrained quantum optimization algorithms. Such a theory would focus on the relationship between problem geometry, feasible manifolds, and algorithmic dynamics. It would seek to classify problems by the kernels they induce and to characterize which classes of algorithms are naturally aligned with which classes of problems.

Quantum algorithms enter this picture as particular dynamical systems with specific invariances and transition rules. And for each class of problems, understanding their capabilities and limitations requires placing them within this structural landscape. The framework for such a structural theory developed through the lens of problem algorithm codesign has been proposed in [2] and reported to yield measurable benefits in [3,8].

Future work in quantum finance optimization is needed to make this perspective precise including formal definitions of kernels, structural preservation, symmetries, and the resulting guarantees about sampling behavior and feasibility rates. Limitations of generic penalty only approaches can be sharpened into no-go results that apply across domains.

Whether quantum methods ultimately provide decisive advantages in finance remains an open question. What has become clear from our discussion so far is that when quantum optimization algorithms for finance are constructed to be aware of problem structure, their behavior becomes more interpretable and their evaluation more meaningful with additional physical mechanisms like fast probability transfer and feasibility guarantees already proved to arise elsewhere [3]. Crucially, these interesting benefits come from respecting structure and aligning algorithms with the geometry of the problems they are meant to solve. For problems where valid solutions form a structured manifold, feasible configurations are related to one another by redistributions. This is a powerful lever that should be used whenever available. In finance, these structures arise naturally and is further endowed by an internal structure with conservation laws.

## References

1. S. Hadfield, et al., “From the Quantum Approximate Optimization Algorithm to a Quantum Alternating Operator Ansatz.” <https://arxiv.org/abs/1709.03489> (2017).
2. C. Onah, R. Firt, and K. Michielsen, “Empirical Quantum Advantage in Constrained Optimization from Encoded Unitary Designs.” <https://doi.org/10.48550/arXiv.2511.14296> (2025).
3. C. Onah and K. Michielsen, “Fundamental Limitations of QAOA on Constrained Problems and a Route to Exponential Enhancement.” <https://arxiv.org/abs/2511.17259v1> (2025).
4. D. Herman et al., “Quantum computing for finance,” *Nature Reviews Physics* 5, 450–465 (2023). DOI: 10.1038/s42254-023-00603-1.
5. D. J. Egger, et al., “Quantum Computing for Finance: State-of-the-Art and Future Prospects,” *IEEE Transactions on Quantum Engineering* 1, 1–24 (2020). DOI: 10.1109/TQE.2020.3030314.
6. R. Orús, S. Mugel, and E. Lizaso, “Quantum computing for finance: Overview and prospects,” *Reviews in Physics* 4, 100028 (2019). DOI: 10.1016/j.revip.2019.100028.
7. D. Herman et al., “A Survey of Quantum Computing for Finance,” arXiv:2201.02773 (2022). DOI: 10.48550/arXiv.2201.02773.
8. C. Onah, F. Roman, and K. Michielsen, “Dataset: Empirical Quantum Advantage in Constrained Optimization” (2025).



# Chapter 7

## From Technical Correctness to Hybrid Pragmatism

*Notes from an early quantum founder*

David Isaac

### 7.1 Entering With the Wrong Assumption

When I first began working in quantum computing, I assumed that technical correctness would eventually translate into value. If a formulation was rigorous, if an algorithm exploited structure that classical methods could not, then adoption, funding, and relevance would follow with time. That belief was common. It was reasonable. It was also insufficient.

The gap became visible when the work encountered real financial constraints. The data was noisier than I expected. Regimes shifted faster than my models wanted to admit. Decision windows were defined by schedules and operational deadlines, not by solver runtimes. Tolerance for opaque reasoning was low—not because people were anti-quantum, but because accountability in finance is concrete. Someone has to sign off on the risk. That reality doesn’t just shape modeling choices; it shapes adoption. In finance, credibility is not a marketing layer. It is a form of capital you spend every time you ask an institution to trust a new component in a live decision system.

I was drawn to “pure” quantum approaches because they were clean. You could write down a Hamiltonian, justify an encoding, and call a solver. In isolation, that kind of work is satisfying. In practice, purity collided with everything outside the physics:

---

This chapter is a featured contribution to the monograph *A Portrait of Quantum Technologies in Finance*, edited by Oswaldo Zapata, PhD, and published by The Quantum Finance Boardroom (2026). All rights reserved by the chapter’s author.

embedding overhead, calibration cycles, queue time, constraints that do not survive discretization, and downstream workflows that demand stable, interpretable outputs.

As those frictions accumulated, the definition of success changed. The questions that mattered were no longer “Is the mapping faithful?” but “Does it produce a usable answer within the decision window?” and “Can we understand and control failure modes?” That shift led me toward hybrid systems—not as a compromise, but as a response to reality.

In finance, a solution that is slightly suboptimal but stable, timely, and explainable often beats one that is theoretically superior but fragile or late. That inversion forces a re-evaluation of what “advantage” means in practice. End-to-end improvements matter more than isolated algorithmic wins. The most valuable quantum contribution is often a targeted improvement to one hard step in a larger pipeline, not a replacement for the pipeline.

What proved hardest to let go of was not any specific tool or technique, but a set of beliefs that were technically defensible yet economically irrelevant: that correctness implies value, that purity implies superiority, and that adoption follows capability. What replaced them was a conditional view of quantum computing in finance—quantum as a selective component within hybrid systems, constrained by time, trust, and integration.

## 7.2 A Concrete Failure and Its Consequences

One of the earliest failures that forced a rethink was an optimization formulation I was convinced was “right.” On paper, it was clean. The objective was clear. The constraints were explicit. The mapping into a quantum-friendly form was defensible. I expected the limiting factor to be hardware scale or noise. Instead, the formulation itself became the bottleneck.

The first crack appeared upon contact with live data. Relationships that looked stable in controlled analysis shifted over short horizons. Small changes in inputs produced outsized changes in outputs. The solution was not wrong in a mathematical sense; it was fragile in an operational sense. In finance, that distinction matters. A model can be internally consistent and still be unusable if it does not degrade gracefully under non-stationarity and noise.

One failure was embarrassingly mundane: feasibility. We took a clean portfolio objective and encoded weights via a binary expansion—exactly the kind of “technically correct” mapping you can justify line-by-line. The budget condition was simple: weights sum to one. In the paper, it can be handled as a quadratic penalty; in our first implementation, we enforced it as a hard equality inside a CQM. The solver came back with the same message repeatedly: no feasible solutions. Nothing about the objective was exotic. The brittleness came from the interaction between a strict equality and weights constructed from binary bits. On paper, it’s a minor modeling choice. In practice, it’s the difference between “the formulation is correct” and “the system produces an answer at all.” That was the moment I started treating softness,

scaling, and fallbacks not as compromises, but as requirements.

My instinct was to tighten the formulation. That made it worse. More constraints increased brittleness. Enforcing exact relationships amplified sensitivity. The more “correct” the model became in theory, the less tolerant it was of the imperfections that dominate real systems. At the same time, explainability eroded. When outputs behaved counterintuitively, the burden was no longer “prove the math,” but “justify the decision,” quickly, to people who had to live with the consequences. In hindsight, this was the first time I felt the evaluation burden directly: it wasn’t enough to show that the solver improved an objective. I needed to show that the improvement transferred to behavior a risk team could live with.

The second failure mode was temporal. Even when the solver produced high-quality candidate solutions, they arrived at the wrong time. Queue delays. Calibration overhead. Post-processing latency. In environments with fixed decision schedules—or under stress—a late solution is functionally equivalent to no solution.

In practice, the order had to reverse. Robustness had to be designed in from the beginning. Approximation was not a concession; it was a prerequisite. Classical pre-processing became necessary to stabilize inputs and reduce dimensionality. Heuristics became necessary to shape the search space before invoking any quantum component. Post-processing became necessary to validate, interpret, and sometimes override solver outputs.

The success criterion changed. The relevant question stopped being “Does this formulation faithfully represent the problem?” and became “Does it behave predictably when the world stops cooperating?” Once that lens was adopted—robustness, latency, interpretability—the move toward hybrid systems ceased to be a preference. It became the only workable path.

## 7.3 What I Stopped Believing, and What Replaced It

After that shift, I stopped thinking in terms of “building a quantum solution.” I started thinking in terms of building something a finance team could actually use. That distinction sounds obvious. It was not obvious to me at the start.

The first belief I let go of was purity. I had treated “hybrid” as a temporary phase—something you tolerate until hardware improves and a fully quantum solution becomes viable. Over time, that framing collapsed. Hybrid systems are not a stepping stone. They are what emerges when you stop optimizing inside a toy model and start optimizing inside a real workflow.

The second belief I abandoned was that capability leads to adoption. In practice, people adopt what they can trust, integrate, monitor, and defend internally. The questions I began hearing were never “Is this quantum enough?” They were practical and procedural. Can we plug it into existing systems? Can we explain it to governance? What happens when it breaks? How do we roll back safely? Those

questions are not peripheral. They define the product.

This shift reshaped how I approached formulations. Early on, I wanted the most faithful model. Over time, I came to prefer models that were tolerant—models that do not fall apart when the data becomes ugly. That meant accepting approximation. It meant softening constraints when exactness produced brittleness. It meant scaling and discretizing in ways that behave better numerically. And it meant recognizing that classical filtering and heuristics are not a failure of imagination. They are the entry cost of workable systems.

I also stopped expecting quantum to replace classical systems. That is not the current reality. The more accurate mental model is that the quantum component functions as a targeted accelerant. Sometimes it helps generate stronger candidates. Sometimes it explores a search space differently than standard heuristics. Sometimes it delivers a better trade-off under a hard time limit. The objective is not to maximize “quantumness.” The objective is to justify why the quantum component belongs in the pipeline at all.

This is where my view of “advantage” became stricter. In theory, advantage is measured against a clean baseline. In practice, the baseline is a messy stack of heuristics and shortcuts that already work well enough—and have years of institutional trust behind them. If you cannot beat that baseline in a way that survives scrutiny, you do not have a product. You have a demo.

## **The architecture of pragmatism**

If we accept that hybrid systems are the standard, what does a production-ready architecture look like? In my experience, it looks like a “classical sandwich”: a robust classical layer before the quantum step, a minimal quantum kernel, and a classical layer after it that makes the output safe to use.

### **The Pre-Processing Layer (the shield)**

This layer’s job is to protect the quantum component from reality. It handles data quality issues, scaling, feature engineering, and it performs the most important step for tractability: dimensionality reduction. It may also include regime-aware logic, not because every workflow needs explicit regime detection, but because most financial problems are non-stationary. In practice, sending raw market data directly into a QPU or quantum-inspired solver is often a reliable way to waste compute and produce brittle outputs.

### **The Quantum Kernel (the spark)**

This is the smallest possible component that can justify being quantum at all. It does not “solve the problem.” It generates a distribution of candidate solutions for a narrow sub-problem—usually a combinatorial optimization step—where classical heuristics are likely to get trapped or to deliver poor diversity under time constraints.

The kernel is not the workflow. It is an accelerant inside the workflow.

### The Post-Processing Layer (the filter)

This layer interprets and validates the output. It checks constraints that were too complex to encode. It tests stability under perturbations. It attaches diagnostics so downstream users can see when the quantum step is behaving unusually. And it has a “kill switch”—an automatic fallback when latency or validity thresholds aren’t met: if the quantum solution is late, invalid, or unstable, the system defaults to a known classical heuristic so business continuity does not depend on the QPU behaving perfectly.

This architecture is less elegant than a clean end-to-end quantum story. It is also closer to what survives contact with production.

## 7.4 The Hidden Tax of Being Early

There is a comforting version of this story in which the only obstacle is technology: qubit counts, noise, error rates. Improve the hardware and everything else will follow. That narrative makes progress feel linear and inevitable. It also ignores a harder reality. Being early carries costs that do not appear in technical roadmaps. Most of these costs show up earlier than founders expect—often as “why aren’t they responding?” moments, or as governance friction that looks like skepticism but is really due diligence.

**The Cost of Credibility Capital.** A quantum startup must establish trust before it can plausibly sell outcomes. In mature categories, you can lead with case studies, benchmarks, and standard implementation patterns. In quantum finance, much of that infrastructure is still forming. Time is spent building narratives to bridge gaps you did not create—why this is not vaporware, why it is not “just classical,” and how it fits into existing governance and risk processes.

**The Cost of a Moving Substrate.** Tooling changes. APIs change. Access models change. The practical shape of what can be run shifts underneath you. Even when the mathematics is stable, the operational details are not. Formulations and workflows must be adapted continuously—not out of indecision, but because the platform itself is evolving.

**The Cost of the Evaluation Burden.** In finance, it is not sufficient to claim an improvement in an objective function. You must show that the improvement transfers to something that matters: stability, risk metrics, live performance, or operational KPIs. The bar is high because the environment is adversarial and noisy. Honest progress is difficult to demonstrate cleanly, which creates pressure to overclaim.

**The Cost of Feedback Latency.** You can do everything correctly and still receive no signal for long stretches. No replies. No closed loops. No validation. In other markets, iteration on distribution is fast. In this space, a single missing dependency—hardware access, data permissions, or a governance review—can stall

work for weeks.

**The Cost of the Comparison Trap.** You are measured against classical systems that are deeply optimized and battle-tested, and against quantum hype you do not endorse. You must earn trust from both sides simultaneously. Sound too excited and you are dismissed as hype. Sound too careful and you are dismissed as irrelevant. Navigating that narrow channel is harder than most people expect.

This is where the hybrid worldview becomes more than a technical choice. It becomes a survival strategy. It enforces focus on what can be controlled: problem framing that fits current constraints, evaluation workflows that are transparent and fair, and systems that degrade gracefully under stress.

## 7.5 The Tensions That Do Not Resolve Cleanly

Adopting a hybrid-first approach clarifies the work, but it does not remove trade-offs.

The first tension is between scientific rigor and commercial survival. Rigor demands careful baselines, controlled experiments, and conservative claims. Survival demands speed: something a customer can touch, scope narrow enough to deliver, progress visible quickly enough to sustain engagement. Lean too far in either direction and the work fails for different reasons.

The second tension is between generality and specificity. Quantum finance discussions often operate at high altitude—optimization, risk, allocation—because the promise is easiest to describe there. Adoption happens at ground level. Products succeed because they solve a specific workflow problem: a constraint that breaks existing solvers, a time limit that forces a different approach, or a decision process that needs better candidates.

The third tension is between novelty and trust. Finance does not reward novelty for its own sake. It rewards methods that survive stress, audit, and repetition. As a result, the most valuable quantum contribution is sometimes the least visible one: a component that improves outcomes without demanding that end users adopt new theory as a prerequisite.

The fourth tension is between near-term value and long-term positioning. Focus exclusively on near-term deliverables and you risk building a consulting business that never becomes a platform. Focus exclusively on long-term advantage and you risk becoming a research project waiting for external timelines to align. The uncomfortable middle path is to ship useful hybrid capabilities now while deliberately designing them to absorb future improvements in hardware.

These tensions pushed me toward a set of operating principles.

First: Measure value where the decision is made, not where the algorithm runs. If it does not improve an end-to-end decision process, it is not a product improvement.

Second: Design around failure modes. Non-stationarity, noise, and time pressure are not edge cases in finance. They are the environment.

Third: Make the quantum contribution legible. Not every detail must be explained, but stakeholders must understand its role, what it improves, and how it can be monitored when it misbehaves.

Finally: Protect credibility. In an early ecosystem, credibility compounds slowly and evaporates quickly. Ambition is necessary. Precision is non-negotiable.

In practice, this is how I navigated those tensions, imperfectly. I time-boxed rigor rather than abandoning it: short, explicit baselines first, then incremental complexity only when the baseline was beaten end-to-end. I forced specificity by choosing one bottleneck per workflow—the step that was both painful for classical heuristics and easy to measure under a real deadline. I treated novelty as a liability until proven otherwise, which meant making the quantum step as small and replaceable as possible, with clear diagnostics and a default fallback. And when the temptation was to ship consulting forever, I tried to productize the reusable parts—the evaluation harness, the constraint checks, the rollback logic—because those are what compound across projects even when hardware timelines don't.

## **A checklist for the buyer (and the builder)**

To navigate these tensions, I developed a checklist for evaluating any proposed quantum financial workflow. If you cannot answer “Yes” to most of these, the project is more likely to be an experiment than a product.

### **Does the classical baseline exist?**

Never compare QPU performance against “zero.” Compare it against the best heuristic currently in production.

### **Is time-to-solution inclusive?**

Does the metric include queueing, cloud latency, solver overhead, and post-processing—not just QPU time?

### **Is the failure mode deterministic?**

If the quantum step fails, does the system output a safe fallback, or does it crash or block a decision?

### **Is the metric aligned with the decision?**

Does it improve the KPI that matters at the decision point—tracking error, CVaR, turnover, slippage sensitivity—rather than an abstract objective?

### **Is the pipeline auditable and reproducible?**

Can you version data, code, and parameters, rerun the workflow, and explain why outputs changed?

### **Does the quantum step have a one-sentence job description?**

Not “it's faster,” but “it generates diverse candidates for X under Y constraint,” or “it improves exploration under a hard time limit.”

### **Is constraint handling honest?**

What is encoded in the model, what is enforced downstream, and what is simply assumed? If that split is unclear, governance will treat it as risk.

A checklist like this does not guarantee success. It does prevent a common failure: building a technically impressive workflow that cannot survive contact with real decision-making.

## **7.6 Closing**

The assumption I started with—that technical correctness would eventually translate into value—was not wrong in theory. It was wrong in sequence.

In quantum finance, the real work is not proving that a quantum method can solve a stylized problem. It is building decision systems that remain useful when the world is noisy, time-constrained, and accountable. Hybrid approaches are not a waiting room for the future. They are the architecture of the present.

If genuine quantum advantage arrives, it will not arrive as a tidal wave that washes away classical finance. It will arrive as a distinct signal inside a noisy stack. The teams that benefit will be the ones who built pipelines that can capture that signal, validate it under governance, and operationalize it inside real decision windows.

# Chapter 8

## A Road Less Traveled

### *How a lawyer ended up working in quantum computing*

Francisco Castro

When you set out on a journey with what at the time was a clear and short-term destination, it's almost impossible to see where the road beyond your immediate goals might lead. I certainly never imagined that I would become involved in quantum computing, let alone help a small, academic spinoff grow into a leading startup in the quantum computing industry.

At the outset, I didn't even know the field existed in any meaningful way. Along the way, I pursued a range of interests that seemed unlikely to intersect with quantum computing at all. What follows is the story of how I eventually arrived at IonQ, the leading and first publicly traded pure-play quantum computing company, and what I have been working on since leaving IonQ, including my current company, Bluzec.

### 8.1 The Start of an Unexpected Journey

If my connection to quantum computing has a beginning, it might be traced back nearly 30 years to my time as a graduate student at Drexel University, where I was working toward a Ph.D. in Electrical Engineering with a focus on applied semiconductor physics. My research area was novel structures and materials for high-speed optical communications. The work was heavily mathematical and computational, aimed at understanding how specific design and processing choices affected system-level performance of photonic devices. Moreover, and unsurprisingly, the theoretical foundations of my work relied heavily on quantum mechanics, electromagnetics,

---

This chapter is a featured contribution to the monograph *A Portrait of Quantum Technologies in Finance*, edited by Oswaldo Zapata, PhD, and published by The Quantum Finance Boardroom (2026). All rights reserved by the chapter's author.

optics, and material science, as well as a heavy dose of laboratory experiments to corroborate the theory.

I found the work deeply engaging and was fortunate to have a doctoral advisor who was both demanding and inspiring, pushing me to understand every aspect of the research with meticulous attention to detail. His daily challenge to me and my lab mates was always “what was your contribution today?” – a not-so-subtle way of reminding us of his lofty expectations to produce meaningful and impactful results.

In the mid-to-late 1990s, quantum computing was far from mainstream, even within academic circles, so it was not something I was remotely aware of at the time. It’s worth remembering that it wasn’t until late 1995 that Dr. Chris Monroe, later a co-founder of IonQ, published the first experimental demonstration of a quantum logic gate on any physical platform, using trapped ions as qubits. That work took place under Nobel Laureate Dr. David Wineland at the National Institute of Standards and Technology (NIST) in Boulder, Colorado, who would later become a technical advisor for IonQ.

At the time of my doctoral studies, my research work, and that of the broader photonics community, was motivated not by quantum computing since that was a niche, generally unknown field, but by the rapid scaling of optical networks to support the early growth of the Internet.

## 8.2 Getting Industrial Experience

Just before completing my PhD, I was offered an extraordinary opportunity to begin my professional career as a research engineer at Lucent Technologies. Lucent had emerged from the 1996 divestiture of AT&T Technologies and included the storied legacies of Western Electric and Bell Labs. Once on board, one of my core projects involved the development and testing of optical-based switching systems, a technical area that would later resonate with IonQ co-founder Dr. Jungsang Kim. More on that below.

After several years at Lucent, I joined Motorola Labs, Motorola’s research arm. In my new role I joined a team working on integrated image-processing systems. These were complex systems that combined all hardware and software aspects of the image-processing chain, from image capture, image and color processing, and ultimately image display. This new role also gave me the opportunity to collaborate closely with the semiconductor fabrication division responsible for making and testing our designs. This move marked a turning point in my career, a change that would play a significant role in eventually being part of the quantum computing industry, though I didn’t realize it at the time.

My team at Motorola was highly innovative, developing many inventions spanning both hardware architecture and software applications. It is no surprise that our research work produced a steady stream of patent filings. As one of the more junior team members, I was asked to coordinate closely with the attorneys handling our patent applications. If you think this was an effort to mentor me on the various

responsibilities of corporate research life, it wasn't, it was delegation. Senior engineers disliked dealing with administrative tasks, but they disliked dealing with lawyers even more. I, on the other hand, didn't. I enjoyed translating complex technology into something the legal team could understand and use in preparing the patent applications and subsequently handling the examination process with the US Patent and Trademark Office. Somehow it came natural to me to serve as the liaison between the research and legal teams.

### 8.3 From Engineer to Lawyer

By coincidence, or perhaps by pure luck, my lab was in a small building that sat directly across from Motorola's corporate tower, home to senior executives and the in-house legal department. At the time, Motorola handled patent prosecution internally and had a large number of patent attorneys working at the corporate tower. For me, that meant that I could easily have frequent in-person interactions with the patent attorneys working on my team's patent filings. Over time, I became increasingly curious about their work and started asking many questions about their daily tasks, the range of technologies they got to work on, and the opportunities that such field provided in terms of career development.

Well, after several years of managing my team's patent filings and countless conversations with the legal team, I made the decision to return to school and pursue a law degree, with the goal of specializing in technology law. It was an unexpected career turn but one that I was excited about and fully committed to do.

One more thing about my time at Motorola, I also met an exceptional engineer in the semiconductor division, one of the best engineers I've ever worked with. He would later move on to an international semiconductor research consortium and then to academic posts at Stanford and UC Berkeley. Our paths would cross again in an unexpected way years later.

### 8.4 The Start of My Legal Career

Surprisingly, my legal career began only a few months into law school. In the first few weeks of school, I was sitting on a torts class when a classmate made a comment about the use of technical experts' advice in certain cases. Curious about why he knew about what seemed like an obscure topic, I spoke with him afterward and learned he worked at the best patent boutique in town. A patent boutique is a law firm that specializes in patent law. He told me his firm was looking for engineers to support a rapidly growing new client and that they provided all the training needed to do the work.

I applied and was hired. Just months after starting law school, I was working full-time at a top patent firm. Balancing law school with demanding legal work was exhausting, but the compressed learning experience that I gained by doing both at the same time was invaluable.

That new client turned out to be another quiet connection to my future at IonQ. The company was developing wireless communications hardware and embedded software, and they needed to rapidly build their patent portfolio to compete with industry giants in what was a very crowded space. The law firm developed a highly specialized approach for managing a high-volume of filings without sacrificing quality, and I was deeply involved throughout the process, an experience that would prove critical later on.

Before finishing law school, I had decided that I was going to move to the Washington, DC, area. I was fortunate that a major opportunity presented itself when a leading national law firm based in Palo Alto, California, recruited me to join their office in Northern Virginia, just outside Washington, DC. This law firm was known as a go-to firm for startups in Silicon Valley and across the world, and it allowed me to get first-hand experience into the startup ecosystem and exposed me to a wide range of fascinating clients and projects.

One unexpected outcome from working at this law firm was a lasting relationship with one of their clients, a successful serial entrepreneur. It happens that we lived in the same town and, after unexpectedly running into each other a few times, we began meeting occasionally for coffee or lunch. He had broad interests and the freedom to pursue them. This relationship, too, would matter later.

I left that law firm, and over time, I moved between different law firms, each expanding my practice in new ways.

## 8.5 Serendipity

Eventually, I joined a law firm based in downtown Washington, DC, which was a perfect fit for my practice and client base. The firm's flexibility allowed me to structure arrangements that helped early-stage startups access high-quality legal services with a cost structure that was suitable for their funding levels.

It was there that I received a call from the serial entrepreneur I had stayed in touch with over the years. He was advising a Maryland-based startup developing software to control specialized hardware and needed a patent attorney fluent in physics, photonics, and quantum mechanics. My background fit the bill.

He arranged a meeting for me to talk to the co-founders and their small technical team at their headquarters in College Park, Maryland. The place was not that far from my office, and I just took the Metro there.

The startup was IonQ. The co-founders were Chris Monroe and Jungsang Kim, which I have mentioned above. Chris was at the University of Maryland, College Park, and Jungsang was at Duke University. IonQ was essentially a spinoff based on intellectual property licensed from both schools.

At the time, the entire company could fit into a small conference room. We spoke at length about technology and about building a patent portfolio that would set the industry's standard. My prior experience building patent portfolios for rapidly

growing companies proved immediately relevant. Moreover, Jungsang was quite impressed at my experience at Lucent working on optical-based switching systems as he had worked on similar projects during his time at Bell Labs and prior to joining Duke. A few days after the meeting I received a call from Jungsang and was hired as IonQ's outside patent counsel.

## 8.6 Becoming IonQ's First In-House Lawyer

Over the next 3 years, I worked closely with the core technical team at IonQ to identify technology for protection. Given my proximity to their offices, I regularly visited their facilities and had meetings with various members of the team. The patent portfolio began to grow rapidly and strategically.

IonQ was a growing client but was not my biggest client. I had a lot of work in California that required me to travel constantly. My law firm encouraged me to move to the San Francisco Bay Area to better serve my West Coast clients. It was a unique opportunity and I considered it thoroughly. After much deliberation and discussions with my family, we decided that it would be a good move and went ahead with it, knowing quite well that if it didn't work out, I was given the chance to return to the Washington, DC office.

After a few months in California my family and I were starting to feel settled when I got an unexpected call from Jungsang to tell me that he was too busy to coordinate the legal work internally and was calling to offer me a job with the company. This was completely unexpected and I needed time to digest.

Keep in mind, I had a thriving and growing legal practice that had taken me years to build. Moreover, I had just taken a big step moving my family across the country, could I make a wild bet to join a quantum computing startup? After talking to the then CEO of IonQ, Peter Chapman, and another round of deliberations with my wife, I decided to accept the offer. I was more afraid of regretting not taking the opportunity than taking it and failing miserably. Again, I had the backing of my law firm. They said to go ahead and try it out, and if it didn't work out, I could always come back to the firm. And that is how I went from being a research engineer, to being a lawyer, to being IonQ's outside patent counsel, and finally to being IonQ's first in-house lawyer.

Oh, and that great engineer from the semiconductor division that I met while at Motorola, he was now leading a seminar covering quantum computing at Stanford University. When I told him about my offer and my decision to take it, he simply said that it was a new, upcoming industry but that he would not bet his career on it. Not exactly words of encouragement, but a sober view of what I was really getting myself into.

## 8.7 The Challenges of a Quantum Startup

IonQ went from a small startup to a huge company in the 5 years I worked for the company. We endured COVID, an IPO and visit to the NYSE to ring the bell, short seller reports, shareholder lawsuits, a fast-growing workforce, new facilities, international expansion, many changes at the executive level, ever evolving technology roadmaps, growing competition, funding rounds, mergers and acquisitions, and several company-wide events to promote common cultural and corporate values, although I missed the one in Disney World.

For me, I grew the legal team and welcome our first general counsel, well three to be exact, they didn't last long for some reason. Expanded our patent efforts ferociously to the point that by the time of my departure we had the largest and most robust patent portfolio in the industry, one of the promises I had made to Chris and Jungsang in our original meeting at the College Park headquarters.

I also got involved in all sorts of things that you typically do not get to see in a law firm, such as product and branding strategy, US and international leases, employment and contracting issues, M&A due diligence, government and commercial contracts, technology export controls, compliance and data privacy issues, and many other issues you only get to experience in a fast-growing startup. After five furious and relentless years at IonQ, I decided to take some time off and figure out what my next step would be.

## 8.8 The Harvard Connection and the Origins of Bluzec

A few years earlier, I was contacted to participate as one of several instructors for a summer class at Harvard Law School covering diverse topics related to AI, blockchain, and quantum computing. The person planning the event, a young and highly energetic lawyer from Madrid, had reached out after researching who were the top lawyers at IonQ since he was looking for someone from a leading quantum computing company to talk about legal issues facing the industry. It was a great opportunity, so I immediately signed up.

Turns out that most attendees were from Europe and quite a number of them were involved in some shape or form with quantum computing. It became a great place to meet a different group of people, all with a common interest in the technology and industry.

I have been invited to the same event for several years now and one of the things that happened is that the young Spanish lawyer convinced me of the many opportunities that exist for promoting the quantum computing ecosystem in Europe and creating connections with the US. Fast forward to my departure from IonQ, and now that I was without any potential conflicts, the Spanish lawyer and I decided to create an entity that would allow us to bring together entrepreneurs, startups, investors, industrial partners, research entities, and government organizations to promote,

educate, fund, and build the quantum computing ecosystem in Europe and in the US. That entity is Bluzec, an ever-evolving consulting startup in its own right that wants to play a key role in the continued development of quantum computing and quantum technologies in general. And that is where I am now, working to make Bluzec a successful startup.



# Chapter 9

## Seeing What Isn't There

### *An epistemology of quantum possibility*

Genevieve Hayman

#### 9.1 The View from Somewhere

Every inquiry begins from a perspective. This is not a limitation to be overcome but a condition to be understood. As a philosopher trained in the phenomenological tradition, I learned early that perception is never neutral—that what we see is shaped by the conceptual apparatus we bring to the looking. Edmund Husserl called the deliberate bracketing of our taken-for-granted ontological assumptions about the world the *epoché*, and while I don't intend to turn this chapter into a phenomenology seminar, I mention it because it has been the orienting habit of my intellectual life. Before asking “what is the world really like,” I tend to ask “what are the structures of my experience that allow me to see this way, and what might that vantage point make invisible?”

This habit has taken me across a range of fields that might appear, from the outside, to have little in common. I have written on the philosophy of time-consciousness, on mirror neurons and embodied cognition, on the metaphysics of temporal experience, and on the epistemology of science. I have studied dynamical systems and complex adaptive systems. I have worked in retirement security, economic policy, and financial research. And now I find myself writing a chapter for a book on quantum technologies and finance.

The thread connecting all of it is epistemic: I am drawn to moments where a framework shifts and reveals possibilities that were previously invisible. Quantum

---

This chapter is a featured contribution to the monograph *A Portrait of Quantum Technologies in Finance*, edited by Oswaldo Zapata, PhD, and published by The Quantum Finance Boardroom (2026). All rights reserved by the chapter's author.

mechanics, and the technologies it is beginning to make possible, represents exactly such a moment—perhaps the most profound one since relativity taught us that space and time are not what they seem. What draws me to quantum is not simply the technology. It is the invitation to see differently.

## 9.2 The Roots of Computation in Logic

To understand what is genuinely new about quantum computation, it helps to understand what classical computation is—and, more importantly, where it came from. Computation did not emerge from engineering. It emerged from logic.

The textbook that shaped my understanding of this lineage is *Computability and Logic* by George Boolos, John Burgess, and Richard Jeffrey (shout out to James Mattingly for helping us slog through it). It is a work of formal logic, but at its heart it asks a profoundly philosophical question: what can be computed? What problems can (or cannot) be solved by mechanistically following an explicit set of rules? The answers to these questions—formalized through the work of Gödel, Church, Turing, and others in the 1930s—did not merely create computer science. They mapped the boundaries of a certain kind of thinkability.

Gödel's incompleteness theorems showed that within any sufficiently powerful formal system, there are true statements that cannot be proven from within the system itself. Church and Turing, working independently, contributed to formalizing the notion of an “effective procedure” (what is often understood now as an algorithm) and, in the course of this work, identified a class of problems that no such procedure could ever resolve. The Church-Turing thesis, which holds that any effectively computable function can be computed by a Turing machine, became the conceptual foundation of the digital age.

Classical computation emerged from efforts to formalize what it means to carry out a procedure mechanistically and to recognize there are boundaries to what such procedures can do. The Turing machine is extraordinarily powerful, but still processes information sequentially, deterministically, one state at a time. Even today's supercomputers, which perform many computations in parallel, are built within the computational paradigm of Turing machines, operating with determinate state transitions governed by explicit rules.

Quantum computation may ultimately be bounded by the limits of computability described by Gödel, Church, and Turing, but it fundamentally reconceives what it means for an effective procedure to explore a problem's possibility space. Within classical computation, the space of possible configurations is vast, but the machine explores it through explicit, stepwise transitions between definite states. Quantum computation does not move through a possibility space in this way—it spreads across it. Employing quantum mechanics, the hardware can access the landscape of possibility, and the algorithm's task becomes choreographing the interference among those possibilities so that the right answer emerges with high probability when the system is finally measured. While quantum computation may not redefine the theoretical limits of computability, it does reimagine how a formal system could

operate.

### 9.3 The Quantum Rupture

Quantum mechanics is often described as counterintuitive, as though it were a puzzle that, with enough effort, might eventually be reconciled with our everyday experience. I think this framing misses the point. Quantum mechanics is not counterintuitive—it is differently intuitive. It describes a world that operates according to its own internal logic, one that happens to diverge from the logic our evolved perceptual systems were built to navigate.

Consider what quantum mechanics actually tells us. A particle does not occupy a single definite state until it is measured; it exists in a superposition of states, each with an associated probability amplitude. Two particles can become entangled such that the state of one instantaneously constrains the state of the other, regardless of distance. These are not metaphors or approximations. They are descriptions of how the world operates at a fundamental level.

The classical world gives us a picture in which things have definite properties and follow deterministic trajectories. The quantum world gives us a picture in which possibilities coexist, interfere with one another, and resolve into definite outcomes only under specific conditions. The everyday physics we take for granted is, in a deep sense, an artifact of our perspective—a classical approximation of a quantum reality.

This is what I mean by an epistemic rupture. It is not simply that we have learned new facts about subatomic particles. It is that the conceptual framework through which we understand what “facts” are has shifted. And quantum computation is the technological expression of this shift. A quantum computer does not just process information faster. It processes information differently, exploiting superposition to explore the vast possibility space, using entanglement to create correlations that no classical system can efficiently produce, and employing interference to amplify the paths that lead to solutions while canceling those that do not.

The move from classical to quantum computation is not an upgrade within the same paradigm. It is a change of paradigm.

### 9.4 Phase Portraits and the Landscape of Possibility

My training in dynamical systems gave me a particular way of thinking about this shift—one that, while not a strict technical equivalence, captures something important about what quantum opens up.

In dynamical systems theory, one of the most powerful tools is the phase portrait. A phase portrait maps the full landscape of a system's possible behaviors: not just where the system is at a given moment, but every trajectory it could follow from every possible initial condition. Fixed points, limit cycles, strange attractors—these

structures reveal the deep geometry of what a system can do. The power of the phase portrait is that it lifts your attention from the particular to the possible. You stop asking “where is the system?” and start asking “what is the space of behaviors available to it?”

This training prepared me for quantum in a way I didn't anticipate. When I first encountered the formalism of quantum mechanics (the state vectors in Hilbert space), it felt like a familiar gesture. Here, too, was a framework that asked you to think not about a single definite state but about an entire landscape of possibility. The probability-based nature of quantum mechanics, and the idea at the heart of quantum computation that one might compute across possibilities simultaneously, resonated with the dynamical systems habit of mind: always look at the full portrait, not just the current position.

I do not want to overstate the analogy. Hilbert space is not phase space; quantum amplitudes are not trajectories in the classical sense. But the conceptual move—from attending to the discrete and actual to attending to the nondiscrete and possible—is shared. And I believe that anyone who has internalized the phase portrait way of thinking has a head start in grasping why quantum computation is not merely “faster classical computation” but something categorically different. It operates in a richer space of possibility, one that classical systems cannot reasonably access.

## 9.5 Complex Systems and the Comfort of Irreducibility

There is another thread in my background that bears on this: my work in complex systems.

Complex adaptive systems (from ecosystems to economies to neural networks) share a set of features that resist classical reductionist analyses. They are nonlinear: small changes can produce large effects, and large interventions can produce nothing at all.

They exhibit emergence: the behavior of the whole cannot be predicted from the additive behavior of the parts. They are path-dependent: history matters, and the same inputs can lead to different outcomes depending on the route taken. And they are characterized by irreducible uncertainty. Not uncertainty born of ignorance, but uncertainty that is structural, baked into the dynamics themselves.

Working with complex systems teaches you a particular kind of intellectual humility. You learn to be comfortable with the fact that some phenomena cannot be decomposed into simpler components without losing what makes them interesting. You learn to hold multiple possibilities in mind simultaneously, to think in terms of distributions and landscapes rather than point predictions. You learn that the map is never the territory, and that useful models are the ones that preserve the right kinds of structure, not the ones that aspire to total accuracy.

This training, too, turned out to be preparation for quantum. Both complex systems thinking and quantum mechanics demand comfort with irreducible uncertainty. Both

require you to take seriously the idea that wholes can behave in ways not derivable from their parts. Both reward the thinker who can hold a space of possibilities open rather than collapsing prematurely to a single answer.

For those who work in finance, none of this should be entirely unfamiliar. Financial markets are complex adaptive systems (check out this piece<sup>1</sup> I wrote on it). They exhibit nonlinearity, emergence, path-dependence, and deep uncertainty as matters of daily reality. The models that work in finance are models that respect this complexity rather than assuming it away. I would suggest that the intellectual habits cultivated by serious engagement with complex systems are precisely the habits that will prove most valuable as quantum technologies enter the financial landscape.

## 9.6 Alpha as Epistemic Advantage

This brings me to the question that, for many readers of this book, is likely the most pressing: what does any of this have to do with generating returns?

My answer is that it has everything to do with it—but not in the way that most discussions of quantum computing and finance tend to suggest. The typical framing focuses on computational advantage: quantum computers will optimize portfolios faster, price derivatives more efficiently, crack encryption sooner. These are real possibilities and they deserve serious attention. But they are also, in a sense, the least interesting part of the story. They treat quantum computation as a faster engine bolted onto the same chassis. They miss the deeper transformation.

In finance, alpha (i.e., risk-adjusted excess return) is fundamentally an epistemic phenomenon. It arises when someone sees something that others do not. It is the return on a different way of looking. Every genuine source of alpha, if you trace it back far enough, rests on an informational or analytical asymmetry: a framework that captures structure in the world that other frameworks miss.

The history of finance is, in this sense, a history of epistemic shifts. Quantitative finance itself was an epistemic shift—the recognition that mathematical models could reveal structure in market behavior that fundamental analysis alone could not. Index-based and factor investing was an epistemic shift. So was the recognition that markets are not always efficient, that behavioral biases create predictable patterns, and that tail risks are systematically underestimated.

Quantum technologies represent the next such shift, and I believe it will be among the most consequential. Not because quantum computers will be faster (though for specific problems, they may permit speedups by many orders of magnitude) but because they operate according to a fundamentally different approach to logic of classical computers. They can exploit quantum mechanics to extract solutions from large, multidimensional state spaces in ways that have no efficient classical equivalent. And they carry the potential to find structure in data that no known classical algorithm can plausibly detect.

---

<sup>1</sup><https://rpc.cfainstitute.org/research/reports/2025/reframing-financial-markets-as-complex-systems>.

But here is the crucial point: to benefit from this, one must actually undergo the epistemic shift. Treating a quantum computer as a faster classical computer is like treating a telescope as a longer arm. It misunderstands what the tool is. The firms and analysts who will generate alpha from quantum technologies will be those who understand what quantum computation actually does, who grasp that it navigates the possibility space in an entirely different way, and who can translate that understanding into analytical frameworks that reveal structure currently invisible to classical approaches.

This is not a purely technical challenge. It is a conceptual one. It requires the willingness to let go of deeply held assumptions about how computation works, about what constitutes a solution, about the nature of the systems being modeled. It requires, in other words, exactly the kind of framework-level rethinking that has always been the source of genuine alpha.

## 9.7 An Invitation

I did not plan the intellectual journey that brought me here. No one decides at the outset to move from Husserl to Hilbert space by way of dynamical systems and pension policy. But looking back, I can see that each step was, in its own way, preparation for the next. Phenomenology taught me to attend to the perspective that shapes every observation. Logic taught me the power and boundaries of computability. Dynamical systems taught me to think in terms of possibility landscapes rather than individual trajectories. Complex systems taught me to be comfortable with irreducible uncertainty and emergence.

And quantum mechanics brought all of these threads together in a framework that is at once mathematically rigorous and philosophically radical. It tells us that the world is richer than our classical models can reasonably compute and that possibility itself has a structure that we are only beginning to explore.

For those in finance, I would offer this: the quantum era is not just about adopting new tools. It is about adopting new ways of seeing. The epistemic shift that quantum demands is uncomfortable and requires sitting with ideas that feel wrong, frameworks that violate intuition, possibilities that seem to contradict one another. But this discomfort is not a sign that something has gone wrong. It is a sign that something genuinely new is becoming available.

The possibilities are there. The question is whether we will learn to see them.

# Chapter 10

## A Dance of the Blind Puppeteer

*The interplay between a classical optimizer and the Hilbert space*

Jacob L. Cybulski

**Abstract.** This chapter explores the intricate training dynamics of quantum machine learning models, framed as an interplay between classical optimization and the vast geometry of the Hilbert space. We move beyond the common trope of Hilbert space “vastness,” focusing instead on a geometry sculpted by quantum entanglement and operational constraints. We elucidate the “Blind Puppeteer” problem: the classical optimizer’s struggle to refine classical parameters without insight into their quantum representation or their role in evolving qubit states. From a purely geometric perspective, we analyze how the high dimensionality of both the parameter and Hilbert spaces culminates in training pathologies like Barren Plateaus, Orthonormal Desert, and the Measurement Sparsity. Ultimately, we argue for the necessity of “quantum-aware” optimizers—tools capable of navigating the subtle curvature of the quantum manifold that standard gradient-based methods ignore.

### 10.1 Introduction

This chapter presents a concise synthesis of selected concepts, methods, and procedural frameworks in Quantum Machine Learning (QML), which often represent substantial entry barriers for researchers and practitioners in applied Quantum Computing (QC).

The motivation for using QML as a problem-solving toolkit is typically two-fold. First, it can directly model quantum phenomena with inherent nondeterminism, high-dimensional feature spaces, and complexity (e.g., quantum chemistry or quantum

---

This chapter is a featured contribution to the monograph *A Portrait of Quantum Technologies in Finance*, edited by Oswaldo Zapata, PhD, and published by The Quantum Finance Boardroom (2026). All rights reserved by the chapter’s author.

physics). Second, it applies Quantum Mechanics as a lens for complex classical problems characterized by imprecision, large volumes of poor quality or aging data, and dissatisfaction with deterministic or stochastic methods—particularly in quantum finance or optimization, where one already contends with complex combinatorial landscapes. In this chapter, we focus on the latter category of problems and potential QML solutions, which commonly arise in business computing, and in particular logistics, supply chain management, and finance.

To support the adoption of QML for business problem solving, foundational quantum computing concepts must be complemented with calculus, data science, information geometry, and an understanding of how classical computation interacts with quantum phenomena. This chapter provides intuition and a conceptual framework for core QML ideas at the boundary of the classical and quantum realms. We focus on developing predictive and decision-support quantum models trained with classical optimization algorithms that, while robust in classical domains, largely ignore the unique geometry of quantum effects. In this context, we portray such an optimizer as a *blind puppeteer*, pulling the marionette’s strings to make it dance, guided only by the intermittent cheers of the audience as the sole, indirect feedback on his artistry.

## 10.2 Strings: PQC Parameters

Quantum circuits consist of qubits, operations, and measurements. Executing a circuit produces a quantum state that evolves under these operations. A standard circuit is static: it hardcodes classical data and parameters, so any change requires building a new circuit. Such circuits cannot be easily adapted or optimized for different analytic goals or data. However, a Parameterized Quantum Circuit (PQC) provides a template that separates the circuit structure and function from its parameters (Figure 1).




Figure 1: Parameterized quantum circuit trained with variational quantum algorithm.

PQCs can be designed to fulfill specific analytic objectives, e.g. to generate predictions or decisions based on data representing past events. Such PQCs typically consist of the following functional sections:

- *Feature map* embedding classical data in the PQC’s input parameters associated with operations responsible for setting the model’s initial quantum state;

- *Ansatz* consisting of qubit entangling blocks and blocks of rotational operations parameterized with circuit weights, which could be trained to evolve the initial circuit state into a target state;
- *Measurements* responsible for decoding the target state of the executing circuit into a classical data form, which is interpreted and returned as output.

The components of these sections may be intermixed, e.g. data embedding blocks can be interwoven with encoding and trainable ansatz blocks (known as data reuploading), or measurements can appear in the midst of the ansatz (known as mid-circuit measurement).

PQC training typically uses a Variational Quantum Algorithm (VQA) (Figure 1) that relies on a classical optimizer, which is equipped with a classical training dataset and a cost function, to iteratively arrive at the optimum values of the circuit weight parameters. The role of a cost function is to aggregate the results of the associated loss function, which calculates the circuit execution error. In this process, the optimizer acts completely unaware of any quantum processes prescribed by the quantum circuit. However, it is aware and in full control of the classical elements of the circuit, i.e. its input and output values, the ansatz weights, and the result of circuit execution.

The classical circuit elements controlled by the “quantum-blind” optimizer are considered the puppet strings.

### 10.3 Stage and Puppet: Hilbert Space and Manifold

To understand quantum machine learning, one must first visualize the “stage” upon which the quantum model performs—the Hilbert space [4]. While we often speak of the Hilbert space as the arena of quantum states, it must be formally treated as a *Projective Hilbert Space* to be physically meaningful. This is achieved by adding an equivalence relation that treats all vectors differing only by a complex scalar as the same state, effectively collapsing the linear space into a multi-dimensional complex projective surface, referred to as the *manifold*. Unlike the familiar three-dimensional Euclidean space of classical physics, the state space of a quantum system is curved, and hence not measured in straight lines, but by the “overlap” or *fidelity* between quantum states.

A PQC acts as a map that translates classical numbers—the parameters—into specific quantum state configurations in the Hilbert space. However, this mapping is rarely one-to-one or uniform. Because of quantum entanglement, some PQC states become inseparable, which creates a *manifold*—a lower-dimensional surface within the massive Hilbert space—upon which the model is forced to reside. The shape of this manifold dictates whether the model can reach the “ideal” state required to solve a problem or whether it is trapped in a geometric dead-end. The geometry of the manifold is constrained by the mathematical laws of superposition, entanglement, and the *Fubini-Study metric* (FS) [5].

Table 1: Layer-by-layer state evolution and observable projections (cf. Figure 2). Notation:  $p_0 \equiv \text{params}[0]$ ,  $p_1 \equiv \text{params}[1]$ ,  $c_i \equiv \cos(p_i/2)$ ,  $s_i \equiv \sin(p_i/2)$ ,  $\phi_0 \equiv p_0 p_1/4$ ,  $\phi_1 \equiv 0.15 p_0$ .

Step	State / Observables	Comment
1	$ \psi_1\rangle = c_0 00\rangle + s_0 10\rangle$	$R_Y(p_0)$ rotates qubit 0 into superposition.
2	$ \psi_2\rangle = c_0 c_1 00\rangle + c_0 s_1 01\rangle$ $+ s_0 c_1 10\rangle + s_0 s_1 11\rangle$	$R_Y(p_1)$ rotates qubit 1; state remains separable.
3	$ \psi_3\rangle = c_0 c_1 00\rangle + c_0 s_1 01\rangle$ $+ s_0 s_1 10\rangle + s_0 c_1 11\rangle$	CNOT swaps $ 10\rangle \leftrightarrow  11\rangle$ ; generates entanglement.
4	$ \psi_4\rangle = e^{-i\phi_0}(c_0 c_1 00\rangle + c_0 s_1 01\rangle)$ $+ e^{+i\phi_0}(s_0 s_1 10\rangle + s_0 c_1 11\rangle)$	$R_Z(\phi_0/2)$ on qubit 0 applies control-dependent phase, adds nonlinear/interactive warping.
5	$ \psi_{\text{fin}}\rangle = c_0 c_1 e^{-i(\phi_0+\phi_1)} 00\rangle + c_0 s_1 e^{-i(\phi_0-\phi_1)} 01\rangle$ $+ s_0 s_1 e^{+i(\phi_0-\phi_1)} 10\rangle + s_0 c_1 e^{+i(\phi_0+\phi_1)} 11\rangle$	$R_Z(\phi_1)$ on qubit 1 adds independent phase, linear/gauge-like warping; final state.
6	$\langle Z_1 \rangle = \cos p_0 \cos p_1$ , $\langle Z_0 Z_1 \rangle = \cos p_1$ $\langle X_0 \rangle = \sin p_0 \sin p_1 \cos(p_0 p_1/2)$	Expectation values: $Z$ -terms phase-invariant; $\langle X_0 \rangle$ encodes nonlinear interference, it contains $\cos(p_0 p_1/2)$ because the $R_Z$ rotation “drags” the state around the $Z$ -axis, modulating the projection onto the $X$ -axis.

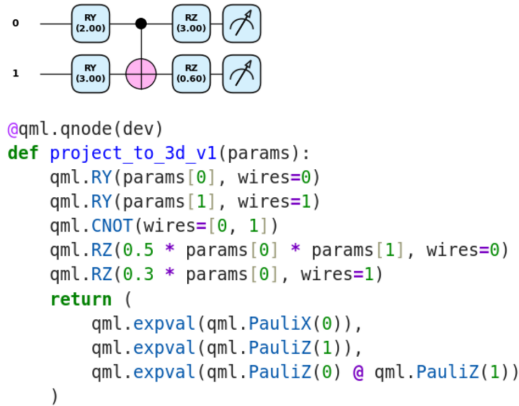



Figure 2: Sample quantum circuit and its implementation in PennyLane.

While *fidelity* measures state similarity, the *Fubini-Study metric* defines the intrinsic geometry of the Hilbert space. By characterizing how a quantum output state changes under infinitesimal parameter shifts, it sets the local curvature of the training surface, “warping” Euclidean parameter space into regions of high sensitivity or stagnation. FS provides the mathematical basis of the *Quantum Fisher Information Matrix* responsible for mapping the curvature of the optimization landscape, hidden from the classical optimizer, but which it nevertheless unknowingly traverses.

Table 2: What aspects contribute to manifold creation

Concept	What it is	Does it define the manifold?
<b>Circuit Ansatz Parameterization</b>	The mapping $(\theta, x) \mapsto U(\theta, x)  0\rangle$	<b>Yes.</b> Mathematically defines the intrinsic manifold $\mathcal{M} = \{ \psi(\theta, x)\rangle\}$ .
<b>Input Encoding</b>	Data $x$ enters as gate parameters, e.g., $R_Y(x_i)$	<b>Yes.</b> Data often defines the input state location on the manifold $\mathcal{M}$ , selecting specific trajectories or submanifolds.
<b>Ansatz Weights</b>	Weights $\theta$ also enter as gate parameters, e.g., $R_Z(\theta_{i1}), R_Y(\theta_{i2}), R_Z(\theta_{i3})$	<b>Yes.</b> Weights define the transformation/rotation of the manifold $\mathcal{M}$ .
<b>Execution and Multiple Shots</b>	Repeated measurements to estimate $\langle O \rangle$	<b>No.</b> Shots <i>sample</i> the manifold; they do not create it. Finite shots introduce statistical noise.
<b>Classical Optimizer</b>	Updates $\theta$ to minimize a scalar cost function	<b>No.</b> Navigates a cost landscape over $\mathcal{M}$ , without knowledge of the underlying quantum geometry.

A quantum model manifold can be derived mathematically from its circuit structure, where the unitary parameterization defines the state evolution, and consequently the model geometry, which exists independently of the circuit execution (see Table 2). As illustrated in Table 1, a quantum circuit (Figure 2) provides sufficient details to predict its step-by-step state evolution. Furthermore, we can predict the warping of the circuit manifold in response to entanglement operations and nonlinearities resulting from  $R_Z$  rotations, which do not alter quantum circuit measurements (in  $Z$  basis), but which accumulate bending the model manifold (in  $X$  axis that is non-commutative with  $Z$  rotations) as parameters change.



(a) Warped circuit manifold with four selected state evolution paths

(b) Optimization cost landscape, which is lacking any details of the model's quantum geometry

Figure 3: The circuit manifold and its cost landscape

Alternatively, we can rely on a quantum processor or simulator to empirically approximate the manifold by sampling the circuit expectation values from finite-shot measurements (Figure 3a). In such an approach, training data are necessary to drive

the geometric reconstruction process. However, data themselves do not “create” the manifold. Instead, data values act as coordinates that select specific trajectories through the manifold (e.g. see state evolution paths in Figure 3a). Apart from entanglements, which are the main agents of shaping the manifold geometry, the ansatz weight parameters are also capable of adjusting the manifold’s curvature.

In practical quantum computing, as opposed to the purely mathematical conceptualization of the Hilbert space, due to decoherence (perturbations to the phase) and noise (perturbations to the magnitude of density matrix elements), data points sampled with NISQ quantum machines do not fall precisely onto the mathematically idealized manifold surface (of pure states). Instead, the empirical manifold forms a cloud of quantum state locations (of mixed states).

In the world of quantum puppeteers, the Hilbert space is the stage for the performance of the manifold—the puppet.

## 10.4 Blind Puppeteer: Classical Optimizer

The actual training of a quantum model is conducted by a classical optimizer – a mathematical algorithm, which adopts the VQA approach but which remains entirely oblivious to the quantum nature of the system it is tuning (Figure 3a).

The optimizer iteratively creates static PQC instances, each with different input samples from the training dataset and different weight parameters. It treats each PQC instance as a black box that, when executed multiple times, returns an approximate output for the chosen inputs. This approximation is compared with the target outputs using a loss function (as implied by the cost function), yielding a single error value, or *loss*. Aggregating these losses over the entire training dataset gives the overall *cost* of a given model parameterization. All possible instantiations of the PQC parameters can be associated with their respective costs, thus forming the *cost landscape* (Figure 3b). Only fragments of the cost landscape are computed to reduce computational costs.

The classical optimizer cannot see qubits, interference, or entanglement, it can only see the cost landscape. Its primary task is to efficiently navigate this landscape to find the geometrically lowest point in the landscape—and consequently the lowest cost and indirectly the optimum circuit parameterization. During its search for the optimum, at each iteration the optimizer is moving blindly through the parameter space by relying entirely on the local surface shape. Optimizers such as SGD and Adam rely on the landscape slope, or *gradient*, to decide the next improvement to the parameter set. However, gradient-free optimizers, such as Nelder-Mead or COBYLA, also suffer from the same blindness; they simply touch the surface for a lower point rather than measuring the pitch of the slope.

The optimizer is the blind puppeteer, sightlessly animating its high-dimensional parameter-strings in search of the optimum dance steps, listening to the erratic clapping of the audience as the noisy feedback for his artistry.

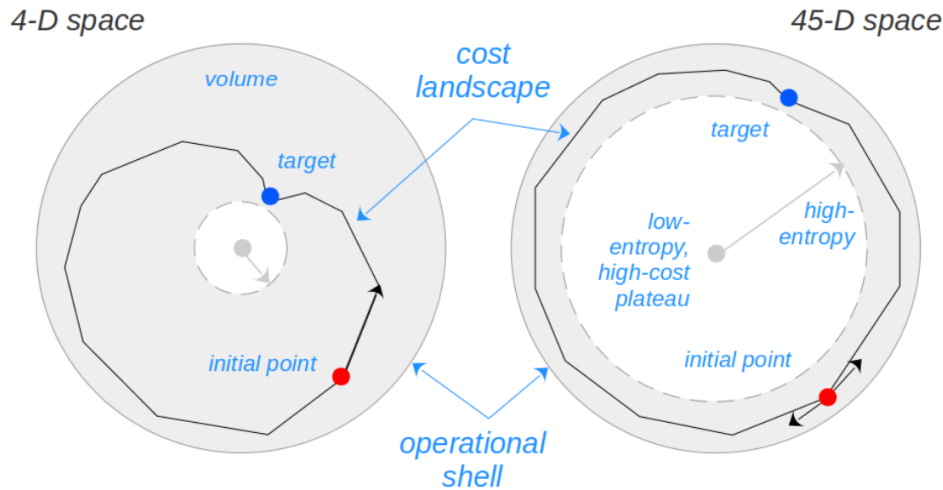


Figure 4: 2D metaphor of a 4-dimensional  $n$ -ball (left), and 45-D  $n$ -ball with a *Barren Plateau* (BP) (right). Note how volume (grey) in  $n$ -ball shrinks with the number of dimensions.

## 10.5 Too Many Strings: The Curse of Dimensionality

The “Curse of Dimensionality” often turns a manageable optimization problem into an intractable one as the number of its parameters or qubits increases, raising the tension between model expressivity (ability to represent data in quantum space) and trainability (capacity to learn and generalize with optimization efficiency) [1]. Higher dimensionality may be required to give the circuit enough expressivity to learn the patterns of a complex dataset, but this typically harms the optimizer effectiveness, and thus the model trainability. To understand why a classical optimizer faces difficulties when training large quantum models, we must look at the geometry of the spaces it is trying to navigate.

Consider the previously discussed quantum model (Figure 2). The representations of its state space (Figure 3a) and its cost landscape spanning the parameter space (Figure 3b) are very different. In this case, the shape of the state space is folded and complex, and the shape of the cost landscape is simple and smooth. However, when their respective dimensions become very high, both shapes will display the same unusual characteristics and problems.

In the explanation of these problems, we will use a well-known concept of a multidimensional unit-size  $n$ -ball as a proxy for the parameter space as well as the Hilbert state space—the  $n$ -ball principles apply to any high-dimensional space. By analyzing the high-dimensional  $n$ -ball properties we encounter a surprising series of paradoxes (cf. Figures 4, 5 and 6).

**The Shrinking  $n$ -Ball Volume Paradox.** In our three-dimensional world, we expect that adding dimensions increases the space available for a solution. However, the volume of  $n$ -dimensional  $n$ -ball with radius  $R=1$  follows a counter-intuitive

trajectory defined by the equation:

$$V_n(R) = \frac{\pi^{n/2}}{\Gamma(\frac{n}{2} + 1)} R^n, \quad \text{and} \quad V_n(1) = \frac{\pi^{n/2}}{\Gamma(\frac{n}{2} + 1)} \quad \text{for } R = 1 \quad (10.1)$$

where  $\Gamma(z)$  is a *Gamma function*, an extension of the factorial function to complex numbers, which is defined as follows:

$$\Gamma(z) = \int_0^\infty t^{z-1} e^{-t} dt, \quad \text{for } \text{Re}(z) > 0 \quad (10.2)$$




Figure 5: Manifestation of an *Orthonormal Desert*. Note how distribution of pairwise distances between states within an  $n$ -ball concentrates around the mean as the dimension increases.

As the dimensionality  $n$  increases, the volume  $V_n(1)$  does not grow indefinitely. It reaches a maximum at  $n \approx 5.25$  and then begins an exponential collapse toward zero. By the time we reach the 45-dimensional space (figuratively represented in Figure 4), the total volume of the parameter space has effectively vanished. The space itself does not become empty, instead the mathematical definition of its volume in high dimensions diminishes the space “interior.”

**The Thinning-and-Expanding  $n$ -Ball Shell Paradox.** Perhaps the most damaging property for an optimizer is the *concentration of volume*. In high dimensions, the volume of the  $n$ -ball is not evenly distributed; instead, it migrates almost entirely to a thin layer near the surface (Figure 4). The volume contained within a surface shell of thickness  $\epsilon \in [0, 1]$  is defined as the difference in volume of the  $n$ -balls of radius 1 and  $(1 - \epsilon)$ , which can be shown to be:

$$V_{shell} = V_{total} [1 - (1 - \epsilon)^n] \quad (10.3)$$

As  $n \rightarrow \infty$ , the term  $(1 - \epsilon)^n$  vanishes. For example:

- In 3D ( $n = 3$ ): The outer 1% of the radius contains only  $\sim 2.9\%$  of the volume.
- In 45-D ( $n = 45$ ): The outer 1% of the radius contains  $\sim 36.4\%$  of the volume.
- In 500-D ( $n = 500$ ): The outer 1% of the radius contains  $\sim 99.3\%$  of the volume.

In our 3D intuition, the outermost 1% of a sphere’s radius is an insignificant sliver, containing less than 3% of its volume. However, as we move to a 45-dimensional model, this same 1% shell expands to claim more than 36% of the total searchable space.

In the  $n$ -ball proxy of the Hilbert space,  $n$  represents the dimension of the state space ( $2^{\#\text{qubits}}$ ), illustrating how a linear increase in physical qubits leads to an exponential ‘thinning’ of the interior; e.g. at 9 qubits, the model state space already exceeds 500 dimensions ( $2^9 = 512$ ). In the parameter space, a circuit of 19 qubits with 9 rotational blocks of 3 rotational operations each ( $R_z, R_y, R_z$ ) exceeds 500 dimensions ( $19 \times 9 \times 3 = 513$ ). After that point, the volume rapidly migrates toward the surface, leaving the interior hollow.

For the optimizer, this creates a serious difficulty. Any movement the optimizer makes is statistically forced to be *tangential* to the surface. It is no longer moving “toward” or “away” from anything—it is merely drifting around a very thin high-dimensional shell.

**Gradient Extinction Paradox.** Because the “interior” has vanished and all content is held in a thin surface shell, the “slope” between an arbitrary pair of points becomes too shallow to detect. This *concentration of gradient* around zero makes any surface within the shell completely flat. In the parameter space, the cost landscape becomes a *Barren Plateau*, i.e. also flat [3] (Figure 4), depriving the classic optimizer of any meaningful gradient to follow.

**Orthonormal Desert and Equidistance Paradox.** In a curious twist of multi-dimensional mathematics, any two random vectors pointing to the surface of the  $n$ -ball are approximately orthogonal to each other (Figure 5, right). This *concentration of angular separation* around orthogonality can be formally justified by examining the dot product of these vectors. Since in each coordinate dimension, the product of the corresponding components is equally likely to be positive or negative, the contributions from individual dimensions tend to cancel out. Consequently, in a high-dimensional space, the sum of these products converges statistically toward zero, leading to expected near-orthogonality between randomly selected vectors.

This makes the majority of the  $n$ -ball volume the *Orthonormal Desert*.

At the same time, it can be shown through computational experiments that a distribution of distances between randomly selected points in high-dimensional  $n$ -ball concentrates close to their mean (Figure 5, left). In fact, since all points in the thin  $n$ -ball crust shell are equidistant from its center, the radius of the ball is 1 and the

points are orthogonal, their pairwise distance must be approximately  $\sqrt{2}$ , creating an incredible *concentration of distance* in high-dimensional space.

From the classical optimizer’s perspective, the Hilbert space is not only very large but also incredibly homogeneous in all directions.

The resulting state space is a highly fractured and directionless void. The classical optimizer is not just stuck in the local minimum, it is trapped on a high-dimensional manifold where the very concept of a path toward the target has been erased by the geometry of the space itself. While the optimization solution exists, the geometry of the  $n$ -ball ensures that the optimizer lacks the “volume”, “direction” and “smoothness” to navigate the  $n$ -ball space leading to it.

In such situations, one of the most deceptive aspects of quantum model training is the appearance of convergence. Because the cost landscape is subject to extreme flattening due to *Barren Plateaus*, the gradient signal often drops below the numerical or statistical noise floor. To a classical optimizer, this lack of signal is indistinguishable from reaching a true minimum. Practitioners may find their models converging quickly to a stable cost value, only to discover that the resulting “optimal” quantum state remains deep within the *Orthonormal Desert*—nearly orthogonal to the target and possessing no actual predictive power.

**Entropy Explosion Paradox.** Because the radial probability of finding a state at a given distance from the center  $P(r) = nr^{n-1}$  vanishes at the origin, the  $n$ -ball center is a region of minimal entropy—a unique, highly ordered state that occupies no volume in the high-dimensional landscape. In contrast, the  $n$ -ball outer shell is a high-entropy mix of states, containing the vast majority of state configurations. Paradoxically, this leads to *concentration of entropy*, where uncertainty becomes confined to an extremely narrow range of space.

In terms of circuit states, the states that are most likely to occupy the vicinity of the  $n$ -ball center are unique, specific, and unlikely to be encountered by the optimizer.

Also, because the volume of the  $n$ -ball is concentrated near the surface, a random initialization of the circuit state (the red dot in Figure 4) is statistically likely to land in a narrow but high-entropy (s)hell. The target state (the blue dot in Figure 4) of the lowest cost is also likely to be located in the same high-entropy region, where all states look equally probable and orthogonal, thus making the target a “needle in a high-entropy haystack.”

The optimizer begins its journey not only far from the target, but its search is also hampered by the huge volume of competing states packed into a narrow sliver of space—a high-entropy labyrinth where every direction looks the same to a blind puppeteer.

**The Statistical Conundrum.** While the geometry of the  $n$ -ball ensures that the directional signal is mathematically extinguished, measurement difficulties create a separate practical bottleneck. This is not merely a matter of diminishing gradients, but a fundamental issue of statistical resolution [2].




Figure 6: Manifestation of *Measurement Sparsity*. Note the sparse distribution of measurement outcomes. As counts (shots) are low, the calculated probabilities become imprecise.

Figure 6 illustrates what happens when the number of qubits ( $n$ ) increases and the dimension of the outcome space consequently grows exponentially ( $2^n$ ). In such situations, a fixed number of measurement “shots” must be distributed across a vast array of potential results. In low dimensions, these shots aggregate into a clear probability distribution. In high dimensions, the same number of shots becomes spread so thinly that the results appear as a sparse, random scatter. This *Measurement Sparsity* means that the “audience’s cheers” are no longer a coherent signal, but a chaotic sequence of individual claps. For the blind puppeteer, the cost landscape becomes “jittery” and unreliable, as the statistical noise of the sampling process effectively drowns out any surviving geometric signal.

In the interplay between the optimizer and Hilbert space, we are confronted not only with the noise of quantum hardware but also with the nature of the circuit’s high-dimensional geometry.

## 10.6 Resolution: Navigating the Metric Mismatch

The fundamental friction in *Variational Quantum Algorithms* arises from a profound metric mismatch. Classical optimizers, such as Adam or BFGS, operate under the assumption of a flat, Euclidean landscape where the distance between two points is measured along a straight line. However, the *Parametrized Quantum Circuit* does not inhabit such a world. Instead, it moves within the *Fubini-Study metric* of the Hilbert space—a highly curved manifold where a small step in parameter space ( $\theta$ ) may result in a massive jump in the state space, or conversely, no meaningful change at all. To resolve this, we must provide the blind “classical-puppeteer” with a geometric compass, translating the optimizer’s Euclidean intuition into the language of quantum geometry (Figure 7).

The resolution lies in the adoption of the “quantum-aware” optimization, specifically the *Quantum Natural Gradient* (QNG). By incorporating the *Fubini-Study metric* tensor directly into the update rule, we effectively even out the curved manifold from the optimizer’s perspective. Crucially, this requires the classical cost function to account for the interplay between the manifold’s curvature and the classical parameter

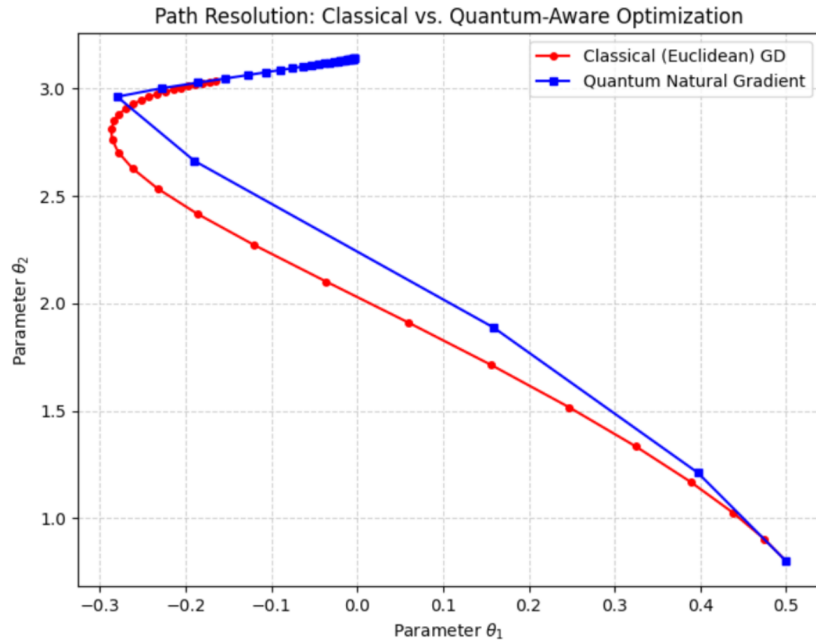


Figure 7: Euclidean vs. QNG optimization.

space. We are not just looking at the shape of the Hilbert space in isolation but how our specific “classical-strings”—the PQC parameters—pull against that geometry. By scaling the gradient by the inverse of the metric, the puppeteer no longer pulls in the dark—every update is weighted by the local sensitivity of the state space relative to its parameters.

As an example, Figure 7 compares the efficiency of the classical and quantum-aware optimization of the circuit presented in Figure 2. Although the cost landscape in Figure 3b suggests a perilous path for any optimizer, Figure 7 (blue) demonstrates how geometric awareness fundamentally alters the optimization journey. By operating within the natural metric of the manifold shown in Figure 3a, the Quantum Natural Gradient effectively collapses the “warped” distances of the parameter space. It identifies the most direct route to the manifold’s high-volume regions, bypassing the inefficient, curved trajectory that the classical gradient (red) is forced to follow due to its lack of geometric context.

The difference in final convergence (Figure 7, top-left), where the Quantum Natural Gradient (QNG) reaches a cost of  $\approx -2.00$  while the classical optimizer stalls at  $\approx -1.992$ , illustrates the *Metric Mismatch*. Although this gap seems small in parameter space, it reflects a fundamental gap in the Hilbert space. In the vastness of the Hilbert space, the 0.008 difference in cost can represent a significant distance in state fidelity, marking the difference between a model that has truly captured the target phenomenon and one that is merely hovering near it. The classical “blind puppeteer” is clearly misled by the flattened Euclidean landscape, mistaking a shallow ridge for the minimum. In contrast, the QNG’s geometric awareness reveals remaining room for improvement in state space, guiding the model to the true global minimum, before which the classical approach stalled.

However, geometry alone sometimes cannot escape the most resilient plateaus—regions

so featureless that even a compass offers no heading. In these cases, we introduce stochastic perturbations, a deliberate “shaking” of the “puppet strings”. By injecting controlled noise or using methods like SPSA, we break the symmetries that cause stagnation. These perturbations act as a kinetic energy boost, knocking the model off barren manifolds and pushing it toward high-volume regions of Hilbert space where global optima are easier to find.

Ultimately, navigating the metric mismatch requires a synthesis of both sight and movement. Whether through the precise lens of the *Quantum Natural Gradient* or the strategic chaos of stochastic shaking, the goal remains the same: to harmonize the classical update with the quantum state. By reconciling the “classical-strings” with the Fubini-Study “quantum-stage,” we transform a blind search into a directed exploration, ensuring that the optimizer’s path respects the complex topography of the quantum landscape.

Much of the insights presented in this chapter were related to the paradoxes of the  $n$ -ball, where the volume is concentrated near the surface and points become orthonormal and equidistant, which are often dismissed as artifacts of purely random systems and not applicable to real-world QML modeling. In QML, however, apparent randomness arises from the *Metric Mismatch*. When a classical optimizer navigates a Hilbert-space manifold using a Euclidean cost landscape, the signal eventually vanishes, pushing the algorithm into stochastic drift. Then, its steps become statistically indistinguishable from random sampling, and the  $n$ -ball randomness requirements are met, trapping the model in the *Orthonormal Desert*. Only “quantum-aware” methods like QNG, which follows the manifold’s geodesic curvature, can preserve the specificity needed to avoid this high-entropy extinction.

## 10.7 Conclusions: The Geometric Imperative

The journey through QML model development is ultimately a story of geometric translation. Success in training does not depend on the power of qubits alone, but on the alignment between the classical parameters and the quantum state space. By understanding the constraints of high-dimensional geometry—from the concentration of distance and measure, to the metric mismatch—quantum researchers and professionals can move beyond trial-and-error model development and begin to design architectures that are geometrically optimized for the quantum world.

This chapter explored the notion of the interplay between a classical optimizer and the Hilbert space, which identified the following key takeaways.

**The Geometry of Measure as a Barrier:** The *Curse of Dimensionality* is a physical transformation of the training landscape.

In business applications, such as those in logistics or finance, where we already face combinatorial “explosions,” the  $n$ -ball concentration of volume ensures that the cost landscape is a featureless *Barren Plateau*. For these fields, training is not just a search—it is an effort to prevent the model’s specialized business logic from dissolving into statistical noise.

**The Metric Mismatch:** Optimization fails when a classical optimizer’s *Euclidean assumptions* clash with the *Fubini-Study metric* of the Hilbert space. Without accounting for the “warped” manifold in quantum space, classical optimization algorithms are essentially “flying blind” through non-linear quantum curvature.

**The Transparency Gap and Sampling Barrier:** The classical optimizer only perceives the quantum model through the lossy statistical filter of measurement. *Measurement sparsity* ensures that as the model scales, the resolution of our “instrument” drops, hiding the vanishing geometric signal beneath the noise floor.

**False Convergence:** Due to the possible presence of *Orthonormal Desert*, cost convergence in parameter space does not always ensure fidelity in Hilbert space. In a financial risk model, we must beware “stagnation disguised as success”, where a model appears stable during training but lacks the state-space precision to capture rare, high-impact tail risks.

**The Necessity of the “Shake”:** Strategic, injected stochastic perturbations are a requirement, not an admission of hardware failure. Unlike passive hardware noise, which erodes the landscape, structured perturbations act as a geometric probe to break symmetries and force the optimizer out of stagnant, low-dimensional “potholes.”

**Quantum Natural Gradient:** Classical optimizers navigate the cost landscape under the assumption of a flat parameter space, often resulting in a metric mismatch with the underlying quantum hardware. Conversely, the Quantum Natural Gradient (QNG) utilizes the Fubini-Study metric and the QFIM to perform optimization directly on the Hilbert space manifold. This ensures that the scale of parameter updates remains consistent with the actual change in the quantum state. Crucially, while QNG optimizes the path through state space, the optimality of the resulting candidates is still defined by a classical cost function, leaving the process tethered to the constraints of the classical task.

The “interplay” between a classical optimizer and the Hilbert space is a fundamental conflict of scale and perspective. We are attempting to use a linear Euclidean tool to navigate an exponential projective manifold. The difficulties discussed—*Barren Plateaus*, the *Orthonormal Desert*, and *Measurement Sparsity*—are not merely engineering hurdles to be overcome with more qubits or faster classical hardware. They are the natural consequences of high-dimensional geometry. To bridge this gap, we must move beyond the “black-box” optimization. By adopting a purely geometrical viewpoint—utilizing the *Fubini-Study metric* to correct our steps and injecting *strategic perturbations* to maintain our signal—we can begin to navigate the Hilbert space with “geometric honesty.” The future of Quantum Machine Learning training lies not in fighting these geometric realities, but in building optimizers that actually speak the language of the manifold they are trying to master.

With these goals conquered, we will finally revel in the spectacle of the razor-eyed puppeteer, no longer groping in the dark, but commanding the curvature of the stage to whip his quantum marionettes into a passionate, purposeful dance.

## References

- [1] Cybulski, J. L. and Zajac, S. (2025). Quantum Modelling of Time Series: Expressivity vs. Trainability.
- [2] Lotan, S., Defienne, H., Talmon, R., and Bartal, G. (2025). Sparsity-Driven Entanglement Detection in High-Dimensional Quantum States. *arXiv preprint arXiv:2511.12546*.
- [3] McClean, J. R., Boixo, S., Smelyanskiy, V. N., Babbush, R., and Neven, H. (2018). Barren plateaus in quantum neural network training landscapes. *Nature Communications*, 9(1):4812.
- [4] Schuld, M. and Petruccione, F. (2021). *Machine Learning with Quantum Computers*. Springer, 2nd ed. 2021 edition.
- [5] Stokes, J., Izaac, J., Killoran, N., and Carleo, G. (2020). Quantum Natural Gradient. *Quantum*, 4:269.

## From the Author

As a child, I grew up among artists in a small two bedroom flat in Warsaw, Poland, where every Thursday the tiny place transformed into a studio hosting a dozen of painters, sketch artists and print makers, occasional sculptors and poets. I was among them painting and drawing and dreaming of doing this for the rest of my life. However, at the same time I had passion for mathematics, so eventually I entered a specialist Mathematical high school, where we were taught to discover mathematical miracles through formal proofs, the process that was creative, requiring a lot of imagination and intuitive navigation between own inner space of abstractions—exactly what you do when you create art. I had a really difficult time choosing between different University faculties, which included fine arts, mathematics, theoretical physics, and informatics. I thought a computer was a perfect medium for expressing abstract ideas and eventually finished computer science in Australia, followed by Masters in AI and PhD at the intersection of natural language processing and software engineering. Then I worked at several Australian universities in computer science, information systems, and business analytics; at the same time undertaking self-study of cognitive science and philosophy. I found immersive 3D data visualization to be a perfect area for me to blend arts, mathematics, and computation into a single expression of truth in data and intuition in the eye of the beholder. At some point, I asked myself what area of research is an extension of my interests but would mature not in a week or five years but in many years, perhaps 50 years into the future. This was how I moved to quantum computing and quantum machine learning, which required a lot of learning, a lot of discipline, and a lot of experimentation. So now I paint my quantum models with the colors of state evolution on the manifold canvas of the multi-dimensional Hilbert space.



# Chapter 11

## A Quantum Journey

### *Exploring the theory, building the future*

Jeremy Green

Quantum computing was not about a fascination with physics or an abstract curiosity about the behaviour of particles. As a self-confessed ‘geek’ with an interest in computing and later cybersecurity. My interest in quantum derived as seeing it as a security problem. Reading articles on Post Quantum Cryptography (PQC) I realised we had a technology that could break some existing cryptography and with that technology. As I dived deeper in my PhD in Computer Science I realised that quantum computers could also be attacked with an assortment of Denial of Service (DoS) attacks this then started me on my journey with Q-SLICE and QUANTA.

My current role as a security architect has seen me working on a variety of systems assessing and adding security which included threat modelling, risk governance and system benchmarking. Adding compensating controls to end of life or OT systems. This has helped me have a wider and broader view of systems and balancing useability with security and business requirements as they don’t always meet in the middle.

Quantum computing presented itself precisely in that way. It wasn’t a technology I approached with admiration or excitement. It was new technology which I approached with a security lens, forming slowly under the cryptographic foundations I have always had an interest in. The more I researched quantum the clearer it became that quantum wasn’t simply another emerging technology to monitor it was a fundamental shift in computational capability with direct implications for confidentiality, integrity and long term trust.

I have always wanted to understand the ‘how’ something works and ‘why’ it works that way. This is the same approach I took with quantum computing. Once I understood how it worked, I needed to understand how it changed adversarial capability. How it

---

This chapter is a featured contribution to the monograph *A Portrait of Quantum Technologies in Finance*, edited by Oswaldo Zapata, PhD, and published by The Quantum Finance Boardroom (2026). All rights reserved by the chapter’s author.

reshaped risk horizons and how it undermined current security. By understanding the adversarial element, I could then start to think about the defensive side and the protection that would be required. After all security is a protective and defensive discipline.

My early career didn't begin in industry; it began in the classroom. I started my career as a computer science teacher in both school and college environments, helping students understand the fundamentals of computing, logic and system design. Teaching programming, web development, networking and system support at various levels and age groups. Those years shaped the way I approach technology. Not as a collection of tools, but as a set of concepts that must be understood, explained. Understanding the foundational concepts to build on the more complex elements. These same foundational concepts are what underpin security especially in the more technical world of cybersecurity. Teaching helped me to break complexity into clarity and that skill would become essential later in my journey. By teaching something you also learn it in detail and deepen your own understanding of the topic or subject. You also get interesting questions, different points of view and understanding. I have always found I learn just as much teaching as my learners do!

From education, I moved into policing as a police IT instructor where I had been volunteering for many years as a police officer. That shift brought me into not only operational IT and systems training but into the world of Open Source Investigation (OSI) and the development of a cybercrime course for detectives. Working with SME's in digital forensics and open source I learnt an awful lot about cybercrime, which led me to the world of cybersecurity. However, it was cybercrime, not cybersecurity, that first taught me the mindset of the adversary. It was the opposite side of the coin, but it completed the picture when I moved into teaching cybersecurity.

This experience led me deeper into the security profession. I transitioned into instructing for major security certification bodies such as CompTIA, EC-Council, ISACA and ISC2. Where I taught practitioners how to defend systems, govern risk and build resilient architectures. Each certification I delivered I had to pass the certification before being able to deliver it. This enables me to soak about a wealth of information across twenty certifications. My focus during this period was on classical security such as cryptographic hygiene, governance frameworks, adversarial modelling, secure networks, logging and monitoring and the operational realities of protecting data. Moving into industry enabled me to take the theory and apply it, with a realisation that sometimes the theory does not work in practice and is either not useable or has to be adapted and changed. It was during my time at Vodafone working on a control roadmap that I first came across PQC, which led to me into the world of quantum and quantum security.

Quantum computing reframed the entire problem. Even before practical quantum machines existed, the *harvest-now, decrypt-later (HNDL)* threat model made it clear that adversaries didn't need a quantum computer today to break tomorrow's secrets. They only needed storage and patience. The threat wasn't hypothetical it was already operational. Data with long term value was already being collected, archived and stockpiled by actors who understood exactly what was coming.

This was the pivot point for me. Quantum wasn't a future technology to monitor

or a distant research topic to keep an eye on. It was a present day risk, exerting pressure on cryptographic assumptions that had held for decades. My background from teaching fundamentals, to understanding adversarial behaviour in policing, to training professionals in cybersecurity and working on large security projects. Meant I recognised the pattern immediately. Quantum was a structural shift in the threat landscape, and it demanded a new way of thinking and adapting or changing current security thinking and concepts for this new opportunity and threat.

## 11.1 Learning Quantum: A Pragmatic, Security First Approach

My entry into quantum computing didn't begin with wavefunctions, Hilbert spaces or the mathematical elegance that physicists might admire. I didn't come to quantum through the door of theoretical physics at all never mind via computer science. I approached it the same way I approached every emerging technology. As a teacher, security architect and adversarial thinker looking for how it worked and then how quantum could be used to attack or be attacked.

My inquisitive nature and wanting to understand the depth as well as the breadth led me to ask myself:

- What can it compute that classical systems cannot?
- How does it impact on cryptographic systems?
- What does quantum bring that we have not got already?
- How could it be used by adversaries?

These questions shaped my learning journey. More so with the initial research I was undertaking with my PhD looking at quantum computing from a security standpoint and what quantum security should look like. To do this I needed to understand quantum computing as a computational model with different scaling behaviours, one that alters the economics of attack and defence.

Once quantum computation is treated as controlled interference rather than abstraction, its impact is explicit. Shor's algorithm directly compromises RSA, ECC and all public-key systems built using this type of cryptography. Grover's algorithm also applies a pressure to symmetric cryptography, compressing effective key strength and invalidating legacy comfort around 128-bit targets.

The more I learned, the more obvious it became that quantum computing wasn't a scientific curiosity. It was also a threat vector with a new class of adversarial capability that demanded the same scrutiny, modelling and governance as any other high impact security risk. My learning path was therefore pragmatic, not theoretical. I focused on attack surfaces, cryptographic fragility and the real world implications of quantum advantage. However narrow or specialised it may be. Although without the same level of real world data due to quantum still being largely around theory as opposed to practical application and knowledge.

Quantum computing is best understood as the manipulation of probability amplitudes rather than classical bits. Classical computers operate on definite states, a bit is either 0 or 1. Quantum systems, by contrast, evolve over superpositions. Weighted combinations of 0 and 1, where each component carries a complex probability amplitude. These amplitudes can interfere with one another, reinforcing some outcomes while cancelling others. This interference behaviour is the source of quantum computational advantage. Despite this richer internal structure, any measurement of a qubit still yields a definite classical result: 0 or 1. The strength of quantum computing doesn't come from science-fiction notions of infinite parallelism. It comes from three very specific physical phenomena that translate into computational advantages:

- **Superposition:** the ability to represent many possible states compactly within a single quantum register. This isn't "doing everything at once"; it's encoding a landscape of possibilities in a way classical systems cannot.
- **Entanglement:** the ability to correlate qubits so strongly that the state of one instantly constrains the state of another, no matter how far apart they are. In computational terms, entanglement creates relationships between data elements that classical systems cannot replicate. Qubits do not share a state; they share a constraint on possible joint outcomes.
- **Interference:** the mechanism that allows quantum algorithms to amplify the probability of correct solutions while suppressing incorrect ones. This is the real engine of quantum advantage: shaping the probability landscape so that the right answers become more likely.

These principles are elegant, but their relevance to security is brutally practical. Quantum computers do not threaten cryptography because they are "faster." They threaten cryptography because they compute differently. They exploit mathematical structures in ways classical machines fundamentally cannot.

Shor's algorithm doesn't succeed because it runs quickly; it succeeds because it uses interference to expose periodicity in number-theoretic problems that underpin RSA and ECC. Grover's algorithm doesn't magically accelerate brute force; it reshapes the search landscape so that the correct key becomes quadratically easier to find.

Once you understand quantum computing as a fundamentally different computational model, with its own scaling behaviours and operational principles. The security implications become impossible to ignore. Quantum systems are not faster versions of classical computers; they implement a distinct method of computation with unique strengths and weaknesses. They do not replace classical architectures but instead provide specialised speedups for particular problem classes that are intractable or slow on classical machines.

From a security perspective quantum isn't just an interesting computational model. It introduces new adversarial capabilities, breaks long standing assumptions and forces organisations to rethink how they protect data across its full lifecycle. The threat is not singular or isolated. It is multidimensional.

Quantum computing exerts pressure in three major areas:

## Cryptographic fragility

The most immediate and widely discussed impact is the fragility of classical public-key cryptography. Algorithms that have underpinned secure communication for decades such as RSA, ECC and Diffie-Hellman, These are mathematically vulnerable to Shor’s algorithm. This isn’t a theoretical weakness; it is a structural one. Once quantum machines reach sufficient scale and stability, these algorithms collapse.

Symmetric cryptography fares better as Grover’s algorithm effectively halves the security strength of symmetric keys, forcing organisations to double key sizes to maintain equivalent protection. AES-128 becomes insufficient. AES-256 becomes the new baseline.

## Temporal asymmetry

Quantum also introduces a profound shift in how we think about confidentiality over time. In classical security models, encryption protects data at the moment it is transmitted or stored. Once encrypted, it is assumed to remain confidential for as long as the cryptographic scheme remains secure.

## 11.2 Quantum Breaks This Assumption

Adversaries can and already do harvest encrypted data today with the intention of decrypting it in the future once quantum capabilities mature. This “harvest-now, decrypt-later” strategy turns long-lived data into a liability. Sensitive information with multi decade value medical records, genomic data, state secrets, intellectual property, financial archives become vulnerable long before quantum computers exist at scale.

### Harvest Now, Decrypt Later (HN DL)

HN DL describes the collection of encrypted data today with the intention of decrypting it once large-scale quantum computers can break classical public key encryption. This is a confidentiality centric threat: the attacker’s objective is to expose information that is currently protected by RSA, ECC, or other vulnerable schemes. HN DL matters most for:

- Long-retention data
- Sensitive archives
- Communications requiring multi-decade secrecy
- National-security and diplomatic material

HN DL is widely recognised, but it is only the beginning.

## Harvest Now, Forge Later (HNFL)

HNFL shifts the target from encryption to digital signatures. Instead of collecting ciphertext, adversaries collect signed artefacts software updates, firmware images, certificates, audit logs, legal documents, financial transactions with the intention of forging signatures once quantum computers can break classical signature schemes. HNFL is an integrity threat. It enables:

- Forged software updates
- Counterfeit certificates
- Manipulated audit trails
- Fraudulent transactions
- Retrospective rewriting of “trusted” records

Where HNFL compromises secrecy, HNFL compromises the truth of digital artefacts.

## Trust Now, Forge Later (TNFL)

TNFL extends HNFL into the domain of trust infrastructure. Systems today rely on signatures, certificates, and identity proofs that will become forgeable in the future. TNFL captures the risk that trust anchors valid today become counterfeit tomorrow. TNFL is an authenticity and trust-chain collapse threat. It enables:

- Impersonation of devices, services, or authorities
- Fabrication of entire PKI chains
- Subversion of firmware-update ecosystems
- Undetectable supply-chain compromise
- Manipulation of safety-critical systems that rely on signed commands

TNFL is particularly dangerous for operational technology (OT), where trust in signed commands directly affects physical processes.

## Deploy Now, Exploit Later (DNEL)

DNEL addresses a different dimension, lifecycle exposure. Many devices deployed today IoT sensors, industrial controllers, medical implants, EV chargers, automotive ECUs will remain in service long after quantum computers can break their embedded cryptography. If these devices cannot be upgraded to post-quantum primitives, they become permanently vulnerable.

DNEL is a cyber-physical safety threat. It enables:

- Impersonation of devices that cannot rotate keys
- Compromise of long lifecycle infrastructure
- Persistent footholds in OT environments

- Attacks on vehicles, energy systems, and medical devices
- Exploitation of unpatchable hardware trust anchors

DNEL is the most structurally difficult threat to mitigate because it arises from design choices made years before quantum exploitation becomes possible.

## Governance and migration risk

The transition to post-quantum cryptography (PQC) is not a software update. It is a systemic migration that touches every layer of the digital ecosystem. Every protocol, every embedded device, every certificate chain, every key exchange mechanism, every cryptographic library. All must be re-evaluated and in many cases, replaced. This introduces governance challenges on a scale the industry has not faced since the early days of the internet. Organisations must:

- inventory cryptographic assets
- identify long lived data
- assess protocol dependencies
- test PQC performance and interoperability
- plan phased migrations
- manage hybrid cryptographic environments
- ensure supply chain readiness

The complexity is immense. The risk of misconfiguration, partial migration or inconsistent implementation is high. And because cryptography is deeply embedded in systems, the migration will take years, not months. The identification of cryptographic assets alone is a huge task for some organisations.

## Building Q-SLICE and benchmarking quantum reality

As my understanding of quantum computing deepened, it became clear that the field suffered from a chronic imbalance with an element of hype and far too little measurement. Claims about “quantum advantage,” “breakthroughs” and “revolutionary performance” were often made without reproducible evidence, consistent benchmarking or any meaningful connection to real-world security impact. Security cannot be built on marketing language or theoretical potential. It must be grounded empirical behaviour. With little to know data with quantum still in its early stages an element of assumptions has to be made. Which ultimately when dealing with uncertainty within security is what an organisation must undertake. However, these assumptions are often built around data and knowledge assimilated over years.

This realisation led me to develop the Q-SLICE a threat model modular designed to look at quantum hardware and quantum-inspired systems with the same rigour we apply to classical security architectures using threat models like STRIDE. Q-SLICE

was built to answer a simple but critical question. What does quantum hardware actually do under load, under noise and under cryptographic pressure?

To answer that, I created a Q-SLICE Threat Harness using QSKIT modules as a small program to a range of foundational characteristics both in a simulator and on IBM Quantum Cloud. This covered elements such as:

- **Coherence** how long qubits retain usable quantum states before decoherence destroys computational value.
- **Fidelity** the accuracy of operations, gates and circuits, and how closely real hardware behaviour matches theoretical expectations.
- **Noise behaviour** the patterns, sources, and propagation of noise across circuits, and how noise interacts with algorithmic structure.
- **Scaling characteristics** how performance, error rates and resource requirements evolve as circuits grow in depth and width.
- **Backend variability** the differences between quantum hardware providers, architectures, and runtime environments, and how these differences affect reliability.
- **Cryptographic workload fragility** how quantum systems behave when executing circuits derived from cryptographic primitives, including those relevant to Shor-style and Grover-style attacks.

## 11.3 The Purpose of All This is Straightforward

Replace hype with measurement. Replace assumptions with evidence. Quantum security cannot be governed by optimistic projections, vendor claims or theoretical limits that assume perfect qubits and noiseless circuits. It must be governed by what hardware actually does. Its fragility, its variability, its scaling behaviour and its real-world impact on cryptographic workloads. Q-SLICE exists to bring clarity to a field clouded by speculation. It provides the empirical grounding needed to make informed decisions about risk, migration and long-term security strategy.

### The human side: learning, misconceptions, and the reality of quantum

As I progressed deeper into the quantum domain, I quickly realised that the technical challenges were only half the story. The other half, often the more difficult half, was navigating the misconceptions, oversimplifications and confident misunderstandings that surround quantum computing. The field attracts hype in a way few technologies do AI is another with similar issues. I encountered the same recurring beliefs everywhere: in classrooms, in industry discussions, in vendor presentations and even in policy conversations.

## “Quantum computers will solve everything”

This idea collapses the moment you examine what quantum machines are actually good at. Quantum computing is powerful, but it is narrowly powerful. It excels at specific classes of problems. Optimisation, simulation and certain algebraic structures but it is not a universal accelerator. Most workloads will never benefit from quantum hardware and many will perform worse.

## “Quantum is decades away”

This belief persists because people equate “breaking RSA” with “quantum maturity.” But quantum risk doesn’t begin at the moment RSA falls. It begins the moment adversaries start harvesting encrypted data for future decryption. It begins when organisations realise their long lived data will outlast their cryptography.

## “PQC is just a drop in replacement”

This is perhaps the most dangerous misconception. Post quantum cryptography is not a simple swap of algorithms. It changes key sizes, message sizes, handshake patterns, performance characteristics and protocol assumptions. It affects embedded systems, constrained devices, certificate chains and supply-chain dependencies. PQC is necessary, but it is far from a ‘find and replace’ as the impact and feasibility over using PQC has to be assessed.

Each of these beliefs falls apart under scrutiny, but they persist because quantum computing sits at the intersection of science, engineering, policy and imagination. People fill gaps in understanding with optimism, pessimism or mythology when it should always be evidence based.

My journey through quantum has therefore been as much about cutting through noise as it has been about learning the technology itself. I’ve had to build a grounded, evidence based understanding that resists hype, avoids fatalism and focuses on measurable reality. That mindset has shaped everything I’ve done since from how I evaluate quantum hardware, to how I design threat models, to how I communicate risk to organisations that must prepare for a future which is arriving unevenly, but undeniably.

## 11.4 My View of the Quantum Future

As I look ahead, my perspective on quantum computing is shaped not by hype or speculation, but by the patterns that emerge when you analyse the technology through a security and evidence driven lens. Quantum computing is transformative, but not in the way popular narratives often suggest. It will not replace classical computing. It will not become a universal machine capable of running every workload faster or more efficiently. It will not render traditional architectures obsolete.

Quantum computing will instead evolve into a specialised accelerator, much like GPUs, TPUs and other domain specific architectures that complement. Rather than replace classical systems. Its strengths lie in narrow but profoundly important domains:

**Optimisation.** Quantum and quantum inspired algorithms have the potential to reshape how we approach complex optimisation problems, from logistics to finance to energy systems. These are areas where even marginal improvements can have enormous real world impact.

**Simulation.** Quantum systems excel at simulating quantum phenomena, something classical machines struggle with fundamentally. This capability will drive breakthroughs in chemistry, physics and biological modelling, enabling discoveries that were previously out of reach.

**Cryptanalysis.** This is the domain with the most direct security implications. Quantum computers will not break all cryptography, but they will break the cryptographic assumptions that underpin global secure communication. This alone makes quantum a strategic capability.

**Materials Science.** The ability to model molecular interactions with high fidelity will accelerate the development of new materials, pharmaceuticals and energy technologies. These advances will have geopolitical, economic and security consequences.

In other words, quantum computing will not be a general purpose replacement for classical systems. It will be a powerful tool in the domains where its computational model aligns with the structure of the problem and can solve a problem a classical computer cannot.

For the security community, this distinction matters. Quantum is not a curiosity to be monitored from a distance. It is a strategic capability, one that will shape national security, economic competitiveness and the long term viability of cryptographic systems. Treating it as anything less would be a profound mistake.

The future of quantum is neither utopian nor catastrophic. It is specialised, impactful and unavoidable. Preparing for that future requires clear thinking, disciplined measurement and a willingness to confront the fragility of our current security foundations across people, process and technology.

## 11.5 The Security Imperative: What Must Happen Next

Understanding quantum computing is only the first step. The real challenge and the real responsibility lie in preparing for its impact. Quantum security is not something that can be deferred until hardware reaches a certain threshold or until a headline announces a breakthrough. By the time quantum machines are capable of breaking classical cryptography, it will be far too late to begin planning. Preparing for quantum requires a structured, disciplined and long term programme. Organisations

must shift from reactive security to proactive security, recognising that quantum risk is already shaping adversarial behaviour and long-term data exposure.

Several actions are essential.

- **Inventory cryptographic assets.** Most organisations do not have a clear understanding of where cryptography is used within their systems. Encryption is embedded in protocols, applications, firmware, APIs and supply chain components. Without a comprehensive inventory, migration becomes guesswork. Identifying algorithms, key lengths, certificate chains and protocol dependencies is the foundation of any quantum readiness strategy.
- **Identify long-lived data.** Not all data is equal. Some information loses value quickly; other data retains sensitivity for decades. Medical records, genomic data, financial archives, intellectual property and government communications fall into this category. These are the assets most vulnerable to HNDL attacks. Organisations must classify data based on confidentiality lifespan, not just operational importance.
- **Adopt post-quantum cryptography (PQC) standards.** The transition to PQC is no longer theoretical. NIST has selected algorithms and vendors are beginning to integrate them. Organisations must begin adopting PQC in a phased, controlled manner. This includes hybrid approaches that combine classical and quantum resistant algorithms to ensure continuity during the transition.
- **Test PQC performance and interoperability.** PQC algorithms behave differently from classical ones, introducing larger key sizes, altered handshake patterns and distinct performance characteristics. Before deployment, organisations need to understand how these differences manifest in practice by examining latency impacts, bandwidth demands, protocol compatibility, behaviour on constrained devices and interoperability across vendors and platforms. This testing phase is essential; without it, the risk of operational disruption becomes unavoidable.

Quantum migration itself is not a patch but a multiyear transformation and organisations must approach it with structured planning rather than ad-hoc updates. Effective roadmaps prioritise high risk systems, define phased deployment strategies, coordinate with supply chain partners, update governance and compliance frameworks and ensure that technical teams receive appropriate training. A well-constructed roadmap turns quantum readiness from an abstract concern into a manageable, long-term programme of work.

At the same time, organisations must continuously monitor the progress of quantum hardware. Capability is advancing unevenly across vendors, architectures and research groups. Making it necessary to track qubit counts, error rates, coherence times, algorithmic developments and cryptanalytic milestones. This monitoring ensures that migration efforts remain aligned with real world progress rather than speculation or hype.

Finally, quantum threat modelling must be integrated into existing risk assessments. Traditional models assume classical computational limits, but quantum capability

alters those assumptions in fundamental ways. Organisations need to consider HNDL, accelerated key search capabilities, hybrid classical quantum attack chains and the supply chain vulnerabilities introduced by PQC adoption. Threat modelling must evolve alongside the technology to remain relevant and effective.

## **QUANTA: mapping quantum adversarial capability**

While Q-SLICE was built to measure the behaviour of quantum hardware under realistic load, noise and cryptographic pressure, it quickly became clear that measurement alone was not enough. Understanding quantum systems empirically is essential, but security requires more than observation. It requires modelling the adversary their capabilities, constraints, incentives and attack pathways. This is where QUANTA emerged.

If Q-SLICE answers the question “What does quantum hardware actually do?”, then QUANTA answers the equally critical question “What can an adversary actually do with quantum capability?”

Quantum security cannot be governed by classical threat models alone. Quantum introduces new computational behaviours, new scaling patterns, and new asymmetries that reshape the economics of attack. QUANTA was designed to capture these shifts in a structured, repeatable and evidence driven manner. Helping to score and then introduce relevant controls to mitigate risk.

## **The purpose of QUANTA**

QUANTA is a quantum-aware cybersecurity control framework and scoring system, operating in the same governance space as ISO 27001 and the NIST Cybersecurity Framework. Its purpose is to define the controls, maturity levels, and organisational capabilities required to manage the risks introduced by emerging quantum technologies.

Where classical frameworks assume stable computational limits and static cryptographic strength, QUANTA introduces quantum specific control domains, including:

- Cryptographic transition and crypto agility controls to manage the impact of quantum-accelerated cryptanalysis.
- Temporal risk and long lived data controls aligned to HNDL, HNFL, TNFL and DNEL adversary models.
- Hybrid resilience controls that account for attack chains combining classical and quantum resources.
- Resource bounded capability tracking to ensure controls scale with realistic, not hypothetical, quantum progress.
- Incremental governance, recognising that quantum capability evolves unevenly and adversaries exploit each stage.

QUANTA provides a structured, evidence driven framework for developing quantum resilient security programmes. At its core, it defines a comprehensive set of control families that span cryptography, system architecture, governance, monitoring, supply chain assurance and lifecycle management. These controls give organisations a coherent way to embed quantum-aware security practices across both technical and organisational domains.

The framework also introduces maturity tiers that assess readiness across classical, hybrid and post quantum states, allowing organisations to understand their current posture and plan their progression. Complementing this, QUANTA includes scoring mechanisms that quantify exposure, preparedness, and compliance with quantum-era requirements, enabling consistent measurement and benchmarking.

A key feature of QUANTA is its lifecycle aware approach, recognising that systems deployed today will remain operational long enough to be exposed to future quantum threats. To address this, the framework defines controls that explicitly manage long lived data risk and deferred exploitation scenarios. It also provides transition pathways that guide organisations from current state cryptography toward architectures aligned with post-quantum standards.

## **Quantum security is not a switch**

It is a programme strategic, long term and essential. Preparing for quantum is not about predicting the exact moment when cryptographic collapse will occur. It is about ensuring that when that moment arrives, organisations are not caught unprepared. The work begins now and the organisations that act early will be the ones that maintain trust, resilience and operational continuity in the quantum era.

Quantum computing is not a distant horizon or a speculative threat. It is a present day pressure on the assumptions that underpin digital trust. My journey through quantum security has been shaped by a commitment to evidence, measurement and adversarial realism. Principles that guided the creation of the Q-SLICE threat model and QUANTA cybersecurity framework. Preparing for the quantum era is not about predicting the exact moment of cryptographic collapse; it is about building the resilience to withstand it whenever it arrives.



# Chapter 12

## Speaking Two Languages

### *Quantum computing's real path into finance*

Joe Ghalbouni

#### 12.1 From Momentum to Responsibility

Around 2020, quantum computing reached a tipping point. Not in the sense of technical maturity—far from it—but in visibility, credibility, and institutional attention. What had lived for decades in academic journals and specialized laboratories suddenly entered boardroom conversations, investment memos, and strategic roadmaps.

For many researchers, this moment represented validation. For me, it felt more like responsibility.

I had always believed that emerging technologies only become meaningful when someone is willing—and able—to translate them into the language of the sectors they are meant to impact. Quantum computing was no exception. Its promises were vast, but its vocabulary was inaccessible. Fidelity, circuit depth, entanglement structure—these were not the terms that would determine whether quantum would matter to real institutions. If quantum computing was going to leave the lab, someone needed to speak both languages fluently: the language of quantum mechanics and the language of industry.

That realization shaped everything that followed.

---

This chapter is a featured contribution to the monograph *A Portrait of Quantum Technologies in Finance*, edited by Oswaldo Zapata, PhD, and published by The Quantum Finance Boardroom (2026). All rights reserved by the chapter's author.

## 12.2 Why Finance Became the Proving Ground

Quantum computing has been associated with many sectors: chemistry, materials science, logistics, energy, pharmaceuticals. All valid. Yet finance stood out early, and not by accident.

First, the nature of financial problems maps naturally onto the classes of problems where quantum approaches are being explored: optimization, sampling, machine learning, and stochastic modeling. These are not foreign abstractions to finance—they are its daily bread.

Second, finance has both the capital and the institutional muscle to experiment early. It can afford to explore emerging technologies before they are fully mature, not out of curiosity, but because early informational or computational advantages compound quickly.

Third—and often overlooked—finance has a long tradition of importing methods from physics and mathematics. Systematic strategies, risk models, Monte Carlo methods, and signal extraction all emerged from this cross-pollination. The idea that techniques developed in one scientific domain could be repurposed for markets is not new. Quantum computing entered an ecosystem already predisposed to abstraction and modeling.

In that sense, finance was not chasing quantum. Quantum was arriving at a sector already primed to engage with it.

## 12.3 The First Reality Check: A Change of Dialect

What surprised me most upon entering the financial world was not skepticism—it was fluency of a different kind.

Quantum, in industry, is not discussed in terms of gate fidelity or entanglement depth. It is discussed in terms of return on investment, scalability, integration, governance, and timelines.

Conversations revolve around how a proof of concept becomes a minimum viable product, how it fits into high-performance computing (HPC) environments, how it interfaces with existing workflows, and how it survives regulatory scrutiny.

This was not a rejection of quantum mechanics; it was a change of dialect.

In academia, success is measured by novelty, correctness, and elegance. In finance, success is measured by robustness, repeatability, and value creation under constraints. The same technology is being evaluated—but against a completely different coordinate system.

Learning to operate in that dialect was not optional. It was the prerequisite for being taken seriously

## 12.4 From Education to Execution

What began as education quickly became execution.

Early engagements focused on explaining what quantum computing could and could not do, cutting through both excessive optimism and reflexive dismissal. That phase did not last long. Curiosity turned into experimentation. One proof of concept became several. Isolated explorations evolved into coordinated efforts across teams.

At that point, something important shifted. The conversation stopped being about “whether” quantum computing would matter and became about “how” and “when.” Proofs of concept were no longer isolated demos; they were expected to show a path to integration, performance comparison against classical benchmarks, and a plausible route to production.

As these efforts matured, an inseparable companion emerged: risk.

## 12.5 When Opportunity and Risk Converge

Quantum computing ceased to be a purely innovation-driven topic the moment it started entering enterprise environments.

Once quantum hardware began to materialize—once quantum computers were no longer hypothetical—it became obvious that their implications extended beyond optimization speedups or modeling capabilities. Security, cryptography, and long-term data protection moved to the foreground.

This realization did not come out of nowhere. Work on post-quantum cryptography had begun years earlier, and the timeline was not accidental. The same year standardized efforts around quantum-resistant cryptography gained momentum, the first cloud-accessible quantum processors appeared. The message was implicit but clear: this was no longer theoretical.

For financial institutions, this reframed quantum computing entirely. It was no longer just about gaining an edge. It was also about managing exposure. Delaying action carried its own risks.

Quantum became both an opportunity and a liability—and strategies had to account for both.

## 12.6 Finance’s Paradox: Pioneering Yet Constrained

From the inside, finance’s relationship with quantum computing is paradoxical.

On one hand, financial institutions are often pioneers. They explore early, invest

in talent, and experiment seriously. On the other hand, they are constrained by their own success. Finance expects returns—and it expects them on relatively short horizons.

That expectation is not unreasonable. But it creates a structural problem: genuine, nuanced quantum expertise often struggles to be heard amid louder, more speculative narratives. In an environment where many technology providers must amplify their claims to secure funding, signal and noise blur together. Institutions searching for real value find themselves facing a haystack with too many needles claiming to be the one.

The result is not rejection, but hesitation.

## 12.7 A Misalignment of Incentives

This hesitation reflects a deeper misalignment.

Quantum technology providers are incentivized to emphasize novelty and long-term potential. Financial institutions are incentivized to prioritize clarity, integration, and measurable outcomes. When vendors focus on “quantumness” rather than usefulness, adoption stalls—not because institutions are conservative, but because they are rational.

Successful technologies rarely succeed by forcing users to abandon existing workflows. Operating systems offer a useful analogy: organizations often stick to older systems not because they are inferior, but because replacing them would disrupt governance, compliance, and operations at scale.

Quantum computing is no different. Introducing it as a radical replacement rather than an incremental enhancement triggers resistance—not fear of complexity, but fear of uncontrolled consequences.

Governance, therefore, is not an obstacle to adoption. It is the condition for it.

## 12.8 The Hybrid Future Is Not a Compromise

One conviction has become unavoidable: quantum computing will not arrive as a standalone revolution. It will arrive as part of hybrid systems embedded within HPC environments.

This is not a retreat from ambition—it is realism. A powerful quantum processor without the surrounding ecosystem, tooling, and human expertise is meaningless. What matters is not raw capability, but applicability.

Understanding what to use quantum for, how to combine it with classical systems, and where it actually provides incremental value matters more than hardware specifications alone. This hybridization is already visible in conversations across

industry and policy circles, and it reflects how complex technologies historically enter production: quietly, incrementally, and through integration rather than disruption.

## 12.9 Talent, Not Qubits, Is the Bottleneck

If there is a true bottleneck today, it is talent.

Not quantum physicists in isolation, and not traditional software engineers alone—but people who understand both quantum computing and the domain in which it is applied. At least initially, this hybrid talent is essential for training internal teams and ensuring that quantum does not become an isolated, expensive capability.

If quantum computing remains usable only by specialized external teams, return on investment will suffer. The cost of maintaining large, disconnected expert groups for a single technology is unsustainable. Long-term value requires internalization, not dependence.

This is not unique to finance, but finance feels it early.

## 12.10 Governance Will Decide the Winners

In the end, technology compliance and governance will likely determine which solutions succeed.

Institutions will choose technologies they can explain, audit, integrate, and regulate—not necessarily those that are theoretically most powerful. Trust, traceability, and control will matter as much as performance.

Quantum computing will follow this same path. Its adoption will not be decided by qubit counts alone, but by how well it fits into existing institutional realities.

## 12.11 The Role That Had to Exist

Looking back, the most important realization was not technical.

There was a missing role—someone able to take the technology from point A to point B. Someone who could open eyes without inflating expectations. Someone who could translate without diluting. Someone who respected integration paths as much as innovation.

Not a gatekeeper. Not a cheerleader. A bridge.

Quantum computing does not need louder voices. It needs clearer ones.



## Chapter 13

# My Roundabout Journey into Quantum

### *How not to do it!*

John Galani

I was born in London and grew up in Paris, receiving a French education. I am however of Greek origin and finished my French high school in Athens in 1991. At that time in Athens — and I am dating myself here — being a "computer geek" carried none of the cachet it does today. Programming and basic circuit design did not impress quite like football, basketball, or social fluency on a dance floor. I therefore kept my love of technology close to my heart rather than pursuing it academically.

I studied economics and business finance in London, falling in love — as a good Greek — with shipping, which I pursued through internships and later at the Master's level. It formed the backbone of the first part of my career.

What followed was, on the surface, a tour through shipping, commodities, real estate, and aviation. The connecting thread, though, was always the same: I gravitated toward situations that required structuring — whether building something new or restructuring something broken. That instinct, more than any single sector, has defined my career.

The financial crisis of 2009 was a wake-up call. It pushed me to revisit the fundamentals of business and value creation, and to ask whether finance had outgrown the scientific discoveries that had powered human progress for millennia.

That question became personal when I was working in Sierra Leone on a mining exploration project for a family office, taking it from structuring to sale over three

---

This chapter is a featured contribution to the monograph *A Portrait of Quantum Technologies in Finance*, edited by Oswaldo Zapata, PhD, and published by The Quantum Finance Boardroom (2026). All rights reserved by the chapter's author.

years. There, I rediscovered how transformative basic technology — in transportation, in medicine — could be on human lives. The lesson grew sharper still when I had the privilege of meeting scientists from the CDC and their European counterparts as Ebola devastated the country and the surrounding region.

After a further stint growing a commodity trading arm across the Middle East and Europe, I wanted to immerse myself again in my first love: technology. I enrolled in a fintech course at Oxford Business School (discreetly, since I had no wish to advertise to my employer that I had such "free time"). It proved a blessing in disguise: two years later I was named COO of the newly minted, NASDAQ-listed fintech I had helped to spin off, where I built and structured the operating teams that carried the company through its listing and into its first years as a public business. From there I moved to a Singapore-based, MAS-regulated platform that required restructuring — the same instinct applied to the other side of the cycle.

I did not want to stop there. During COVID I picked the most improbable topic I could find, simply because I could not get my head around it. Enter a certain cat — Schrödinger's — which may still be both alive and dead. What is certain is that the course I followed at Delft University was an eye-opener, particularly in week one, when my professor told us to stop trying to "understand" superposition or entanglement. Curiosity may have killed that cat, but I had nine lives and pressed on.

From there I began showing up at conferences and tentatively speaking with scientists, mindful that I was neither one of them nor knowledgeable about quantum. What I did know was business — investing, plans, startups, and structuring — and I had already watched a fair number of technology companies rise and fall across cycles and geographies.

I have always felt like a foreigner, in every country I have lived in and every sector I have entered. With time, I have come to see that outsider's vantage point as a quiet asset: it forces you to listen first and to ask the obvious questions that insiders no longer think to ask. My affection for quantum is, I suspect, partly inherited. My father understood early that computers were the future and installed the first one on a vessel — though he himself never grew comfortable with a PC. I watched radio operators on those same vessels lose their livelihoods when satellite communications arrived, unable to move on. That memory has kept me close to whatever is coming next — partly out of professional discipline, partly so I can remain relevant to my children.

I became a member of UKQuantum, the UK's industry body for quantum computing, and along the way met genuinely interesting people, among them Oswaldo Zapata — of whom more shortly.

By the time that Singapore restructuring was drawing to a close — the platform on stable footing — I felt that AI demanded a serious working understanding: not how to spend money on it, but how to make money from it. I took several courses in 2023 and 2024 and have since put that learning to work on a number of projects, including a large AI build for investment funds. Just as importantly, I have used it to teach my son.

Quantum, however, is where my conviction sits. If I was late to the most successful fintech wave and arguably late to AI, I am early to quantum. My heart is there, my mind is set, and I want to make it count.

How does someone like me — not a scientist, a newcomer to a niche field — go about it? As a mentor at City University, I work with people entering the job market with no prior knowledge, experience, or contacts; I drew on the same playbook for myself.

First, I looked closely at quantum, clearly separating the longer-term hardware prospects from the nearer-term opportunities in sensing, networking, and security.

Second, I studied the investment ecosystem and how these startups reach Series A. There are strikingly few venture investors active in the space, and fewer still with real technical fluency.

Third — because curiosity, like our journeys, should never end — I kept learning. By the time this is published I should have completed my second formal course, this time with MIT.

Having mapped the landscape and spoken with a number of VCs, I selected one in which I both invested personally and brought in outside investors.

Investing, though, is a starting point rather than an end point. What I would most like to contribute to quantum is what I contributed to fintech a decade ago: the operational structuring that turns a technical breakthrough into a business — the teams, the processes, the discipline to scale. Restructuring has taught me what good structure looks like; the harder and more rewarding work, I think, is structuring on the way up. That is where I hope to be useful here.

Ever since I met Oswaldo, I have felt he understood my journey, my passion, and my limitations, while always encouraging me further. I am delighted that he and his colleagues are leading The Quantum Finance Boardroom — bringing insiders and outsiders together, drawing in people like me with a passion for quantum and expertise in finance, and advancing a field we all love and cherish. This journey is only at the beginning. People of every background and level are welcome. We are going places, and I am enjoying every moment.



## Chapter 14

# Quantum Is the Next AI Moment, Only Bigger

*Why business leaders and finance teams need quantum literacy now*

John Riley III

### 14.1 The AI Déjà Vu Moment

AI blindsided business teams that didn't build AI literacy early. When ChatGPT rolled out to the public in November 2022, it hit many like a punch in the mouth. The challenge now is that leaders and business units are still trying to implement AI across their organizations while navigating the internal shifts required to adopt it effectively. Most are still trying to get their legs under them, excited about the upside but dizzy with how fast it's moving. Even though AI has been around for decades, this easy access to it for everyday use shocked a lot of people. With anything new, the fastest way to understand it is to get educated. Instead, a lot of folks, including business leaders, stuck their heads in the sand.

Historically, many organizations followed the old school mindset: "Don't be an early adopter." But over the past 10 to 15 years, that thinking has started to shift. Emerging tech is moving faster than anything we have seen in our lifetimes. Some leaders and business teams are finally realizing that being early adopters isn't just smart, it's necessary. It's the only way to get ahead of the curve instead of constantly playing catch up.

In June 2024, the United Nations General Assembly declared 2025 the International

---

This chapter is a featured contribution to the monograph *A Portrait of Quantum Technologies in Finance*, edited by Oswaldo Zapata, PhD, and published by The Quantum Finance Boardroom (2026). All rights reserved by the chapter's author.

Year of Quantum Science and Technology (IYQ 2025). To me, that's a big signal. Quantum isn't just living in a hype cycle anymore. It's getting real momentum. While a lot of people are still blinded by AI, and honestly just trying to catch up to it, countries and companies were already building quantum strategies in the background. I got caught off guard by how fast quantum was moving a few years ago, until I started researching what companies were working with it and talking to people in the quantum space across different countries. That's when it hit me: you can't be complacent in emerging tech anymore. Something else is always on the way, and right now, that's quantum.

With AI becoming mainstream only a few years ago, many haven't grasped the lesson about the importance of preparing early for emerging tech advancements. Out of all the disruptive emerging technologies that have come along in the past few years, like blockchain and AI, quantum is going to be the one that overshadows them with the fast pace it's moving at. This won't be a punch in the mouth like AI did to everyone, but a bullet train that will blow past them and suck the air out of their lungs. If they don't have their ticket ready by having early awareness, education, and preparedness to hop on that train, it's not stopping or slowing to wait for them to be comfortable before they are sitting down in their seat.

As many organizations try to adopt their digital transformation strategies, they need to already be starting to look at how quantum is going to fit in their organizations and preparing their teams with basic education. Quantum won't be used for everything in business. Everyday tools like spreadsheets, accounting, and routine reporting will still run best on classical systems, and we're not going to see a quantum computer sitting on everyone's desk.

The time is now for leaders and business teams to be open to a new mindset shift and to be very intentional about being proactive in preparing for quantum, and not waiting to see if they have more time like they did with AI.

## 14.2 Making Quantum Easy to Digest

To start understanding what quantum is and what advancements it will bring doesn't require you to be a quantum scientist to understand its value. I've seen over the years how many companies and leaders will rush into a new technology and expect their technical teams to own it. With quantum, understanding the business case of where it will best fit first in your organization is essential for being able to prepare for it early. Quantum will not replace your classical computers currently being used, but they will be leveraged in a hybrid approach.

When it comes to giving a definition of what quantum is, you can get various definitions depending on who you are speaking with. Some will resonate and some won't based on the person. This has happened with many of the emerging technologies that have come around in the past few years, like blockchain and AI. I think this has been a reason why many people who are not technical have challenges understanding new technologies. They are given a technical definition they can't connect with and will shut down and say, this must be something the IT department needs to focus on.

Understanding the difference between classical computers and quantum computers helps with connecting to the definition of what quantum is. Classical computers are based on what they call binary digits, or bits. A bit holds either a 1 or 0, never both. All modern day devices rely on this binary logic. It's like a switch. It's either off or it's on. This approach is fast and reliable based on the current technology available.

Quantum computing brings a new approach that wasn't practically accessible until recently. Quantum uses qubits. A qubit can be 0, 1, or both at once (superposition). For the right problems, that lets quantum systems explore many possibilities in parallel, far faster than classical computers, and even supercomputers, can today. That opens the door to solving complex problems that could take decades on classical systems. In 2019, Google's quantum processor ran a specialized task in about 200 seconds that Google estimated would take a leading supercomputer far longer. Whether or not you like the phrase "quantum supremacy," the point stands. Quantum can do certain things differently, and that gap is closing as hardware improves.

Advancements are being made every month by quantum companies, getting us closer to what they call Q Day.

Here's how I explain quantum to business leaders in a way that's easy to picture. Imagine you walk into a library with six trillion books and you have to find one specific title. If you search one book at a time, it could take you a thousand lifetimes. Now imagine you walk in with six trillion versions of yourself. Each one grabs a different book at the same time. One of you finds it immediately. That's the intuition. Quantum can evaluate many possibilities in parallel for the right problem.

I think a big component of understanding quantum easily is trying to connect with a definition that will resonate with you in a way that makes you curious to want to learn more. If the definition is too technical for you or confuses you, ask more questions. Try to find one that you will be able to tell others so they can understand it easily as well.

### **14.3 Convergence Imperative: AI, Quantum, Cyber & Blockchain Are Becoming One Strategy**

Quantum won't show up alone. It's converging with what every finance organization already cares about. Over the past few years, financial services have been early adopters of emerging tech like blockchain, AI, and cybersecurity because of strict regulations, the need for faster and more accurate analysis of massive data sets, and the mandate to protect that data. It's also a reminder that these technologies shouldn't live in silos. They need to work together to unlock more powerful solutions to business challenges.

These four technologies are starting to "stack" on top of each other, and when they do, they will unlock real world advantages. Quantum combined with AI is like giving AI a new turbo tool for the hardest problems. Things like finding the best route for logistics, optimizing supply chains, improving portfolio decisions, or running deeper simulations for drug discovery and advanced materials. Cyber and quantum represent

the urgency play. Powerful future quantum machines could crack today’s common “public key” locks. Post Quantum Cryptography (PQC) is the practical fix, allowing new locks designed to hold up even in a quantum era. This isn’t theoretical anymore. The U.S. government has an active PQC migration program. The U.S. National Security Agency (NSA) has published its quantum resistant CNSA 2.0 requirements for national security systems, and the UK and Germany have published official guidance and timelines and recommendations to move organizations toward PQC. Finally, blockchain and quantum are about long term trust. If quantum changes digital signatures, it impacts wallets, signing, and key management. The benefit is upgrading how we protect identity and transactions now so data and records still verify as authentic years down the road, especially for long lived data like government records, healthcare data, and critical infrastructure logs.

With 2025 being the International Year of Quantum Science and Technology, the awareness push can’t stop at a one year banner. A lot of folks are already treating 2026 as a “Year of Quantum Security” (#YQS2026), because the security conversation is getting louder and the preparation work is real.

Earlier I mentioned how finance organizations have led the way in adopting emerging technology, and quantum is the next one they have been focused on. The global financial sector is undergoing a profound technological transformation, driven by the convergence of the top three major disruptive technologies: blockchain, artificial intelligence (AI), and quantum computing. While blockchain and AI are already integrated into core operations, from payment systems and fraud detection to personalized banking, quantum remains in the exploratory and pilot phase.

Financial institutions are primarily focusing their quantum efforts on complex optimization problems, risk modeling, and, critically, developing quantum safe cryptography to protect against future threats.

## 14.4 Leading Financial Institutions in Tri-Technology Adoption

The following table summarizes the adoption of Blockchain, AI, and Quantum Computing across ten major global financial companies. The data highlights a competitive landscape where institutions are strategically investing in quantum to gain a potential long-term advantage in areas requiring massive computational power, such as portfolio optimization and complex risk analysis.

This isn’t just a finance story: it’s a signal for every industry. The real strategic advantage comes when these emerging technologies cross-pollinate and work together. That’s where the full impact shows up, stronger outcomes, faster innovation, and a much better return on every tech investment.

Company	Blockchain Use	AI Use	Quantum Team Start	Quantum Use Cases
<b>JPMorgan Chase</b>	Onyx platform, JPM Coin, IIN	450+ use cases, Gen AI for back-office	2018	Portfolio optimization, option pricing, risk management, quantum-safe cryptography.
<b>Barclays</b>	R3 consortium, DLT pilots for trade finance	Operational efficiency, fraud detection	Summer 2017	Quantum-resistant encryption, portfolio optimization, settlement batching.
<b>Wells Fargo</b>	Digital Cash for internal transfers	AI assistant, credit scoring, risk modeling	2019	Risk analysis, derivatives pricing, cybersecurity.
<b>BBVA</b>	Crypto custody, asset tokenization	Financial health app, fraud detection	2019	Portfolio optimization, credit scoring, currency exchange optimization.
<b>Citi</b>	Citi Token Services for cash management	Fraud detection, internal process automation	2020	Portfolio optimization, fraud detection, anomaly recognition.
<b>Goldman Sachs</b>	GS DAP (Digital Asset Platform), tokenization	Gen AI for developers, document analysis	2020	Derivative pricing (Quantum Monte Carlo), portfolio optimization.
<b>HSBC</b>	Orion tokenization platform, FX Everywhere	AI-powered fraud detection, customer service	2022	Algorithmic trading (first quantum-enabled bond trade trial), cybersecurity, NLP.
<b>Mastercard</b>	Provenance solution, Multi-Token Network (MTN)	Decision Intelligence for real-time fraud	2023	Loyalty and rewards optimization, feature selection for AI models.
<b>Visa</b>	B2B Connect, crypto-linked cards	Advanced Authorization for fraud prevention	2015 (PQC research)	Post-Quantum Cryptography (PQC) to secure payment networks.
<b>Standard Chartered</b>	SC Ventures, Zodia Custody, trade finance DLT	Credit risk, Anti-Money Laundering (AML)	Sep 2025 (Project Quanta)	Quantum-powered applications for financial markets and risk management.

## 14.5 Quantum Isn't "10 or 20 Years Away" Anymore

Many people, when I ask them if they are preparing for quantum, give me a glazed look, like "Are you kidding me? Our focus is on AI." They still think quantum is 10 to 20 years away. It doesn't bother me because I was the same way a couple of years back when someone said I should start looking into how fast quantum is moving. I can remember laughing to myself and saying, "Yeah, whatever. Quantum isn't on anyone's radar." That all changed when I started doing some research on my own and was caught off guard by how many countries and companies were focused on quantum.

In 2017, I came across blockchain and was so excited about this new disruptive technology that was going to change the world. I started taking courses and reading articles to get a stronger understanding of it. The biggest impact I got was when I started reaching out to people internationally on LinkedIn who had anything in their

title that said blockchain and asking if I could speak with them to understand what they were doing in the space. That approach gave me so many new perspectives that allowed me to share those stories from those conversations with others who were new to the blockchain space and how it was being adopted globally.

I decided to take that approach with quantum and was blown away by people I spoke with from countries like Japan, Chile, Switzerland, the UAE, India, Germany, Spain, France, and the UK, just to name a few. Hearing their stories and expertise was exciting and a little scary at the same time. Even though I consider myself new to the quantum space, it validated that quantum is moving faster than even they were expecting, just in the last five years, with the continuous breakthroughs that have been happening. It proved to me that people really need to start preparing earlier than ever to avoid getting left behind. As AI is still blinding many people, it is proving that quantum awareness is needed more than ever.

I saw this happen with blockchain back in 2017. Everyone was trying to understand what Bitcoin was, but many weren't being educated on the underlying technology and its broader benefits. As hype ramped up in 2019, leaders got confused about what to focus on, and everything got lumped into "blockchain is just Bitcoin," which pushed a lot of people away. To this day, organizations are still waking up to where blockchain can help, like supply chain and healthcare. Now with AI, it's happening even faster. Leaders and business teams see the value, but many are still struggling to put the foundations in place, such as education, training, and governance, so they don't get lost as AI speeds up.

With quantum, you're seeing the same pattern starting to form. There is a rush to get people "ready" without the basics in place first. The key is awareness and foundational education early. Starting sooner helps us avoid the mistakes of the past, when new tech was forced on people before they were prepared. Quantum will reward the organizations that prepare early, but for those of us in the quantum space, we also have a responsibility. Don't drag people straight down the rabbit hole of qubits, entanglement, and tunneling. Start with business value and relatable use cases. It will also take strong, coordinated collaboration across education, government, and the broader ecosystem, paired with smart investment, to ensure everyone is included and businesses can ease into quantum before they are forced into the deep end.

In the US, a handful of states have started to prepare early and are taking the initiative to prepare for quantum and its future. Here are a few:

### **Colorado (CO)**

- Colorado is accelerating early quantum adoption through statewide incentives, shared facility planning, and a strong commercialization pipeline anchored by Elevate Quantum.
- The state is also pushing quantum workforce readiness through coordinated education in through K-12 and training efforts.

### **Illinois (IL)**

- Illinois is building serious quantum infrastructure through the Illinois Quantum

& Microelectronics Park (IQMP), designed to take quantum from lab to market at scale.

- The Chicago Quantum Exchange and national-lab partnerships give IL a strong research-to-industry advantage.

### **New Mexico (NM)**

- New Mexico is leaning into its national-lab strength by aligning state strategy with federal partners and positioning itself as a quantum commercialization hub.
- The state’s “tech hub” momentum signals a direct push to convert R&D into workforce growth and startup formation.

### **Maryland (MD)**

- Maryland is branding itself as the “Capital of Quantum,” pairing major public-private investment with deep federal and university partnerships.
- With dedicated startup infrastructure and national-security-adjacent research assets, MD is pushing quantum from innovation to deployment.

### **South Carolina (SC)**

- South Carolina is formalizing its quantum ecosystem through the SC Quantum Association and targeted workforce development programs.
- The state is positioning quantum as an economic growth driver by aligning training, academia, and industry partnerships.

### **Massachusetts (MA)**

- Massachusetts is advancing quantum adoption by investing in shared computing infrastructure and supporting industry-academic partnerships.
- With an established deep-tech ecosystem, MA is turning quantum research strength into scalable innovation and commercialization.

### **Tennessee (TN)**

- Tennessee launched the Tennessee Quantum Initiative with \$20M in state funding, signaling a statewide commitment to quantum innovation and workforce development.
- EPB in Chattanooga is adding a \$4M NIST-backed investment to expand quantum workforce development tied to the region’s quantum network momentum.
- Oak Ridge National Laboratory (ORNL) is advancing hybrid quantum-classical computing, partnering with major compute players to integrate quantum into high-performance computing workflows.

Even as everyone races to adopt AI, these states are already preparing for the next wave because quantum and AI will be joined at the hip. That only happens when leaders get educated early, see what's coming, and move before the scramble starts.

## **Florida (FL)**

I live in Florida, and over the last six years we've seen tremendous growth as the state has evolved into a major innovation hub for emerging technology. A lot of that momentum started around 2018, when blockchain and Bitcoin were surging and Miami hosted the Bitcoin Conference, bringing international attention and accelerating interest from tech companies and investors. Now, with AI at the forefront, that momentum has only grown, positioning Florida as an increasingly attractive destination for both builders and capital. We've even seen companies with Silicon Valley roots expand into the state, including D-Wave Quantum, Anaplan, ServiceNow, and LeverX.

Now quantum is planting its flag in Florida, and it's catching many people in the region off guard because it hasn't been on their radar. This shift didn't happen overnight, but momentum has been building statewide, led in large part by Matt Cimaglia, Managing Partner at Quantum Coast Capital (QCC). QCC has been ahead of the curve, investing early to support quantum companies and startups. They realized early that they needed to build a stronger ecosystem before they could deploy money or investment. They've been clear that their mission is to work alongside researchers, institutions, and technologists to advance responsible quantum innovation across the U.S. and allied nations.

To support that vision, Matt has built real momentum across Florida by meeting with city leaders, state officials, and state universities to make the case for preparing for quantum, even while most organizations are still focused on AI. Over the past two years, he's helped bring these stakeholders together and get them aligned. Those conversations began to take root, and in October 2025 he brought the Quantum Beach conference to West Palm Beach, Florida. That convening helped bring major universities across the state together to agree to an MOU establishing the Florida Alliance for Quantum Technology, something Florida has never done before. It's a clear signal that quantum has to be a team sport, and that collaboration is exactly what it takes to build a real ecosystem, supporting companies and developing a quantum ready workforce.

Matt was instrumental in helping bring D-Wave into Palm Beach County and supporting the move of their headquarters to Boca Raton, FL, into the same building where IBM's first computer was built. That momentum helped set the stage for connecting D-Wave with Florida Atlantic University (FAU). In January 2026, FAU announced a \$20M agreement to purchase and install a D-Wave Advantage2 quantum system on its Boca Raton campus, with deployment expected later in 2026. That's a landmark moment, Florida's first on campus quantum computer at a university, and it puts the region on the global map in a real way. I believe Quantum Coast Capital's approach and the strategic relationship building Matt has driven across Florida is a strong blueprint for other regions looking to build a thriving quantum ecosystem.

## 14.6 Quantum Readiness: Move Early, Move Smart, Don't Get Left Behind

Quantum isn't a "flip the switch" technology. Adoption will be phased, and the companies that start early will have the advantage. The first step is understanding how quantum fits your business strategy, transformation roadmap, and real-world use cases. A quantum roadmap gives everyone clarity on what the goal is they should focus on, why it matters, and what investment it will take to do it right.

It's especially important not to chase the hype that always comes with new technology. Don't make the mistake of playing checkers for a quick win, treat this like chess. Be strategic and stay focused on long-term readiness. Here's a high-level quantum adoption roadmap that can help guide you on how to get started.

### **Phase 1 – Awareness & education (year 1)**

Focus on awareness and education in the first year. This helps the C-Suite and business leaders understand what's coming, reduces uncertainty, and lowers the anxiety that often shows up with new technology adoption.

### **Phase 2 – Pilots & talent development (years 1–2)**

In years one and two, prioritize small pilot projects and talent development. Start with a pilot tied to a real business case where quantum could create value, and make sure it aligns with your digital transformation strategy. In parallel, invest in building internal capability, developing or hiring the skills needed to form an internal quantum team that can carry adoption forward.

### **Phase 3 – Strategic integration into R&D and convergence (years 2–4)**

Between years two and four, focus on integrating quantum strategically into R&D especially where it converges with AI, cybersecurity, and other emerging technologies you've adopted. Some organizations may want to move faster but do it carefully. The AI landscape will continue shifting, and premature integration can force unplanned changes later.

### **Phase 4 – Operational adoption & scaling (years 4–6)**

Between years four and six, move toward operational adoption at scale. This phase is about confirming organizational readiness, rolling out capabilities responsibly, and building a plan to scale. Vendor and platform alignment becomes critical here—there won't be one quantum platform that "wins it all," so choosing the right partners

based on your specific business needs will matter.

These timelines aren't fixed they'll move based on how fast your organization is progressing. Just remember start early, and you control the pace.

## **14.7 Quantum is Here: What Are You Waiting For?**

If there's one thing I believe is essential, it's building international relationships in the quantum space. It shouldn't matter what country you're in. The more you connect with leaders, researchers, and businesses across regions, the more you'll accelerate your learning and build the relationships that will be pivotal as you move forward. I've said it before. Quantum is a team sport. If we collaborate, we can all win and be better prepared for the day quantum reshapes the world in ways we can't fully imagine yet. Don't wait for perfect clarity. Start building literacy, relationships, and a roadmap now.

# Chapter 15

## The Quantum in Saudi Arabia

### *From entanglement to engagement: a Saudi physicist's journey and perspective*

Khulud Almutairi

#### 15.1 The Nonlinear Path to Quantum

##### 15.1.1 The spark at King Saud University

My fascination with physics grew from an earlier and more fundamental instinct: the desire to solve problems. It was during my third undergraduate year at King Saud University in Riyadh, in a course on quantum physics, that I first encountered the idea that would come to define my intellectual life quantum entanglement, the phenomenon Einstein once dismissed as “spukhafte Fernwirkung”, or “spooky action at a distance.” The notion that two particles could remain correlated across any distance, sharing their physical states instantaneously, did not strike me as a mere theoretical curiosity. It seemed to me a profound reimagining of how information itself might travel through the universe. Where others saw a paradox to be reconciled, I saw an architecture to be understood.

After completing my bachelor's degree, I was hired as a teaching assistant at King Saud University, an appointment that reflected both my academic promise and the Kingdom's emerging investment in building domestic scientific capacity. My hunger for deeper engagement with entanglement led me to King Abdulaziz City for Science and Technology (KACST), where I sought collaboration with the research group of Prof. Zbigniew Ficek. The arrangement required working at a distance; cultural

---

This chapter is a featured contribution to the monograph *A Portrait of Quantum Technologies in Finance*, edited by Oswaldo Zapata, PhD, and published by The Quantum Finance Boardroom (2026). All rights reserved by the chapter's author.

norms at that time made close physical proximity between genders in laboratory settings difficult to navigate. I came to understand my own position in a curiously apt metaphor: I saw myself an entangled particle, sharing quantum information across a separation, connected to the research I needed despite physical distance. The collaboration yielded its own measurable outcome: a first-author paper in *Physical Review A* on generating two-photon entangled states in a driven two-atom system, a foundational milestone that established my credentials in the international quantum optics community. I would later observe, with satisfaction, how profoundly the Kingdom's research culture had transformed, KACST now welcomes and trains students of all genders equally, and I count myself among its researchers.

### 15.1.2 The Canadian formation shift

What began at King Saud University proved to be only an introduction. To probe quantum information more deeply, I traveled to the University of Calgary in Canada, where I pursued my master's degree under the supervision of Prof. Barry Sanders, a figure of considerable stature in quantum information theory. My research there addressed a scheme for generating nonlinear optical phase shifts by storing light pulses as collective atomic excitations in a two-component Bose-Einstein condensate, a mechanism with direct relevance to optical quantum computing. The work produced a second publication in *Physical Review A*, but its deeper significance lay in what it taught her about the gap between theoretical possibility and experimental realization.

That gap is what drew me next to Edmonton, at the University of Alberta, where I made my first experimental turn. There, I designed and fabricated my first nanofabricated chip from a silicon wafer, from initial design through cleanroom implementation, as part of a nano-optomechanical system for gas sensing using micro gas chromatography. Seeing quantum effects manifest at scales I could hold in my hands, rather than equations on a whiteboard, reshaped my understanding of the field's practical demands. During those years, I won first prize in the NanoNexus competition or a sensing project, a recognition that confirmed my instincts about the industrial relevance of quantum-enabled devices.

Then, in 2018, the geopolitical rupture between Canada and Saudi Arabia forced a Quantum jump, an abrupt end to my work in Edmonton. I returned to the Kingdom, carrying with me the technical formation of two Canadian institutions and an unfinished doctoral trajectory. The disruption was sharp, but not permanent, diplomatic relations have since normalized, and Saudi students once again pursue education in Canada.

### 15.1.3 Building the quantum community

After one year back in Saudi Arabia, I moved to King's College London to continue my experimental nanophotonics search. Within months, the COVID-19 pandemic shattered laboratory operations worldwide, and personal circumstances compelled my return to the Kingdom once more. What might have ended as a truncated academic itinerancy instead became the pivot point for a new vocation: not merely studying

quantum phenomena, but building the human infrastructure required a nation to comprehend and deploy them.

In late 2024, I co-founded the Saudi Quantum Technology Association, a nonprofit organization dedicated to expanding national quantum literacy. The association was supervised by the Ministry of Communications and Information Technology (MCIT), and deliver many activities, contributing in LEAP conference as exhibitor, partnered with IBM to deliver the Qiskit Challenge, and also I participated in World Quantum Day 2025, where I moderated panel to explore the quantum ecosystem in Saudi Arabia.

I subsequently left the Saudi Quantum Technology Association to focus on QSaudi Arabia, the Kingdom’s chapter of QWorld, the global nonprofit advancing open quantum education. The chapter was pre-lunched in 2025, having delivered the QBronze174 introductory workshop in November 2025; it expanded its official mandate in 2026 with additional programs aimed at raising quantum awareness across the general population. Looking back at this trajectory, from an undergraduate captivated by entanglement to a national advocate for quantum literacy, I see a symmetry that I could not have anticipated. The entangled particle, sharing information across distance, became the citizen-scientist building bridges of understanding across a nation. My personal trajectory mirrors, in miniature, the Kingdom’s larger transformation: from a distant observer of quantum revolution to an active participant shaping its regional terms.

## 15.2 The Quantum Kingdom: Strategy, Policy, and Architecture

### 15.2.1 Vision 2030 and the Five Pillars

Saudi Arabia’s quantum strategy did not emerge from a vacuum; it was anchored from the outset to Vision 2030, the Kingdom’s overarching framework for economic diversification and knowledge-economy development. The World Economic Forum’s assessment of the national quantum pilot notes that this anchoring was deliberate: “Saudi Arabia anchored the blueprint to Vision 2030 priorities from the outset,” positioning quantum technologies as enablers for “knowledge-intensive job creation in priority sectors such as energy, healthcare and finance”.

What distinguishes the Kingdom’s approach is the velocity of execution. Where many nations have produced quantum strategy papers that languish in bureaucratic review, Saudi Arabia moved from policy articulation to tangible on-soil deployment within approximately eighteen months. The strategy now spans five confirmed pillars that have matured from aspiration to active implementation: talent development, research capacity, hardware innovation, commercialization infrastructure, and public engagement. Each pillar has a designated institutional champion, and each has progressed from whiteboard to budget line.

## 15.2.2 The institutional landscape

The governance of Saudi Arabia’s quantum ecosystem is distributed across six primary institutional actors, each with a distinct mandate and a specific contribution to the national architecture. The following table maps these institutions, their core functions, and their interdependencies:

<b>Institution</b>	<b>Core Quantum Mandate</b>	<b>Key Initiative</b>	<b>Strategic Partners</b>
KACST	National quantum coordination; stakeholder mapping; ecosystem governance	Quantum Valley co-development (announced April 2025)	Aramco, SDAIA, IBM, Pasqal
SDAIA	Quantum-AI convergence; data infrastructure; professional upskilling	Quantum Valley partnership; quantum bootcamps (Jan 2026)	KACST, KFUPM
RDIA	National moonshot definition; R&D goal setting	Three moonshots declared Feb 2025: cognitive cities by 2040, fault-tolerant quantum by 2045, AGI by 2050	National Alliance for Quantum
C4IR Saudi Arabia	Global policy piloting; international framework adaptation	World’s first WEF Quantum Economy Blueprint pilot (2024–2026)	WEF, KACST
MCIT	Digital agenda integration; talent challenge execution	Quantum Challenge with IBM and Saudi Quantum Technology Association (Nov 2025)	IBM, SQTA
SAMA	Regulatory engagement; financial-sector quantum use-case development	Developing quantum computing use cases in banking risk analysis	Saudi banking sector

The table reveals a deliberate division of labor. KACST functions as the ecosystem’s central node, coordinating stakeholder alignment and co-developing the planned Quantum Valley with Aramco and the Saudi Data and AI Authority (SDAIA).

SDAIA, for its part, has embedded quantum computing within its broader data and AI mandate, most visibly through a five-day quantum computing bootcamp launched in January 2026 in collaboration with King Fahd University of Petroleum and Minerals (KFUPM), targeting professionals in computer science and related disciplines.

The Research, Development and Innovation Authority (RDIA) provided the timeline’s most dramatic punctuation mark in February 2025, when acting head announced quantum computing as one of three national “moonshot” goals, targeting a scalable, fault-tolerant quantum computer by 2045. This declaration placed Saudi Arabia in a small group of nations, alongside the United States, China, and the European Union, that have articulated explicit sovereign targets for fault-tolerant quantum capability, though the distance between a 200-qubit neutral-atom machine and the millions of physical qubits required for fault tolerance remains measured in generations of scientific progress, not merely years.

Centre for the Fourth Industrial Revolution (C4IR) Saudi Arabia occupies a unique position as the world’s first pilot implementation of the World Economic Forum’s Quantum Economy Blueprint. Over a structured multi-phase process, C4IR translated the global framework into context-specific national pathways, producing five strategic lessons that are now referenced in quantum policy discussions well beyond the Kingdom: the imperative of talent development; the necessity of hardware access; the bridge between innovation and commercialization; the cultivation of decision-maker awareness; and the foundational importance of governance and security. These lessons have become a de facto reference architecture for other nations contemplating national quantum strategy design.

The Saudi Central Bank (SAMA) represents the newest institutional entrant, but potentially the most consequential for one of the Kingdom’s priority sectors. SAMA has begun developing formal quantum computing use cases in banking risk analysis, marking the first active regulatory engagement with quantum applications in the Saudi financial sector. The techniques under exploration align with globally mature quantum finance applications, Monte Carlo simulation acceleration for Value at Risk (VaR) modeling, credit portfolio optimization, and systemic risk scenario analysis, domains where quantum advantage is considered achievable within the current decade using near-term hardware in hybrid classical-quantum architectures.

### 15.2.3 A global-first policy experiment

Saudi Arabia’s role as the reference implementation for national quantum strategy design carries implications beyond its borders. The WEF pilot revealed something that theoretical strategy documents rarely capture: the sequencing of interventions matters as much as their content. The Kingdom’s decision to pursue stakeholder mapping, talent challenges, and hardware deployment in parallel, rather than sequentially, created a governance architecture that is horizontally integrated but vertically shallow. Every pillar is active; no pillar is yet deep.

The 2045 moonshot, in this context, functions less as a prediction than as an organizing horizon. A fault-tolerant quantum computer by 2045 would require

sustained progress through multiple generations of qubit technology, error correction schemes, and cryogenic engineering. The declaration’s value lies not in its achievability but in its directional force: it signals to researchers, investors, and international partners that Saudi Arabia’s commitment is structural, not opportunistic.

## 15.3 Building the Quantum Stack: Hardware, Infrastructure, and Industrial Partnerships

### 15.3.1 The Aramco–Pasqal quantum computer

The physical manifestation of Saudi Arabia’s quantum ambition arrived on November 24, 2025, when Aramco and Pasqal announced the deployment of a 200-qubit neutral-atom quantum computer at Aramco’s Dhahran data center, the Middle East’s first quantum computer dedicated to industrial applications. The system, based on Pasqal’s programmable neutral-atom architecture arranged in two-dimensional arrays, is not a research curiosity housed in a university basement; it is an operational machine processing problems relevant to energy, materials, and industrial optimization.

Pasqal’s localization strategy in the Kingdom extends beyond hardware placement. Pasqal Arabia, the company’s dedicated Saudi entity, is led by General Manager Florent Verthuy, and the partnership architecture includes joint research programs, talent development pipelines, and a sustained commitment to building regional application expertise. This is not a vendor-customer relationship; it is a co-development arrangement in which knowledge transfer is an explicit contractual pillar.

Aramco’s position in this ecosystem is structurally unique among national oil companies globally. No oil company has deployed on-premises quantum hardware at this scale while simultaneously co-founding a national quantum valley and hosting an emulator (Dammam 7Q) among the region’s largest. Aramco is simultaneously an energy giant, a quantum infrastructure anchor, and a national technology deployer. This triple role creates efficiencies, the company can test quantum algorithms against its own operational datasets, but it also raises governance questions about access neutrality, particularly as the Quantum Valley takes shape and other institutions seek comparable compute resources.

### 15.3.2 KAUST quantum foundry

If the Aramco–Pasqal machine represents Saudi Arabia’s quantum present, the KAUST Quantum Foundry represents its bid for quantum sovereignty over the longer term. Launched on January 21, 2026, the foundry is the Kingdom’s first shared-access platform for quantum hardware fabrication, a facility designed to produce reproducible, commercial-grade quantum devices rather than one-off laboratory prototypes.

The Foundry’s architecture addresses what has historically been the most severe

bottleneck in quantum hardware development: the transition from heroic, irreproducible laboratory fabrication to standardized, scalable manufacturing. Its defining features include Process Design Kits (PDKs) for reproducible quantum device fabrication; shared access to KAUST’s Class 1000 cleanroom for industry and academic partners, dramatically reducing the capital barrier to hardware prototyping; and a multi-platform scope encompassing superconducting, photonic, and hybrid quantum architectures. Commercialization pathways run through KAUST’s National Transformation Institute, creating structured routes from foundry output to market-ready technology.

This multi-platform approach constitutes what might be termed a “leapfrog” strategy. Most nations that have built quantum fabrication dominance, the United States in superconducting qubits, Japan in photonics, the Netherlands in trapped-ion systems developed single-platform expertise before expanding. Saudi Arabia, by contrast, has chosen platform-agnostic fabrication from the outset. The strategy maximizes optionality: if superconducting qubits prove more amenable to error correction, the foundry can pivot; if photonic platforms achieve superior coherence times, the infrastructure is already in place. But breadth carries a risk. The talent required to design superconducting circuits differs substantially from that required for photonic integration, and a multi-platform mandate may dilute expertise of accumulation in any single architecture. The foundry’s success will depend on whether the Kingdom can generate sufficient specialized human capital across multiple quantum platforms simultaneously, a challenge that returns the narrative, inevitably, to the talent pillar.

### 15.3.3 Complementary infrastructure

The Aramco–Pasqal machine and the KAUST Foundry do not operate in isolation. They are flanked by a broader infrastructure layer that deepens the Kingdom’s quantum stack. The Aramco–IBM Innovation Hub, established under a memorandum of understanding signed in October 2022, explores applications of hybrid cloud, artificial intelligence, and quantum computing across materials science, supply chain, sustainability, security, and digitization. The hub is broader than a pure quantum play; IBM’s portfolio spans AI and classical high-performance computing, but its quantum component serves as the primary commercialization bridge for IBM’s quantum tools within the Kingdom.

At NEOM, the futuristic city rising on the Red Sea coast, a different kind of quantum infrastructure is taking shape. NEOM’s collaboration with Arqit, initiated in December 2021, embeds quantum-resistant cryptographic protocols into the city’s foundational architecture. This is anticipatory infrastructure: NEOM is being designed as quantum-secure before it is fully built, a rarity in global urban development where most cities retrofit security rather than architect it in from bedrock.

Finally, the Quantum Valley itself, announced by KACST in partnership with Aramco and SDAIA at World Quantum Day 2025 remains in planning and conceptual design as of early 2026. When realized, it is intended to serve as the Kingdom’s centralized hub for quantum hardware fabrication, testing, prototyping, and commercialization,

complementing the KAUST Foundry’s academic-industrial hybrid model with a national-scale facility. The physical ground has yet to be broken. This gap between declaration and construction is not unique to Saudi Arabia, quantum valleys in Europe and Asia have faced similar timelines, but it underscores the distinction between policy momentum and built reality.

Taken together, these three layers, the operational machine at Dhahran, the fabrication facility at KAUST, and the planned valley yet to be constructed, describe a nation that has chosen to build quantum infrastructure across the full stack simultaneously. It is an approach that prioritizes speed and coverage over depth in any single layer. Whether that gamble yields sovereign quantum capability or scattered underutilized assets will be determined not by the machines themselves, but by the people trained to operate them.

## 15.4 The Human Qubit: Academic Research and Talent Infrastructure

### 15.4.1 Research institutions

The Foundry etches qubits into silicon, and the Pasqal machine arranges neutral atoms in tweezers of light, but neither facility produces the minds that must design the algorithms, interpret the noise spectra, and engineer the error correction that turns fragile quantum states into reliable computation. Saudi Arabia’s academic quantum infrastructure has expanded significantly, but its depth remains uneven across the disciplines a sovereign ecosystem requires.

At King Abdullah University of Science and Technology (KAUST), the quantum research portfolio spans computing, sensing, photonics, and communications with a breadth that matches the Foundry’s multi-platform mandate. Collaborations with Saudi Aramco and SABIC extend quantum research into industrial contexts where problems are defined by pipeline corrosion rates and catalytic efficiencies rather than by abstract Hamiltonian elegance. The National Transformation Institute provides the commercialization bridge, addressing a persistent failure mode in academic research: the gap between publication and product.

King Fahd University of Petroleum and Minerals (KFUPM) approaches quantum technology from the Kingdom’s energy identity. The university hosts advanced quantum communication laboratories and offers specialized master’s-level programmes in quantum information and quantum computing, with curricula designed to feed graduates directly into energy-sector quantum applications. A dedicated research group in quantum photonics pursues integrated photonic circuits and optical neural networks technologies with immediate relevance to quantum sensing and secure communications. KFUPM’s partnership with the SDAIA Academy on professional upskilling creates a channel through which working engineers can acquire quantum literacy without leaving their posts.

The most strategically urgent research anchor, however, may be the quietest. King

Saud University’s Center of Excellence in Information Assurance (COEIA) conducts dedicated research in post-quantum cryptography, the discipline of designing encryption algorithms that resist attack by both classical and quantum computers. While KAUST and KFUPM build the tools of quantum opportunity, COEIA builds the defense against quantum threat. In a financial system where trust is the tradable commodity and Islamic finance operates across jurisdictions with varying regulatory maturity; cryptographic integrity is not merely a technical concern; it is a systemic one.

### 15.4.2 The talent pipeline

If the research institutions form the deep groundwater of national quantum capability, the training programmes are the pipes that must carry it to the surface, and those pipes are, by every assessment, too narrow for the demand.

The most visible recent interventions have been compact, intensive, and IBM-centered. The Ministry of Communications and Information Technology’s Quantum Challenge, held in November 2025, delivered five days of Qiskit-based quantum programming training in partnership with IBM and the Saudi Quantum Technology Association. The Tuwaiq Academy, operating as the Kingdom’s national technology talent upskilling platform, has embedded quantum computing within its broader curriculum through a long-standing IBM partnership providing hands-on training to students and professionals alike. In January 2026, the SDAIA Academy opened a five-day quantum computing bootcamp in collaboration with KFUPM, targeting professionals with five hours of structured training each day.

These programmes share a common design philosophy: short-duration, high-intensity, and vendor-aligned. They create an initial cohort of Saudis who can write a quantum circuit and interpret a Bloch sphere. What they do not yet create, what no Saudi programme has yet created at scale, is the quantum engineer who can design a cryogenic control stack, or the error-correction theorist who can adapt surface codes to a new hardware platform, or the quantum chemist who can translate molecular Hamiltonians into gate sequences. These specializations require years, not days, and mentors who have done the work themselves. The C4IR-WEF Blueprint identified talent development as the foremost strategic lesson from the Kingdom’s quantum pilot; the lesson is being learned, but its implementation remains measured in workshop certificates rather than graduate degrees.

My personal view is that I see this constraint from both sides: as a researcher who needed trained colleagues, and as an educator trying to produce them. The hardware arrived before the workforce to operate it. Every 200-qubit machine, every cleanroom, every planned valley amplifies the urgency of a problem that cannot be solved by procurement. It is a problem of human formation, and it moves on to the slow time of education.

## 15.5 The Quantum-Finance Nexus: Opportunity, Urgency, and the Saudi Advantage

### 15.5.1 Market landscape and sector dynamics

The Saudi quantum computing market stood at an estimated USD 33.7 million in 2025, projected to reach USD 279.3 million by 2032 at a compound annual growth rate of 35.3%, based on P&S Market Research. For an economy whose sovereign wealth fund deploys tens of billions annually, this might seem barely worth a regulator’s memo. But market sizing in quantum computing is deceptive. The figures measure direct spending on hardware, software, and services; they do not capture the value of portfolios optimized by quantum algorithms, transactions secured by quantum-resistant protocols, or systemic risks avoided by quantum-enhanced stress testing. In that broader frame, the quantum finance nexus is not a niche; it is a structural transformation of how financial value is created, protected, and transmitted.

<b>Metric</b>	<b>Value</b>	<b>Period / Notes</b>
2025 Market Baseline	USD 33.7 million	Foundation year
2032 Market Projection	USD 279.3 million	End of forecast horizon
CAGR (2026–2032)	35.3%	Among highest regional technology growth rates
Riyadh Market Share	45%	Policy, financial, and corporate HQ concentration
BFSI Sector Share	~30% (largest category)	Leading vertical by revenue
IT & Telecom CAGR	35.4%	Fastest-growing sector; QKD and QRNG interest
Trapped Ions CAGR	Highest among technology categories	Communications and cryptography applications
Eastern Province CAGR	35.5%	Fastest-growing province; Aramco quantum deployment

The table reveals a market that is both nascent and explosive. Riyadh’s 45% concentration reflects the capital’s aggregation of policy authority, financial institutions, and corporate headquarters, efficient for coordination but creating a geographic monoculture that concentrates risk alongside decision-making. The BFSI sector’s roughly 30% share makes it the largest revenue vertical, yet this figure underrepresents its true strategic weight. A SAMA regulatory mandate would cascade through all banks and trade-finance platforms, amplifying the BFSI figure’s systemic significance well beyond its nominal proportion.

On the technology trajectory, trapped ions represent the fastest-growing hardware

category, driven by applications in quantum communications and cryptography. The IT and telecommunications sector's 35.4% CAGR reflects accelerating investment in quantum key distribution (QKD) and quantum random number generation (QRNG) for secure communications infrastructure. These are defensive preparations for a threat environment that shifted dramatically in March 2026.

### **15.5.2 The Saudi Central Bank and regulatory engagement**

SAMA's entry into the quantum conversation marks a threshold. The regulator has begun developing formal quantum computing use cases in banking risk analysis, making it the first Saudi financial authority to treat quantum not as a distant abstract threat but as a present operational tool. The domains under exploration align with globally mature quantum finance applications: Monte Carlo simulation acceleration for Value at Risk (VaR) modeling, credit portfolio optimization, and systemic risk scenario analysis techniques where quantum-classical hybrid architectures are expected to demonstrate advantage within the current decade.

The timing is significant. SAMA's exploration of quantum opportunity coincided with a global event that transformed the threat side of the ledger. In March 2026, Google Quantum AI, in collaboration with the Ethereum Foundation and Stanford University published a whitepaper demonstrating that elliptic curve cryptography (ECC-256), the mathematical foundation protecting Bitcoin, Ethereum, and the majority of digital financial signatures worldwide, could be broken with fewer than 500,000 physical qubits. This represented roughly a twenty-fold reduction from prior estimates. A parallel paper from Oratomic and Caltech researchers suggested the requirement might fall as low as approximately 26,000 qubits under certain architectural assumptions. The two papers compressed the threat timeline from "decades away" to "conceivable within a scientific generation."

For SAMA, this creates a dual-urgency window rarely seen in regulatory history: the same institution must simultaneously explore quantum computing as a tool for competitive advantage in risk modeling and prepare the banking sector against quantum computing as an existential threat to cryptographic integrity.

## **15.6 Conclusion: Foundations Extraordinary, Building Just Begun**

### **15.6.1 The honest assessment**

Three years ago, the Kingdom's quantum ecosystem existed largely in presentation slides and conference statements. As of mid-2026, the inventory of built reality is substantive: a 200-qubit neutral-atom machine operating at industrial scale in Dhahran, a hardware foundry with shared cleanroom access at KAUST, the world's first completed national pilot of the WEF Quantum Economy Blueprint, a named quantum programme at SAB one of the Kingdom's largest banks, two active

nonprofit associations advancing public quantum literacy, and a regulatory authority actively exploring quantum use cases in financial risk analysis. This is not a nation experimenting with quantum as a peripheral curiosity. It is a nation embedding quantum into its industrial, financial, and educational infrastructure with deliberate velocity.

Yet honest accounting must also include what remains unbuilt. The talent pipeline produces awareness at a rate that outpaces expertise. The Quantum Valley, announced at World Quantum Day 2025, remains in planning and conceptual design. SAMA has not issued post-quantum cryptography guidance for financial institutions, even as the March 2026 research compressed the threat timeline to a horizon measured in single-digit years. And the 2045 moonshot, a scalable, fault-tolerant quantum computer, remains not a plan but a direction, separated from present capability by generations of scientific progress.

The distinction between policy momentum and built reality, a recurring motif in this account, has not been resolved. It has merely shifted the scale. Where once the gap lay between “having a strategy” and “having a machine,” it now lies between “having a machine” and “having the people who can extend it beyond its initial purpose.”

### 15.6.2 The decisive decade

Saudi Arabia’s quantum ecosystem is in its preparation phase. The hardware is being built. The policy frameworks are being tested. The financial sector is beginning to engage. But the critical question, whether quantum will be built in Saudi Arabia, by Saudi people, for Saudi systems, on Saudi terms, remains open. Its answer will be delivered by the accumulated choices made over the next decade in laboratories where Saudi graduate students debug their first cryogenic circuits, in classrooms where quantum mechanics is taught in Arabic as naturally as in English, in bank vaults where post-quantum protocols are migrated with methodical care, and in data centers where quantum-classical hybrid algorithms process the first Saudi-origin optimization problems at scales impossible for classical computation alone.

The foundations are extraordinary, and the building has barely begun. The answer to whether the Kingdom completes its quantum transformation will be written not in reports or roadmaps, but in the laboratories, classrooms, bank vaults, and data centers of the Kingdom over the decisive decade to come.

# Chapter 16

## Quantum Is a Public Good

*Translating quantum innovation into solutions for business and society*

Lionel Martellini

### 16.1 Professional Foundation

Looking back on my journey into finance, what stands out is that it was never driven by a primary interest in markets or investments. It was rather, from the beginning, an intellectual journey shaped by a fascination for mathematics and for the possibility of understanding complex systems through formal reasoning. In other words, I came to finance somewhat by detour. What attracted me was not the financial industry itself, but the opportunity it offered to apply mathematical tools to real-world problems involving uncertainty, risk, and decision-making.

After graduating from one of the leading French business schools, I felt a strong desire to pursue graduate studies in mathematics and statistics and eventually enrolled in a PhD in finance at UC Berkeley, where I was exposed to a rigorous vision of financial economics. I had the great privilege of having the late Mark Rubinstein as my PhD advisor. Mark was not only an exceptionally deep financial economist – perhaps best known for the Cox–Ross–Rubinstein discrete-time option pricing model, among many other outstanding contributions – but also a man of remarkable culture and intellectual breadth. To this day, I remember him once telling me that his favorite book in his entire (and extensive) library was a French book, *The Count of Monte Cristo* by the French author Alexandre Dumas. He ranked it even above Shakespeare’s plays, for which he had the greatest admiration.

---

This chapter is a featured contribution to the monograph *A Portrait of Quantum Technologies in Finance*, edited by Oswaldo Zapata, PhD, and published by The Quantum Finance Boardroom (2026). All rights reserved by the chapter’s author.

The early years of my academic career in the United States, at the University of Southern California at Los Angeles, followed by my return to EDHEC, were decisive in shaping my perspective. Very early on, I became convinced that finance, as a discipline, suffered from a structural tension: on the one hand, a strong theoretical framework grounded in elegant mathematical models; on the other, a real-world environment characterized by frictions, behavioral biases, and institutional constraints. Over the years, I became increasingly aware of the limitations of an ivory-tower approach to finance—one that emphasizes academic publication over engagement with practitioners, whose perspectives are critical to identifying and addressing real-world investment challenges.

This conviction led me to co-found the EDHEC-Risk Institute, with a very clear ambition: to contribute to reshaping best investment practices through the development of what I have often called “useful finance”—finance that is not only theoretically sound, but also practically relevant. Over time, this vision translated into a number of concrete initiatives that had a tangible impact on the asset management industry. One of the most significant contributions was our work on factor investing and smart beta strategies. At a time when capitalization-weighted indices were still largely dominant, we developed a research agenda aimed at providing efficient and robust alternatives based on well-documented risk premia.

This work ultimately led to the creation of Scientific Beta, a venture designed to bridge academic research and industrial implementation. Through this academic work and related industrial initiative, we were able to influence industry practices in a meaningful way—promoting a more disciplined, evidence-based approach to investment, and contributing to the broader adoption of smart factor investing as a more efficient way to harvest risk premia in equity markets.

At the core of my work has always been a focus on investment solutions and a paradigm that could be described as “precision finance”—a framework in which portfolios are no longer built through the mass distribution of standardized products, but through the design of tailored investment solutions aligned with specific investor objectives, constraints, and risk profiles. This shift implies moving away from a product-centric logic, often driven by marketing and scale considerations, toward a solution-oriented approach grounded in economic relevance and robustness. In such a framework, the focus is not on selling pre-packaged exposures, but on engineering portfolios that efficiently deliver desired outcomes—whether in terms of income generation, risk control, or long-term capital preservation—while explicitly accounting for uncertainty, estimation risk, and implementation constraints. This includes issues such as retirement investing, risk management, and long-term asset allocation.

My recent appointment as Research Director at the CFA Institute Research Foundation also provides me with a unique vantage point on the evolution of the global investment industry. It has reinforced my conviction that finance is at a turning point, where traditional approaches are increasingly challenged by technological transformations.

At the same time, this reflection on the evolving nature of finance has also led me to revisit a longstanding intellectual thread in my own journey which has been running in parallel to my academic career in finance: a deep and persistent fascination for

physics. For a long time, this remained a personal pursuit—something I explored outside of my professional life. In retrospect, it is clear that this dual interest in finance and physics was not entirely accidental. Both fields are concerned, in their own ways, with understanding complex systems under uncertainty through suitable mathematical structures.

## 16.2 The Quantum Pivot

If there was a “pivot” toward quantum physics and quantum technologies, it did not happen overnight. It was the result of a long intellectual maturation, and perhaps also of a form of fidelity to an earlier, more fundamental curiosity about the nature of the physical world that I have developed as a child.

A key turning point occurred during my sabbatical at Princeton in 2011-2012. Immersed in an environment immensely rich in the history of theoretical physics, with Albert Einstein, John Von Neumann or Robert Oppenheimer among the leading figures, I found myself reconnecting with a long-standing intellectual attraction. Attending two doctoral-level courses one in general relativity, and the other in cosmology by Paul Steinhardt remains, to this day, the most intellectually exhilarating experience of my career.

This experience led me to pursue a PhD in astrophysics, and later to participate in the international LIGO collaboration that achieved the first detection of gravitational waves. Being part of this effort, even at a modest level, was an extraordinary privilege—an experience that profoundly reshaped my relationship to science.

But the real turning point came later, during a subsequent sabbatical at MIT in 2022-2023, where I had the privilege of interacting with, and learning, from exceptional scholars such as Seth Lloyd at MIT or Jacob Barandes at Harvard, and where I became deeply immersed in the foundations of quantum mechanics.

My work focused in particular on the problem of time in quantum systems, and more specifically on what is known as the time-of-arrival problem—a fundamental question that, surprisingly, remains unresolved within the standard formalism of quantum mechanics. This led to a series of publications in leading physics journals, including work on time-of-arrival distributions for both continuous and discrete quantum systems, and the development of a general framework to address these questions without departing from the standard quantum formalism.

What fascinated me in this line of research was not only the technical challenge, but the conceptual depth. The fact that such a basic question—“when does a particle arrive?”—does not have a universally accepted answer within quantum mechanics reveals something profound about the limits of our current understanding of the physical world. At the same time, these questions are not purely abstract. They have direct implications for quantum information processing, including quantum computation and quantum communication, where the notion of time plays a critical role in algorithmic execution and information transfer. Different tasks such as circuit execution times, error detection latencies, or communication and synchronization

protocols, may naturally call for different notions of arrival time.

While my work has been primarily foundational so far, there are several concrete domains in which operational notions of time-of-arrival distributions may play a substantive role, for instance in applications to information propagation in locally interacting quantum systems. The so-called Lieb-Robinson bounds<sup>1</sup> rigorously constrain the maximal velocity of information spreading but they provide worst-case envelopes rather than operational access times. Arrival time distributions associated with specific detection protocols could instead characterize when excitations or logical information become experimentally accessible at a remote site. In this sense, these concepts complement velocity bounds by introducing a probabilistic notion of information arrival time, which may help enhance available estimates based on circuit depth time performance in many body systems.

It was at this point that the connection between my work in physics and my background in finance became evident. On the one hand, I was working on foundational questions in quantum theory and applications in quantum technologies. On the other, I had spent decades studying complex optimization and decision problems in finance and their applications in financial engineering. The realization was that we are entering a new phase—the second quantum revolution—where these foundational insights are beginning to translate into technological capabilities. Quantum computing is not just about faster computation. It is about a fundamentally different way of representing and manipulating information. And this has immediate and profound implications for fields like finance. In that sense, the “spark” was not a single moment, but the gradual realization that two intellectual trajectories I had pursued independently for many years were, in fact, converging toward a common frontier.

From this journey that has allowed a somewhat dreamy teenager to reconcile a passion for physics with a professional career in finance, I take away one central lesson: the importance of remaining, in some sense, faithful to one’s childhood dreams. This fidelity strikes me as a valuable guide in the seemingly modest yet extraordinary adventure that is a human life. Quantum physics teaches us that, at the heart of the fabric of the universe, lies a vast wave function—a generative source of infinitely many possibilities. In much the same way, human life can be seen as a continuous journey of self-transcendence, growth, realization, and the unfolding of these latent potentials. Just as it takes time, vision, talent, and sustained effort for an artist to transform what exists only in potential—clay for a sculptor, a blank page for a writer, or a blank canvas for a painter—into reality, each of us must build our own destiny step by step, assuming full responsibility for being the sole architect of this long and patient process through which our life ultimately takes shape. All of this is beautifully captured in Nietzsche’s famous injunction: “Become who you are. Do what only you can do.”

---

<sup>1</sup>Lieb, E.H.; Robinson, D.W. The finite group velocity of quantum spin systems. *Commun. Math. Phys.* 1972, 28, 251-257.

## 16.3 Identifying the Void

The creation of the EDHEC Quantum Institute emerged from a very clear observation: despite the rapid acceleration of quantum technologies, there is a profound gap between technological development and economic understanding.

On the one hand, we are witnessing massive investments, strong governmental support, and rapid progress in quantum hardware and algorithms. On the other hand, there is a striking lack of clarity among business leaders, policymakers, and even academics about what quantum mechanics tells us about the world, the physical world, and how quantum technologies are going to impact our economies.

“Shut up and calculate” is not a viable option. This famous expression, due to David Mermin, reflects a crude pragmatic injunction according to which physicists should simply apply the unambiguous mathematical rules of quantum mechanics to obtain correct predictions, instead of trying to interpret the meaning of the theory. While this approach has proven remarkably effective in advancing the field, it deliberately sets aside deeper questions about what quantum mechanics actually tells us about the nature of reality. And for the development of quantum technologies, such an attitude is no longer sufficient. If these technologies are to fulfill their promise, they must be understood—not mystified. They are not based on magic, but on our ability to harness the logic of the physical world at its most fundamental level. What may appear counterintuitive is not irrational; it reflects a deeper structure of reality that we are only beginning to grasp.

In that sense, we are at a stage comparable to that of our ancestors when they first learned to master fire. We are aware of its transformative potential, we can generate sparks, and we are beginning to experiment with its applications. But we are still learning how to control it, stabilize it, and deploy it safely and effectively at scale. And we need to be able to communicate and explain to non-physicists what these new technologies are all about. The challenge ahead is therefore not only technological, but also conceptual. It requires moving beyond fascination and toward understanding—building a coherent framework that allows scientists, engineers, and decision-makers alike to engage with quantum technologies in a rigorous and informed manner. Only then can we transition from isolated demonstrations to robust, real-world impact.

In other words, while physics and engineering departments are doing an amazing job at pushing the boundaries of our knowledge in quantum science and technologies, there is a lack of institutions capable of translating quantum innovation into business-relevant insights. There is also a need for independent analysis to distinguish between what is feasible today, what is plausible in the medium term, and what remains purely aspirational.

In finance, this gap is particularly problematic. The industry is both a potential beneficiary of quantum technologies and a potential victim, especially in the context of cybersecurity and post-quantum cryptography. It is somewhat difficult for decision makers in the financial industry to understand these threats and opportunities, and how best to prepare for these technological advances. There is a certain tendency,

which we call “quantum washing”, to attribute near-term economic or operational benefits to quantum technologies either through implicit assumptions that are not currently satisfied or by focusing on problems chosen primarily for their amenability to quantum advantage rather than for their genuine economic relevance. Claims of universal speedups or systematic outperformance that often mix long-term theoretical possibilities with present-day capabilities makes it particularly challenging for non-experts to assess the reality of what a genuine quantum advantage can mean.

Beyond these technical considerations, there is a deeper issue: the lack of quantum literacy among decision-makers. This is where I believe business schools have a critical role to play. Training engineers and physicists is, of course, essential. But it is equally important to train business leaders who are capable of understanding the strategic implications of quantum technologies—who can ask the right questions, make informed decisions, and avoid both excessive skepticism and blind enthusiasm.

This represents a fascinating and non-trivial challenge. How do you teach quantum concepts—often perceived as abstract, counterintuitive, even “mysterious”—to students with a business background? How do you strike the right balance between conceptual rigor and accessibility?

This challenge is at the heart of the EDHEC Quantum Institute, which aims at bridging quantum science and real-world impact to unlock transformative solutions for business and society by leveraging EDHEC existing areas of expertise and existing strategic partnerships. We have launched research programs on foundations of quantum mechanics, quantum technologies in finance and insurance, and quantum technologies in information and communication, in partnership with EURECOM, an internationally renowned academic center for research and education in digital sciences. In a second step, we envision expanding the scope to other business domains, such as quantum technologies for AI, for the new space industry or for the biotech industry, which are already strongly impacted by the 2nd quantum revolution.

The Institute was designed as a bridge structure, with four core pillars: research, education, innovation, and outreach. Its ambition is obviously not to compete with physics departments, but to operate at the interface between science and application. The education component is central. The talent gap in quantum technologies is large and new educational programs are needed to train experts that can combine knowledge in quantum physics and business applications. In a nutshell, the education pillar of the EDHEC Quantum Institute is designed to equip future business leaders, engineers, and policymakers with the foundational knowledge and strategic insight required to navigate the rapidly evolving landscape of quantum technologies. By integrating quantum science with management education, the initiative addresses the growing need for interdisciplinary talent capable of bridging deep tech and real-world applications. For this, we develop programs that demystify quantum technologies, provide a clear conceptual framework, and connect these concepts to real-world use cases.

## 16.4 The Future Landscape

Quantum computing is widely expected to bring major improvements to financial problem-solving. In the area of portfolio optimization, quantum algorithms such as the Quantum Approximate Optimization Algorithm (QAOA) can solve complex construction problems involving high-dimensional constraints more efficiently than classical algorithms. In risk analysis and financial simulations, quantum-enhanced Monte Carlo methods offer the potential for faster and more accurate results, especially in pricing derivatives and calculating Value at Risk. Additionally, quantum machine learning (QML) is being explored as a means to detect patterns in large financial datasets, enabling better fraud detection, credit risk assessment, and trading strategy development. Quantum communication also has significant implications for finance. Quantum Key Distribution (QKD) can enable ultra-secure communication, which is critical for transmitting sensitive financial data. As quantum computers evolve, they will also pose a threat to existing encryption systems. Preparing for this shift through post-quantum cryptography will be essential to ensure data privacy and security in future financial infrastructures.

In the near term, the use of quantum computing in finance will be focused on hybrid approaches, combining classical and quantum methods. The objective for the time being is not to achieve a clear quantum advantage, but to build capabilities, identify use cases, and develop an operational understanding of the technology. In the medium term, as hardware improves, we can expect more significant advances in areas such as optimization, simulation, and machine learning. At the same time, the emergence of quantum threats to existing cryptographic systems will force a transition toward post-quantum infrastructures.

EDHEC has a long-standing reputation for academic leadership and research in data-driven finance, which offers a solid foundation for engaging in quantum finance. The EDHEC Quantum Institute is developing research on portfolio construction, asset pricing, and risk management, areas where quantum technologies are expected to have the most impact.

Establishing a dedicated research cluster focused on quantum finance is also essential to establish thought leadership. As the quantum finance field is still in its early stages, we ambition to influence the direction of research and standard-setting within the industry. As a thought leader, we also ambition to play a key role in policy development and industry regulation. The ethical, legal, and social implications of quantum technologies are still being defined. Through informed research and engagement, the EDHEC Quantum Institute can contribute to establishing responsible frameworks for their use in finance.

But beyond these technical developments, I believe the most important challenge is conceptual. Quantum mechanics is often presented as mysterious, counterintuitive, even incomprehensible. This narrative, while appealing, is ultimately counterproductive. It creates barriers to understanding and opens the door to confusion—and sometimes even to forms of “quantum mysticism” that have no scientific basis.

As already mentioned, one of our key ambitions at the EDHEC Quantum Institute is

therefore to clarify the concepts, and help economic decision makers understand that quantum mechanics, while profoundly different from classical physics, is not magical. It is based on a coherent logical structure, and it can be understood—at least to a meaningful extent—by non-specialists. This effort of clarification is essential, not only for educational purposes, but also for the development of the field itself. A discipline cannot mature if its foundational concepts remain obscure to those who are expected to use it. In this context, we intend to explore the possibility of developing certification frameworks, which would allow professionals to demonstrate a certain level of understanding of quantum technologies and their business implications. Such certifications could play a role similar to what the CFA designation has played in finance—providing a common language, a shared set of standards, and a signal of competence.

The EDHEC Quantum Institute ambitions to develop onsite and online executive programs aimed at professionals seeking to understand the strategic implications of quantum disruption. In parallel, we are developing an Innovation Hub with a dedicated quantum track within the TechForward incubator created by EDHEC and EURECOM, which is meant to support early-stage ventures that apply quantum technologies to business, finance, communication, security, health, logistics, and other industries. By connecting startups to cutting-edge research from EDHEC and EURECOM, the incubator aims to foster ventures grounded in both scientific rigor and commercial relevance. A special track is devoted to student-led startups emerging from the Institute’s master and MBA programs.

In a way, the EDHEC Quantum Institute is an attempt to address a fundamental inefficiency: the fact that technological revolutions often outpace our collective ability to understand and integrate them. More broadly, our ambition is to contribute to the structuring of the emerging field of quantum finance. This involves not only developing new methods and applications, but also defining standards, best practices, and evaluation frameworks. It also means engaging with industry, policymakers, and the broader public to ensure that the development of quantum technologies is aligned with societal needs.

Again, my core belief is that quantum is a public good, too important to be reserved for a small circle of initiated experts. Its implications reach far beyond physics, shaping industry, policy, security, and education. In addition to equipping us to engage with emerging quantum technologies, engagement with quantum principles offers a further intellectual dividend, one that is particularly useful at a time when public discourse is increasingly shaped by simplistic, one-sided views that distort and caricature what is inherently subtle and layered. Quantum theory is not only a powerful scientific and technological framework; it is also a welcome celebration of complexity, an invitation to contemplate reality and to understand that only a multiplicity of perspectives allows us to grasp its full richness.

# Chapter 17

## My Finance Journey Into Quantum

### *A portfolio manager's perspective*

Maylix Brianto

#### 17.1 When Curiosity Changes the Direction of a Career

Some moments quietly change the direction of a career. Not through a formal decision, not through a carefully planned strategy, but through curiosity — a spark that awakens attention and invites exploration.

For me, one of those moments happened during a networking lunch in Geneva. At first, it seemed like any other gathering: professionals from finance, technology, and academia sharing ideas about innovation and the future. The conversation shifted seamlessly from global markets to artificial intelligence, from economic trends to scientific breakthroughs, and then, almost naturally, someone introduced the topic of quantum technologies.

I remember the warmth of the room, the hum of conversation, and how ideas seemed to ripple from person to person. One colleague pulled out a tablet to show a simulation of a quantum algorithm. Another spoke of the potential for cryptography to secure an entire digital infrastructure against threats we hadn't yet imagined. At first, it all felt distant, the kind of discussion that belongs to physicists, laboratories, and equations far removed from daily portfolio management. Yet, as the conversation unfolded, I felt a spark of recognition — not in the mathematics itself, but in the philosophy underlying it.

---

This chapter is a featured contribution to the monograph *A Portrait of Quantum Technologies in Finance*, edited by Oswaldo Zapata, PhD, and published by The Quantum Finance Boardroom (2026). All rights reserved by the chapter's author.

Quantum theory suggests that reality is rarely linear or deterministic. Multiple possibilities can exist simultaneously, coexisting until observation collapses them into a single outcome. Probabilities intertwine, complexity becomes the natural state of things, and uncertainty is not an obstacle but a dimension to navigate.

In that moment, it struck me that portfolio managers have always operated in a quantum-like mindset. Every day, we confront uncertainty. Markets are not predictable systems; they are living ecosystems influenced by policy, technology, geopolitics, and human behavior. When we allocate capital, we are never betting on a single future. We are managing probabilities, constructing a framework in which multiple potential outcomes coexist and are balanced according to risk, opportunity, and long-term vision.

That lunch planted a seed of curiosity that would quietly grow into a passion for exploring how emerging technologies — especially quantum computing — could shape financial modeling, portfolio construction, and risk management. More importantly, it reminded me that true innovation in finance rarely emerges from staying within familiar boundaries. It arises at the intersection of disciplines: between science and markets, between technology and human judgment, between curiosity and responsibility.

This chapter is not a technical treatise on quantum physics. I am not a physicist, nor do I claim to be one. Rather, it is the reflection of a portfolio manager who has learned to think across dimensions, embrace complexity, and imagine possibilities beyond conventional frameworks. It is a story of curiosity, of discovery sparked by conversation, and of how a simple lunch in Geneva opened the door to a new way of thinking about investing, innovation, and the future of financial decision-making.

Perhaps, as you read these pages, you will discover that the world of quantum thinking is not as distant from finance as it may first appear.

## 17.2 My First Lessons in Finance

To understand why that moment mattered to me, one has to go much further back — to Venezuela. My father worked at the Venezuelan Central Bank. When I was a child, I often accompanied him through its corridors. I remember the marble floors, the quiet conversations, the seriousness that seemed to surround the building.

At the time, I did not understand macroeconomic policy or monetary stability. But I sensed that decisions made inside those rooms carried weight far beyond the walls of the institution. They influenced people's lives: the price of bread, the stability of families, the future of a country.

Those early impressions planted a seed of respect for economics and finance. I learned, even before formal education, that numbers are never just numbers. They tell stories about societies, hopes, and challenges. Growing up in Venezuela also meant learning resilience. Economic uncertainty was not theoretical; it was lived experience. Observing that environment taught me adaptability, curiosity, and a deep awareness that financial systems are human systems. Perhaps without realizing

it, those early experiences prepared me for a career defined by uncertainty — the very environment where portfolio managers operate every day.

### 17.3 The Beginning of a Global Journey

Driven by curiosity and a desire to understand financial systems more deeply, I pursued studies in banking and finance. During those years, I had the opportunity to represent my university at national and international conferences — experiences that exposed me to a wider world of ideas and perspectives.

After graduation, I returned to the Venezuelan Central Bank while completing a master's degree in finance. Research, conferences, and publications became part of my academic life. What I discovered during that time was something that continues to guide my career: finance is not purely technical. It is also ethical. Financial systems influence societies. They shape opportunities, stability, and prosperity. Understanding them requires not only analytical skill but also responsibility.

My journey eventually took me beyond Venezuela. I moved to Grenoble, France, to pursue an MBA in International Management. Living and studying in a new language and culture required humility and perseverance. It was also a moment of intellectual expansion — learning how different countries approach business, governance, and strategy. Years later, returning as a guest speaker and Professor in Grenoble Ecole de Management felt like a full circle. It reminded me that education is not only about acquiring knowledge; it is about passing curiosity forward.

### 17.4 From Washington to Geneva

My career then took another international step when I joined the Inter-American Development Bank in Washington, D.C. There, finance met development. Numbers connected directly with human impact. Projects were not only financial models; they were roads, infrastructure, education systems, and opportunities for communities.

But life had another destination waiting: Geneva. Geneva is a fascinating place for finance. It is global yet intimate. It brings together wealth management, international institutions, technology, diplomacy, and academia.

It was here that I entered the world of private banking. At JP Morgan Private Banking, I worked within the hedge fund trading support team. Those years introduced me deeply to the universe of alternative investments — hedge funds, complex strategies, operational processes, and risk management frameworks. The intellectual challenge fascinated me. It also inspired me to pursue the CAIA certification, a program that deepened my expertise in alternative investments and connected me to a global network of professionals. Over time, my involvement with the CAIA Association grew into a leadership role within the Geneva Chapter, where I have had the privilege of organizing events, moderating discussions, and contributing to the development of the alternative investment community. Those experiences reinforced

something I believe strongly: knowledge grows when it is shared.

## 17.5 Building Bridges in Wealth Management

My professional path continued within wealth management institutions in Geneva. At Société Générale Private Banking, I worked with clients across Latin America and Iberia, strengthening relationships with independent asset managers and entrepreneurs. This role allowed me to connect my cultural roots with my professional expertise. It also reminded me that finance is deeply relational. Behind every portfolio is a family. Behind every investment strategy is a story.

Later, at Banque SYZ, I joined a team focused on developing relationships with ultra-high-net-worth clients in Latin America. The experience strengthened my understanding of international capital flows, entrepreneurial wealth, and the strategic needs of global families. Eventually, I transitioned into the role that felt most aligned with my interests: multi-asset portfolio management. Today, working within an external asset management and multi-family office environment in Geneva, my role combines strategy, analysis, and client dialogue. Our mission is not only to generate returns but to build portfolios that reflect long-term values, aspirations, and legacies. It is also within this role that my curiosity about quantum technologies found a natural home.

## 17.6 Discovering the Quantum Frontier

Quantum technologies represent one of the most fascinating scientific frontiers of our time. Yet what captured my attention was not the physics itself, but the mindset it represents. Classical computing processes information in binary form — zeros and ones. Quantum systems introduce the possibility of superposition and entanglement, allowing information to exist in multiple states simultaneously.

For physicists, this opens extraordinary computational possibilities. For investors, it opens something equally interesting: a new way of thinking about complexity. Markets are complex systems influenced by countless variables — economic policy, technological innovation, geopolitical shifts, behavioral psychology. Traditional models often simplify these interactions in order to make them computationally manageable. Quantum computing promises the ability to explore far more complex scenarios. In fields such as portfolio optimization, risk modeling, and derivatives pricing, quantum algorithms may one day help financial institutions analyze enormous datasets and identify patterns that classical computing struggles to capture. While many of these applications remain in development, the direction is clear: quantum computing is moving from theoretical exploration toward real-world experimentation.

## 17.7 Exploring the Quantum Investment Frontier

Curiosity alone was not enough. That spark from Geneva had to become action. Over time, I began exploring how frontier technologies could translate into real investment opportunities. Quantum finance, I realized, was not about mastering qubits or coding algorithms. It was about thinking differently, approaching uncertainty with probabilistic rigor, and integrating complex, emerging opportunities into portfolios that balance stability with innovation.

I began reading papers, attending conferences, and meeting the founders and teams pioneering quantum computing, cryptography, and advanced algorithms. Every discussion felt like stepping into a new continent, where the landscape was largely uncharted and full of potential. I allocated capital to early-stage private equity funds focused on quantum technologies — not as a speculative gamble, but as a deliberate exploration aligned with long-term vision. Every report, every investment thesis, every scenario analysis became a lesson in marrying curiosity with responsibility.

Quantum finance reshaped how I think about risk. It reminded me that uncertainty is not a limitation; it is a dimension to navigate. Classical assumptions often simplify a complex reality; quantum thinking invites us to embrace complexity, explore multiple pathways, and design strategies that anticipate alternative outcomes. For a portfolio manager, this is more than an intellectual exercise — it is a practical philosophy.

## 17.8 Navigating Quantum Opportunities Across Continents

Quantum adoption does not look the same everywhere. Europe and Switzerland, for example, benefit from strong innovation ecosystems, world-class universities, collaborative tech hubs, and regulatory frameworks that support frontier experimentation. Investors can leverage venture capital opportunities, specialized research networks, and high-quality data to integrate quantum technologies responsibly into portfolios.

In Latin America, the landscape is different. Emerging markets carry higher risks but also unique opportunities. Governments in Brazil, Chile, and Mexico have started incentivizing research and innovation, creating fertile ground for transformative applications in energy, logistics, and finance. Here, investors must balance ambition with prudence — combining rigorous due diligence with local insights and a deep understanding of market dynamics.

Understanding regional nuances is critical. Quantum finance is not abstract; it is a framework for translating emerging technology into actionable, responsible strategies that reflect local realities and global foresight. The lessons from these diverse regions reinforce a key principle I apply across all portfolios: knowledge without context is incomplete.

Across continents, I see myself as a bridge connecting ideas, capital, and talent. My roots in Latin America, combined with years of experience in Geneva and Europe,

allow me to understand both the nuances of emerging markets and the expectations of sophisticated investors. I can translate local insights into actionable strategies for UHNW individuals, family offices, and institutions, while bringing global perspectives to frontier opportunities. This dual lens — grounded in cultural understanding and professional expertise — enables me to identify opportunities that others might overlook, assess risk with precision, and foster collaborations that span regions and sectors. In practice, this means connecting quantum innovators in Brazil or Mexico with European venture networks, introducing institutional investors to early-stage frontier technologies, and ensuring that every strategy is both locally informed and globally relevant. In a world where knowledge alone is not enough, being a bridge — between continents, between finance and technology, and between human ambition and responsibility — becomes a decisive advantage.

## 17.9 Women, Next-Gen, and the Quantum Ecosystem

Frontier finance thrives on diverse perspectives, and quantum technologies are no exception. Over the years, I have seen how curiosity, mentorship, and representation can shape not only careers but entire ecosystems. As a woman navigating alternative investments and the emerging world of quantum finance, I have felt both the challenge and the opportunity of stepping into spaces historically dominated by others. Each time I have shared the stage — whether moderating panels, speaking at conferences, or teaching students — I have been reminded of the power of visibility and the ripple effect it creates.

I have had the privilege of speaking at events celebrating women in finance, including panels with the 100 Women in Finance network, where conversations go beyond titles to focus on impact, innovation, and leadership. In the world of frontier technologies, I have also contributed to discussions on Women in Web3, highlighting how digital innovation and decentralized finance intersect with human-centered values. These experiences have reinforced a key insight: encouraging women to engage with frontier technologies is not only about equity — it is about enriching the ecosystem with creativity, collaboration, and fresh ways of seeing possibilities.

Equally important is the next generation. Students, early-career analysts, and emerging investors bring curiosity, openness, and fearless creativity. They ask the questions I once asked myself: “How can we navigate uncertainty responsibly?” “What tools allow us to translate complex opportunities into actionable strategies?” As a professor in international economics and finance at business schools, I have the privilege of guiding these conversations in classrooms and workshops across borders, encouraging students to think probabilistically, explore new technologies, and connect theoretical frameworks with practical, responsible investing. Seeing their energy, curiosity, and ambition is a daily reminder that the future of finance will be defined not only by technology but by the people who dare to imagine it differently.

Quantum finance itself is a platform for collaboration, education, and empowerment. It is a space where ideas meet action, where research partnerships intersect with

practical investment strategies, and where the human dimension — values, judgment, and foresight — is as critical as the technology itself. For women and the next generation, participation in this ecosystem is both an opportunity and a responsibility: to contribute insights, to mentor, to challenge assumptions, and to help shape a financial world that is inclusive, ethical, and forward-looking.

Whether through mentorship programs, educational workshops, or global panels, I have seen firsthand how sharing knowledge amplifies impact. Quantum technologies are complex, but their potential is human. By bridging finance, technology, and education, we create a community that does not merely follow trends but actively shapes them, ensuring that innovation is accompanied by wisdom, responsibility, and inclusivity.

The ecosystem is alive, global, and vibrant. Europe’s innovation hubs, Latin America’s emerging markets, and Asia’s growing tech clusters each present unique opportunities. In every region, women, students, and emerging investors are part of the equation — essential contributors to a frontier that thrives on diverse perspectives and bold ideas. My presence in conferences, panels, classrooms, and networks is not only about sharing what I know but about inviting others to step forward, to explore, and to build the future together.

## 17.10 Investor Guide: Engaging with Quantum Finance with Curiosity, Rigor, and Vision

Investing in quantum technologies is not about following the latest trend or chasing hype. It is, at its core, about curiosity, discernment, and a long-term vision — qualities that distinguish thoughtful investors from those simply reacting to headlines. For institutional investors, ultra-high-net-worth individuals, family offices, or tech-focused funds, the frontier of quantum finance requires both rigor and imagination. One cannot separate opportunity from responsibility, nor innovation from strategy.

The first step is understanding the landscape. Quantum finance today is multi-dimensional. It spans private equity funds targeting early-stage quantum computing startups, venture investments in quantum cryptography, collaborations with research institutions developing applied quantum algorithms, and partnerships exploring hybrid classical-quantum applications. Each path carries unique risk profiles, timelines, and potential impact. Knowledge, not speed, is your most important asset. Rushing into investments without comprehension risks missing both nuance and opportunity.

Integration is equally critical. Quantum investments are rarely standalone bets; they perform best as part of a broader, thoughtfully diversified portfolio. Combining frontier technologies with traditional allocations and alternative strategies allows investors to explore innovation without compromising stability or long-term objectives. For family offices, this integration also offers the chance to align financial ambition with values: portfolios that are intellectually exciting can also reflect environmental stewardship, social responsibility, and a commitment to the broader impact of capital deployment.

Due diligence is essential and non-negotiable. This goes beyond reading pitch decks or financial statements. It means meeting the founders, understanding the science underpinning the technology, evaluating scalability, and assessing potential market adoption. It also means identifying advisors who speak both the language of finance and the language of qubits — professionals capable of bridging the gap between highly technical innovation and portfolio strategy. These advisors help translate complexity into actionable insights, ensuring that investments are not only forward-looking but also responsibly evaluated.

Patience, perhaps, is the most underrated yet critical element. Quantum finance is a marathon, not a sprint. The most transformative opportunities often emerge over years, not quarters. Investors who cultivate patience — who balance curiosity with discipline — are best positioned to capture the promise of this emerging frontier. In practice, this patience manifests in ongoing monitoring, iterative learning, and an openness to recalibrate strategies as technology, markets, and regulatory landscapes evolve.

Ultimately, investing in quantum technologies is about more than returns. It is about participating in a frontier that reshapes how we think, how markets operate, and how possibilities are measured. It requires vision, courage, and humility — the recognition that even as we explore uncharted territory, our decisions carry real-world consequences. The investor who approaches quantum finance with curiosity, rigor, and long-term vision does more than allocate capital: they become an active participant in shaping the future of innovation, finance, and society itself.

## **17.11 Quantum-Enhanced Alternative Investments: Bridging Frontier Technology and Portfolio Strategy**

Investing in alternative assets has always been about balancing curiosity with discipline. Hedge funds, private equity, and structured products offer opportunities beyond traditional equities and bonds, but they are also riddled with complexity, uncertainty, and risk. Over the years, I've learned that managing alternative investments successfully is as much about mindset as it is about metrics — understanding human behavior, market dynamics, and the interplay of global forces. Introducing quantum thinking into this mix transforms both perspective and practice.

I recall a meeting with a private equity fund in Switzerland that was exploring early-stage quantum computing startups. At first glance, the pitch felt abstract: algorithms, superposition, entanglement — words more common in physics papers than in portfolio reports. But as we delved deeper, I realized the opportunity was not merely technological; it was philosophical. The fund's approach mirrored my own investment philosophy: probabilistic thinking, scenario planning, and careful integration into a diversified strategy. The founders were not promising certainty; they were offering structured exploration, guided by rigorous analysis. For me, this was a familiar language, now amplified by frontier technology.

Quantum concepts like superposition and entanglement provide metaphors for portfolio design. In alternative investments, outcomes are rarely binary. A hedge fund strategy may succeed in one market regime and underperform in another. A private equity investment may yield extraordinary returns, or it may take years to mature. By thinking quantum, I embrace multiple potential outcomes simultaneously. Risk is not a single number to be minimized; it is a landscape of possibilities to be understood, mapped, and navigated. This mindset changes how I evaluate alternative opportunities, especially those at the cutting edge of technology.

Another lesson comes from data and scenario analysis. Quantum computing promises the ability to process vast, interrelated datasets far beyond classical capacity. While these capabilities are still emerging, the principles apply now. In due diligence for alternative investments, I consider multiple variables simultaneously: macroeconomic trends, sector innovation, competitive positioning, and regulatory changes. Quantum-inspired thinking encourages me to run multiple scenarios in parallel, anticipating not only likely outcomes but also less probable, high-impact events. This approach is particularly valuable when allocating capital to frontier technologies, which are inherently uncertain but potentially transformative.

Integration is key. Alternative investments are most effective when part of a broader multi-asset strategy. Private equity, venture capital, and hedge funds provide diversification and access to non-linear returns, but they also introduce illiquidity, complexity, and concentration risk. Quantum thinking reminds me to model interactions between these assets probabilistically, considering correlations under different market regimes. It encourages patience, iterative evaluation, and flexibility — qualities that are as important as financial metrics when navigating frontier opportunities.

For family offices and UHNW investors, the philosophical implications are equally relevant. Investing in alternative assets with a quantum lens is not about chasing novelty. It is about aligning capital with curiosity, foresight, and long-term vision. When we discuss allocations to private equity funds developing quantum cryptography or sustainable energy technologies, the conversation transcends numbers. It becomes a conversation about values, impact, and legacy. It is an invitation to participate in innovation responsibly, blending intellectual curiosity with financial stewardship.

One concrete example comes from evaluating a venture capital fund focused on quantum-inspired logistics solutions in Latin America. The team was tackling supply chain optimization using algorithms that classical computers could not efficiently process. My due diligence involved not only financial and technical assessment but also cultural and regional understanding. Could these solutions scale across emerging markets? How resilient were the teams to operational challenges? How did local regulations impact implementation? By bridging my experience in Latin American markets with frontier technology insight, I was able to recommend an allocation strategy that balanced ambition with prudence.

Collaboration is another pillar of this approach. Quantum-enhanced alternative investing thrives where diverse perspectives intersect. At conferences, I've seen physicists, portfolio managers, and entrepreneurs collaborate to explore real-world applications. In panel discussions with women in finance or students at international business schools, I emphasize that the human element is as critical as the technology.

Understanding motivations, behavior, and vision is indispensable for responsible investment in frontier technologies.

Finally, patience and adaptability define success in alternative investing with a quantum lens. Unlike liquid markets, alternative assets often require multi-year horizons, iterative evaluation, and the ability to pivot as circumstances evolve. Quantum thinking reinforces the importance of embracing uncertainty, maintaining curiosity, and continually recalibrating strategies. It reminds me that capital allocation is not a prediction of a single outcome but a participation in a complex, probabilistic ecosystem.

In conclusion, applying quantum thinking to alternative investments is not about mastering algorithms or qubits. It is about mindset, methodology, and integration. It is about embracing complexity, anticipating multiple outcomes, and connecting frontier technologies to practical portfolio strategies. It is about guiding investors — institutions, family offices, and UHNW individuals — to allocate capital responsibly, thoughtfully, and creatively. Most importantly, it is about translating curiosity into action: building portfolios that are not only resilient and diversified but also positioned to participate in shaping the future of finance and technology.

## 17.12 Quantum and Sustainability: Investing in the Future Responsibly

Sustainability is no longer an aspiration. It is a necessity. Across industries, regions, and investment strategies, the question is no longer whether environmental and social impact matter — it is how to integrate them intelligently, rigorously, and creatively. When we bring quantum thinking into sustainability, the possibilities expand exponentially. The lens of quantum finance does not only sharpen our view of risk and opportunity; it reframes the way we measure, evaluate, and act on long-term impact.

I first realized this connection at a conference in Geneva, where a team presented quantum algorithms applied to energy grid optimization. Classical models struggled to balance supply, demand, and carbon emissions in real time. Quantum-inspired approaches, by processing multiple variables and scenarios simultaneously, offered the potential to reduce inefficiencies and minimize environmental impact at scale. As a portfolio manager, it struck me immediately: this was not just about science or technology; it was about translating innovation into tangible societal benefit.

Sustainable investing, like alternative investing, is inherently probabilistic. Climate scenarios, regulatory evolution, and technological breakthroughs all influence outcomes. Quantum thinking reinforces the mindset needed: embrace complexity, anticipate multiple pathways, and prepare portfolios to adapt as new information emerges. The same principles that guide frontier technology allocation — rigorous due diligence, scenario modeling, and patient observation — are invaluable when evaluating green infrastructure, energy transition projects, and ESG-driven private equity.

In practice, integrating quantum-informed sustainability requires both strategy and imagination. For family offices and UHNW investors, it means looking beyond traditional ESG metrics and considering how emerging technologies can amplify impact. For example, private equity funds using quantum computing to optimize battery storage, smart grids, or carbon tracking are not speculative side projects; they are investments where scientific innovation directly aligns with sustainability outcomes. Allocating capital here demands curiosity, discernment, and confidence — knowing that the path is not linear but guided by careful evaluation and long-term vision.

Regional context is equally important. In Europe and Switzerland, strong regulatory frameworks, research institutions, and public-private partnerships create fertile ground for quantum-enabled sustainability solutions. Investors can participate through venture funds, impact-oriented private equity, or structured multi-asset products that combine traditional returns with environmental and social objectives. In Latin America, the story is different but equally compelling. Emerging markets present unique opportunities for transformational projects — solar farms in Chile, carbon-efficient logistics networks in Brazil, or sustainable agriculture in Mexico. Here, patience, local insight, and a deep understanding of regulatory and cultural nuances are essential to translate frontier innovation into measurable impact.

I also see a human dimension to sustainability investing that resonates with my philosophy. Engaging with women-led teams, emerging entrepreneurs, and socially driven innovators is not just about representation — it is about unlocking new perspectives and creative solutions. Speaking at Women in Web3 and 100 Women in Finance panels, I've seen firsthand how curiosity, collaboration, and diversity amplify results. Young professionals and students in my international economics and finance courses bring fresh eyes to old problems, asking questions like: “How can quantum computing reduce carbon footprint in supply chains?” or “How do we balance risk, return, and societal impact simultaneously?” Their questions are not abstract; they are invitations to innovate, to rethink the role of capital in shaping a sustainable future.

Due diligence in quantum-enabled sustainability investments goes beyond financial metrics. It requires assessing technology readiness, operational scalability, and societal impact. Founders, engineers, and scientists become part of the evaluation process, and investors must translate technical possibilities into practical, responsible allocation decisions. This is where my role as a bridge becomes critical: connecting innovators with institutional capital, aligning vision with prudence, and ensuring that ambition is matched by rigorous analysis.

Ultimately, quantum thinking amplifies the essence of sustainable investing: seeing the interconnections, embracing complexity, and preparing for multiple futures simultaneously. Just as quantum algorithms model superposition, a thoughtfully designed portfolio can balance financial returns with measurable environmental and social outcomes. This perspective transforms investing from reactive to proactive — from chasing short-term returns to co-creating the systems that will define the next generation of wealth and impact.

For investors willing to step into this frontier, the message is clear: approach quantum-

enabled sustainability with curiosity, patience, and rigor. Align allocations with long-term objectives, collaborate with technical and scientific experts, and never underestimate the value of perspective — both regional and human. In doing so, portfolios become not only financially resilient but also ethically responsible and socially meaningful. The opportunity is to invest in the future in a way that creates lasting impact — for clients, for society, and for the planet.

Quantum and sustainability together illustrate a principle I've carried throughout my career: innovation without responsibility is incomplete. Whether in private banking, family offices, or institutional portfolios, we have the privilege and the duty to deploy capital where it can generate both wealth and positive change. By bridging frontier technology, probabilistic thinking, and human-centered values, we can build strategies that are as visionary as they are disciplined, as creative as they are accountable.

As we look ahead, the question for investors is not whether to explore quantum or sustainability — it is how to do so thoughtfully. This is where experience, perspective, and curiosity converge. It is where a global understanding of markets, technology, and human behavior translates into portfolios that are innovative, resilient, and purposeful. And it is where the journey — from the corridors of central banks to the frontier of quantum-enabled sustainable investing — continues to unfold, one deliberate, curiosity-driven decision at a time.

## 17.13 Conclusion: Curiosity, Conviction, and the Future of Investing

Reflecting on this journey, one truth has become clear: finance is more than numbers, and technology is more than tools. Both are, at their heart, about people, decisions, and possibilities. My path from the corridors of the Venezuelan Central Bank, through classrooms in Grenoble, to the innovation hubs of Geneva, has taught me that curiosity — the willingness to ask questions, to explore beyond what is known — is the most valuable asset an investor can carry.

Quantum finance represents a frontier not only of technology but of thinking itself. It challenges us to see multiple possibilities at once, to embrace complexity without fear, and to navigate uncertainty with elegance. In this sense, it is both a metaphor and a practical framework for portfolio management: we allocate capital not to predict the future, but to participate responsibly in shaping it.

For investors, this means more than understanding qubits or algorithms. It means approaching frontier opportunities with rigor, patience, and vision, balancing innovation with stability, curiosity with discipline. For family offices and UHNW investors, it is an invitation to integrate forward-looking technologies into portfolios that reflect long-term values, aspirations, and societal impact. For institutional and tech investors, it is a call to engage deeply, leverage networks, and cultivate partnerships that bridge finance and frontier science.

Yet perhaps the most important lesson is personal. Quantum finance, like any

frontier, thrives where human creativity, empathy, and collaboration intersect with knowledge. Whether mentoring young investors, speaking at a conference, or advising a client, my goal remains the same: to translate complexity into clarity, to empower confident decisions, and to inspire curiosity that transforms potential into action.

As you close this chapter, I hope you carry a sense of possibility. The quantum frontier is not distant; it is here, woven into the choices we make, the strategies we design, and the futures we imagine. It is both a challenge and an invitation: to invest responsibly, to innovate thoughtfully, and to embrace the unknown with curiosity and courage.

This is the story I bring to every meeting, classroom, and stage: a story of curiosity, conviction, and the belief that finance, when approached with heart and intellect, can shape not only wealth but the future itself.

If you are ready to explore this frontier, to invest boldly yet responsibly, and to imagine what lies beyond conventional boundaries, I would be honored to join you on that journey.



## Chapter 18

# Standards for Quantum Technology in Finance

### *Mapping a fragmented global standardisation landscape*

Melissa Hernández

I earned my Law Degree from the University of Costa Rica. After graduating, I continued my studies while practising law. And completed a specialisation in Notary and Registry Law, an LLM in Alternative Dispute Resolution, and another in Economic Law with a focus on International Trade. Looking back, I see a common thread: I always wanted to understand how systems are built, not just how they work in practice.

My early career followed that same goal. At the Costa Rican Chamber of Commerce, I did more than manage contracts. I helped shape legislation by advising the Board of Directors, coordinating with national chambers, and contributing to important bills. Later, at Banco BCT, I managed compliance and legal risk in the banking sector. Each job taught me something new about banking, finance, corporate structure, data protection, and governance. When I joined BLP Legal in 2019, my focus shifted. I advised Prival Bank on organisational and compliance matters, standardised credit documents, and restructured legal workflows. I found myself working where law, finance, and institutional design meet. My questions became more about what the rules say and less about what the rules have not yet addressed.

Therefore, in 2021, I chose to move to the Netherlands and enrolled in the Master's in International Technology Law at Vrije Universiteit Amsterdam. This was not a break from my past work, but rather a new direction. The program focused on privacy, data governance, and emerging technologies. Its curriculum aligned with

---

This chapter is a featured contribution to the monograph *A Portrait of Quantum Technologies in Finance*, edited by Oswaldo Zapata, PhD, and published by The Quantum Finance Boardroom (2026). All rights reserved by the chapter's author.

my previous experience and introduced me to a new language and a future-oriented perspective.

But everything came together in January 2024, when I became a PhD Fellow at Leiden University’s Faculty of Law. The research question I had been developing over years of legal practice finally found its place: What governance frameworks does the financial sector need for the quantum era?

## 18.1 Why Quantum Finance?

Quantum technology is not a general idea, which is why it is important to address it now. Once quantum computers become powerful enough, they will have the potential to break the cryptographic protocols that protect the global financial system. They will also be able to improve many financial applications in ways that classical computers cannot. Preparing to move to quantum-safe cryptography and setting standards for its regulation are limited opportunities, which I call the quantum readiness problem. The next regulatory challenges are already being shaped by developments in physics.

## 18.2 What Standards Mean in Quantum Finance

Currently, standards for quantum technology in finance encompass several categories of guidelines. Formal technical standards are established by organisations such as the International Organisation for Standardisation (ISO), the European Telecommunications Standards Institute (ETSI), the Internet Engineering Task Force (IETF), the US National Institute of Standards and Technology (NIST), and the ANSI Accredited Standards Committee X9 (ASC X9). Supervisory expectations are set by entities including the Basel Committee, the European Central Bank, and national regulatory authorities. Additional sector-specific guidance is provided by groups such as UK Finance, the Financial Conduct Authority (FCA), and the G7 Cyber Expert Group. Best practices are further shaped by major vendors, consultancies, and industry associations.

Collectively, these guidelines address both defensive measures, such as transitioning to post-quantum cryptography (PQC), maintaining cryptographic agility, and ensuring system resilience. They also encompass value-generating applications, including the secure implementation of quantum algorithms in risk modelling, trading, and optimisation. In practice, these standards define what constitutes good practice for financial institutions aiming to mitigate quantum-related risks and leverage emerging opportunities.

Currently, most standardisation efforts focus on mitigating quantum risks to existing financial systems by adopting post-quantum cryptography (PQC), rather than on establishing standards for the application of quantum algorithms in business contexts. The PQC process led by the National Institute of Standards and Technology (NIST), along with related initiatives in ISO, ETSI, IETF, and ASC X9, is defining

cryptographic standards for regulated sectors, including finance. In contrast, guidance on the use of quantum technologies for modelling and optimisation primarily appears in research publications and industry reports, rather than in formal standards.

## 18.3 Quantum Standards Summary

This document summarises key current quantum technology standards relevant to the financial sector. The document offers a general summary of current and emerging standards activities in this area and does not seek to be exhaustive or to cover all existing or ongoing initiatives.

Current standardisation efforts in quantum technology primarily focus on cryptography, with particular emphasis on transitioning from RSA and ECC to post-quantum cryptography (PQC). Broader standards for quantum computing, including performance metrics, benchmarking, quantum machine learning, and simulation, remain in the early stages of development. The financial sector is at the forefront of PQC adoption, driven by the harvest-now-decrypt-later threat, requirements for long-term data storage, and recent regulations such as DORA, NIS2, the Cyber Resilience Act (CRA), and PCI DSS 4.0. The standards landscape remains fragmented across various organisations and countries. Consequently, internationally active firms must comply with at least six standards bodies and four regulatory regimes, as no unified “quantum finance” standard currently exists.

## 18.4 Cryptographic Standards Layer

The cryptographic standards layer encompasses general standards for cryptographic algorithms and protocols that are not sector-specific. Standards, regulations, and procurement guidelines within the financial sector reference this layer. Key organisations involved are NIST, ETSI, ISO/IEC JTC 1, IEEE, CEN/CENELEC JTC 22, ITU-T, IEC SEG 14, and various national standards bodies.

**NIST (United States)** The NIST Post Quantum Cryptography standardisation effort is the primary global reference for PQC algorithms, with selected schemes and migration guidance published since 2022. NIST’s Post Quantum Cryptography Standardisation Project, launched in 2016 and concluding its third round in 2022, has published the first three Federal Information Processing Standards (FIPS) specifying quantum-safe primitives.

FIPS 203, *ML KEM, Module Lattice Based Key Encapsulation Mechanism*, FIPS 204, *ML DSA, Module Lattice Based Digital Signature Standard*, and FIPS 205, *SLH DSA, Stateless Hash Based Digital Signature Standard* were published in August 2024. Two further standards remain in the pipeline.

For migration planning, three NIST documents are as critical as the FIPS standards. SP 800 131A Rev.3, *Recommendation for Cryptographic Key Generation* defines algorithm transitions, including the move toward quantum-resistant algorithms

for digital signatures and key establishment. NIST IR 8547 (Initial Public Draft, November 2024) *Transition to Post-Quantum Cryptography Standards* provides the transition timeline, recommending the deprecation of RSA 2048 and ECC 256 after 2030, with full disallowance of all quantum-vulnerable asymmetric algorithms after 2035. The NSA's *Commercial National Security Algorithm Suite 2.0 (CNSA 2.0) Cyber Security Advisory (CSA)* makes these dates binding for US National Security Systems, influencing the private sector through supply chain requirements.

**ETSI (Europe)** Within the European Telecommunications Standards Institute, two structures are relevant. The Technical Committee on Cybersecurity Working Group on Quantum Safe Cryptography (TC CYBER WG QSC) handles PQC, and the Industry Specification Group on Quantum Key Distribution (ISG QKD) handles QKD. Beyond NIST-derived primitive algorithm, ETSI's QSC group plays a leading role in practical migration guidance, emphasising the need for crypto agility and long-term protection of data exposed to harvested-now-decrypt-later attacks. ETSI explicitly lists securing financial and banking transactions as a core motivation for quantum-safe standards and has published materials to assist organisations in planning their transition.

Current published deliverables of direct relevance to financial infrastructure include ETSI TS 103 744 V1.2.1 on *CYBER; Quantum-Safe Hybrid Key Establishment*. The technical reports TR 103 965 V1.1.1 on *CYBER, Quantum-Safe Cryptography, Impact of Quantum Computing on Cryptographic Security Proofs*. TR 103 966V1.1.1 on *CYBER; Quantum-Safe Cryptography, Deployment Considerations for Hybrid Schemes*. And TR 103 967V1.1.1 on *CYBER; Quantum-Safe Cryptography, Impact on Quantum Computing on Symmetric Cryptography*; and the migration-focused documents TR 104 016V1.1.1 on *CYBER; Quantum-Safe Cryptography, a Repeatable Framework for Quantum Safe Migrations*; and TS 104 015 V1.1.1 on *CYBER; Quantum-Safe Cryptography, Efficient Quantum Safe Hybrid Key Exchanges with Hidden Access Policies*.

A further significant work item is the Authenticated Quantum-Safe Hybrid Key Establishment (AQSHKE), project, funded by the European Innovation Council and the SMEs Executive Agency (EISMEA). As an European standard (EN) -rather than a technical report or a technical specification- AQSHKE could be subject to a future Commission standardisation request. A parallel work item on crypto agility in software-based quantum safe migrations is also underway, addressing requirements that have surfaced in DORA's Regulatory Technical Standards (Delegated Regulation (EU) 2024/1774) and in the NIS2 Implementing Regulation (EU) 2024/2690.

**ISO/IEC JTC 1 (international)** The Joint Technical Committee 1 of ISO and the International Electrotechnical Commission (IEC) is the primary international route for adopting standards by reference into national law. Three structures matter in the quantum context: WG 14 for terminology, SC 27 for security techniques, and SC 42 for Artificial Intelligence.

*ISO/IEC 4879:2024, Information technology -Quantum computing -Vocabulary*, produced by JTC 1 WG 14 (now Quantum Information Technology, transferred to the new IEC/ISO JTC 3), is the anchor terminology standard for quantum computing. Within SC 27, Working Group 2 (Cryptography and Security Mechanisms) governs

the inclusion of PQC primitives into the international cryptography catalogue through its Standing Document 8 (SD8) framework. The existing ISO/IEC 14888-4:2024 *Digital signatures with appendix-Part 4: Stateful hash-based mechanisms*, series on digital signatures and ISO/IEC 18033-1:2021 *Encryption algorithms-Part 1: General*, series on encryption algorithms are expected to be amended rather than replaced.

SC 27 Working Group 3 handles security evaluation. It updated *ISO/IEC 19790:2025* on security requirements for cryptographic modules, and *ISO/IEC 24759:2025* on test requirements for cryptographic modules. WG 3 is also producing a Technical Report (ISO/IEC AWI TR 25544) on the effect of specific transmission media on the security evaluation of QKD .

**IEEE SA** The Institute of Electrical and Electronics Engineers Standards Association runs several working groups on quantum technologies. These groups produce more architecture and benchmarking-oriented deliverables than cryptographic specifications, although projects such as IEEE P1943 *Post-Quantum Network Security* and P1947 *Quantum Cybersecurity Framework* are under development. The most consequential ongoing project is IEEE P7130 *Standard for Quantum Technologies Definitions*, which aims to provide a common vocabulary for technical and regulatory work.

**CEN/CENELEC JTC 22** The European Committee for Standardisation (CEN) and the European Committee for Electrotechnical Standardisation (CENELEC) jointly operate JTC 22 on Quantum Technologies, with Working Group 3 focusing on quantum computing and simulation and Working Group 4 on quantum communication and cryptography. WG 4 plays a central role in adapting global work on quantum communication and QKD to the European context.

**IEC SEG 14** The IEC Standardisation Evaluation Group 14 (SEG14) on Quantum Technologies delivered a quantum standardisation roadmap in 2023, recommending the creation of an international joint technical committee; a recommendation that led to the establishment of ISO/IEC JTC 3 on Quantum Technologies in 2024. JTC 3 now coordinates internationally with regional bodies such as CEN-CENELEC JTC 22 and with IEEE working groups, particularly on sensing and metrology. CEN-CENELEC/JTC 22 published its first deliverable, TR 18202:2025 *Layer model of Quantum Computing*, covering the universal gate-based quantum-computing model across multiple physical systems.

## 18.5 The Financial Sector

The financial sector adapts the general standards described above into sector-specific rules for banking, securities, and payments. In the 2024-2025 period, the financial sector made substantial progress in financial PQC guidance. Four categories of actors operate in parallel: 1) formal standards development organisations such as NIST, ETSI and ISO/IEC, which produce the underlying algorithmic and technical specifications. 2) Global political regulatory bodies such as the G7 Cyber Expert Group, the Financial Stability Board (FSB), the Basel Committee on Banking

Supervision (BCBS), and the Bank for International Settlements (BIS). 3) Sector-specific quantum forums producing operational guidance, such as Europol’s Quantum Safe Financial Forum (QSFF), the Financial Services Information Sharing and Analysis Center (FS-ISAC) PQC Working Group, and the World Economic Forum’s Quantum Security workstream. 4) Contractually-binding technical standards such as the Payment Card Industry Data Security Standard (PCI DSS), which apply to any entity processing payment-card data.

**ISO TC 68 and ANSI X9** *ISO Technical Committee 68 (TC68)* covers banking, securities and other financial services through three subcommittees: *SC 2* on financial services, security, *SC 8* on reference data for financial services, and *SC 9* on information exchange for financial services. The Accredited Standards Committee X9 (ASC X9), accredited by ANSI, serves as both the TC 68 Secretariat and the US Technical Advisory Group, which means in practice that a significant share of international financial services cryptography standards is first developed within the X9 framework and subsequently brought into the ISO process.

ISO TC 68 has framed quantum computing primarily as a cryptographic risk to financial infrastructure and as a driver of the migration to PQC. Its 2021 working-group paper, “*Quantum Computing and the Financial Services Industry*”, outlined the threat of Shore’s algorithm to public key cryptography (RSA and ECC). Following NIST’s publication of FIPS 203, 204 and 205 in August 2024, ISO/IEC JTC 1 SC 27 is now integrating these algorithms into its catalogue, and ISO/TC 68 is developing sector-specific implementation guidance. SC 27/WG 2 also maintains its Standing Document 8 (SD 8) on post-quantum cryptography, which functions as the principal reference for ongoing work in this area.

The Accredited Standards Committee X9 (ASC X9) is the ANSI-accredited body that sets technical standards for the US financial services industry. ASC X9 published an Informative Report *Quantum Computing Risks to the Financial Services Industry* (ASC X9 IR-F01-2022), which provides a detailed overview of quantum computing, the cryptographic risks, expected timelines, and migration steps for financial institutions. In 2025, ASC X9 released a *Post Quantum Cryptography Financial Industry Readiness Needs Assessment* to guide financial services management to facilitate a smooth transition to PQC.

Within the X9 framework, several concrete instruments are particularly relevant:

- ANSI X9.69-2023, *Framework for Key Management Extensions*, updated in 2023 to incorporate provisions enhancing resilience against quantum-enabled attacks on symmetric key management.
- Draft X9.146 (Quantum TLS), specifying how the NIST PQC algorithms integrate into the TLS protocols for financial industry use, with the aim to provide financial institutions a hybrid migration path for protecting data against future quantum computer attacks.
- Technical Report ASC X9 TR 50-2019, Quantum techniques in Cryptographic Message Syntax One of the earliest X9 deliverables on integrating quantum-aware techniques into financial messaging. Actual version ANSI X9.73-2023 Cryptographic Message Syntax (CMS).

**BIS, central banks** The Bank for International Settlements has emerged as the central convening body for PQC work at the central bank level. Its key contributions are two BIS Papers issued by the Monetary and Economic Department: BIS Papers No.149 (2024) *Quantum computing and the financial system: opportunities and risks* which provides the analytical framing of the dual nature of quantum technologies for finance; and BIS Papers No.158 (2025) *Quantum readiness for the financial system: a roadmap*, the first systematic phased migration framework directed at the financial system. While not setting binding milestones, the roadmap proposes a phased approach, awareness and inventory in 2025-2026, planning in 2026-2028, and major migrations in the late 2020s.

**Europol Quantum Safe Financial Forum** The Europol Quantum Safe Financial Forum (QSFF) operates under the European Cybercrime Centre (EC3) and brings together European financial institutions, law enforcement agencies and policymakers. Its 2025 *Quantum Safe Financial Forum- A Call to Action* is currently the clearest pan-European signal, urging financial institutions to prioritise the transition to quantum-safe cryptography and to coordinate on timelines and technical choices.

**G7, FS ISAC, FSSCC and WEF** The G7 Cyber Expert Group’s September 2024 *Statement on Planning for Opportunities and Risks Associated with Quantum Computing* integrated quantum readiness into the G7’s regular cyber resilience supervisory cycle. On the industry coordination side, the *Timeline for Post Quantum Cryptographic Migration* white paper, published in 2025 and authored jointly with the FS-ISAC Post-Quantum Cryptography Working Group, the Canadian Forum for Digital Infrastructure Resilience (CFDIR) Quantum-Readiness Working Group, and the Quantum Safe Financial Forum, introduced the “crypto procrastination” framing now widely cited in migration scheduling.

The World Economic Forum’s January 2024 *Quantum Economy Blueprint* and its joint WEF-FCA report *Quantum Security for the Financial Sector: Informing Global Regulatory Approaches*, remain relevant non-SDO coordinating documents, although they have been complemented and partially superseded, at least in operational and technical detail, by the 2025 BIS roadmap.

**FSB, BCBS and PCI DSS** The Financial Stability Board’s Cyber Lexicon (2018, revised 2023) and the Basel Committee on Banking Supervision’s *Principles for Operational Resilience* (2021) provide the vocabulary and conceptual frameworks through which quantum risk is increasingly being mapped into prudential supervision, even though neither document explicitly mentions quantum technologies.

The *Payment Card Industry Data Security Standard (PCI DSS) v4.0.1* strengthens generic cryptographic-governance requirements which, while not explicitly addressing post-quantum migration, provide a practical foundation for quantum-readiness planning in cardholder data environments. Although PCI DSS is technically an industry standard rather than a public regulation, its enforcement by payment networks via acquirer contracts makes it functionally binding on every entity that processes payment card data, and consequently a powerful lever for cryptographic migration in practice.

## 18.6 The EU Regulatory Layer

The EU regulatory layer is the mechanism through which the horizontal and sectoral standards described above become legally binding on financial entities operating in the European market. Six instruments matter in current force; a seventh is expected shortly.

**Commission Recommendation 2024/1101 and the NIS Cooperation Group PQC workstream** In April 2024, the Commission issued *Recommendation (EU) 2024/1101 on a Coordinated Implementation Roadmap for the transition to Post Quantum Cryptography*. The instrument is formally non-binding but politically anchored through the NIS Cooperation Group, which established a dedicated PQC workstream co-chaired by Germany (BSI), France (ANSSI) and the Netherlands. Part 1 of the Roadmap as Version 1.1, published in mid-2025, articulates a phased timetable, with Member States expected to define and launch national PQC strategies by the end of 2026, to begin widespread implementation of PQC solutions across critical public infrastructure by 2030, and to complete the transition by 2035.

**DORA** *The Digital Operational Resilience Act (Regulation (EU) 2022/2554)* has applied to financial entities since January 2025. The PQC “hook” sits in Commission Delegated Regulation 2024/1774, which adopts the regulatory technical standards on ICT risk management. Under these RTS, financial entities are required to remain abreast of cryptanalytic developments, explicitly including “threats from quantum advancements” and to provide for the updating of cryptographic technology. In practice, this is the closest thing to a binding PQC obligation in EU financial services law without the words “post quantum” appearing in the text.

**NIS2 and the Implementing Regulation** Directive (EU) 2022/2555 (NIS2) applies to key entities across critical sectors, including parts of the financial ecosystem not covered by DORA, as well as the digital service providers and trust service providers on which financial entities rely. Commission Implementing Regulation (EU) 2024/2690 sets out technical and methodological requirements for those service providers, including state-of-the-art cryptographic policy, the implementation of cryptographic agility mechanisms enabling rapid algorithm replacement, and alignment with international standards, including ISO/IEC 27001, ISO/IEC 27002 and ETSI EN 319401 and technical specifications, such as CEN/TS 18026:2024. In combination with DORA, this makes PQC migration and crypto agility a live compliance issue for a broad set of financial sector entities and their providers.

**Cyber Resilience Act** Regulation (EU) 2024/2847, the *Cyber Resilience Act*. It applies essential cybersecurity requirements to products with digital elements placed on the EU market. Its compliance architecture relies on harmonised standards developed under standardisation request M/606, channelled mainly through CEN/CLC JTC 13 and the new ETSI EUSR. Although the ETSI AQSHKE and CEN/CENELEC JTC 22 outputs are not currently included in M/606, these may inform future harmonised standards if the Commission issues complementary mandates. In that case, manufacturers of financial-adjacent embedded hardware would be required to demonstrate compliance.

**GDPR** Article 32 of the *General Data Protection Regulation (GDPR)* on security of processing is increasingly interpreted by data protection authorities to include expectations for quantum readiness, given the publication of the NIST PQC standards and the harvest-now-decrypt-later threat.

**ENISA guidance** The European Union Agency for Cybersecurity (ENISA) has produced reports, including *Post Quantum Cryptography: Current State and Quantum Mitigation (2021)* and the *NIS2 Technical Implementation Guidance (2025)*, recommending the adoption of quantum-safe cryptography for long-term data protection. ENISA documents are not binding, but operate as interpretative guidance on what constitutes “appropriate” security under EU cybersecurity law.

**EU Quantum Strategy and the forthcoming Quantum Act** The European Commission published its *EU Quantum Strategy* on July 2025, aiming to build a resilient quantum ecosystem by 2030. A *EU Quantum Act* is expected to be proposed in Q2 2026, structured around three objectives: coordinating R&I investment, scaling industrial capacity, and reinforcing supply-chain resilience and governance.



# Chapter 19

## Building Quantum-Safe Finance

*Securing infrastructure, assets, and trust in the hybrid attack era*

Olga Mamlyga

### 19.1 Introduction

My journey into the world of quantum technology began at the intersection of curiosity and necessity — a desire to understand not just the theory of tomorrow’s technologies, but their real-world impact on business, security, and finance. This quest led us to establish Quantum Scouts, a company dedicated to navigating the complexities of emerging technologies and making them accessible, relevant, and actionable for innovators and decision-makers alike.

At its core, Quantum Scouts champions practical innovation: translating advanced concepts in quantum mechanics and secure communication into tools and strategies that address pressing challenges in cybersecurity, infrastructure resilience, and digital transformation. This perspective resonates deeply with me, particularly as I explore how quantum technologies — from secure quantum networks to resilient encryption — are reshaping financial systems, risk models, and strategic frameworks in ways we are only beginning to grasp.

This perspective inevitably leads to a more uncomfortable question: are we focusing on the right solutions at all?

Much of today’s discourse around quantum readiness — particularly in finance and banking — is framed around Post-Quantum Cryptography (PQC) as a future

---

This chapter is a featured contribution to the monograph *A Portrait of Quantum Technologies in Finance*, edited by Oswaldo Zapata, PhD, and published by The Quantum Finance Boardroom (2026). All rights reserved by the chapter’s author.

safeguard, a problem to be solved “in time.” Yet the deeper one engages with real-world systems, geopolitical dynamics, and operational security, the clearer it becomes that this framing is incomplete. Quantum resilience is not a single upgrade, a standards roadmap, or a cryptographic replacement. It is a systemic challenge — technological, political, and societal — that cannot be postponed without consequence.

It is from this vantage point that a follow-up on PQC becomes not premature, but necessary.

Why conduct a follow-up on PQC before the finance and banking sector has even implemented it? In a nutshell: PQC is not a complete solution for achieving post-quantum resilience or for keeping data and data centres secure. The challenge of shielding data, devices, values, and wealth in an impenetrable way is far more complex, and relying solely on PQC is a dangerous, costly, and ultimately destructive misinterpretation.

Returning from the WEF in Davos 2026, let me share the true concerns voiced by leaders in finance, industry, and politics: the classic approaches no longer work. The large-scale distribution of fake news via the internet has reshaped the world and hollowed out the foundations upon which our societies were built. Vast resources are being poured into AI and quantum technologies by both benevolent and malicious actors, providing the tools to exploit the uncertainties of the emerging world order. Meanwhile, the disruptive and egocentric politics of a few players have eroded the international rules and frameworks that once ensured maximum stability and peace.

When we spoke about politics and society four years ago, the resonance was limited and often negative. At Davos — although we had deliberately held back — the response was finally aligned, as top leaders had begun to catch up. We are neither smarter nor better for having recognized these developments earlier. Rather, we were actively building real-world quantum and AI solutions — not writing academic papers or conducting isolated lab experiments, but working as industry practitioners with a combined experience of hundreds of years. This practical engagement is what allowed us to understand early on the disruptive and destructive political and economic impact these technologies would have on the world.

The entire banking and finance sector is built upon RSA encryption, and according to the latest publications by IBM, a sufficiently powerful quantum computer will be able to break every RSA key by 2029. Considering the time required to implement new technologies in the banking sector — typically seven to twelve years — it already appears too late to shield data, transactions, wealth, and, ultimately, the stability of our industries and societies. Consequently, the traditional implementation approaches will not be sufficient.

This leads to another undeniable fact: even if PQC were implemented before 2029, successful attacks would still be likely. PQC remains vulnerable for at least two reasons (we have identified more, but we do not publish vulnerabilities—we design shielding solutions):

- NIST’s recommendation for PQC is to apply post-quantum encryption on top of classically encrypted data, as there is no certainty that PQC alone could withstand a hybrid attack carried out by a powerful quantum computer in

conjunction with one or multiple high-performance classical computers.

- PQC schemes have already been broken via side-channel attacks prior to the final round of the standardization evaluation.

Why is creating post-quantum encryption so difficult? Of the 82 algorithms submitted in 2017, only 69 were considered complete, and just three have survived to date, declared by NIST as safe and suitable. However, even these carry an asterisk: they must be handled with caution because we have no certainty regarding hybrid attacks or the potential synergy of classical AI and quantum AI orchestrating cyberattacks against PQC.

Moreover, hackers do not play by the rules of academia, and academic research has failed to identify every possible attack vector. Industry practitioners are not exempt from this shortcoming; many — if not most — have struggled in this area as well. Malicious actors, often motivated by more than mere criminal intent, have embedded malware into nearly every system in the world. These actors are more concerning than traditional cybercriminals because their objectives go beyond financial gain: state-sponsored malware is designed as a tool of warfare. The details of these threats cannot be shared here and are partly restricted by law.

We do not wish to downplay the importance of PQC. Yet it is essential to shatter the illusion that PQC alone is sufficient to safeguard sensitive data, or the values and wealth that depend on it.

The high expectations for QKD, championed by academia and some major telecommunications companies, were reconsidered following a decision by the U.S. Department of War (a name our team finds unfortunate). Years ago, at DALO (the Danish Ministry of Defence Acquisition and Logistics Organisation) in Copenhagen, we had the opportunity to explain to the relevant U.S. Navy representatives that QKD was not suitable for deployment in defense. In 2026, QKD was officially deemed unfit for military purposes and was subsequently discontinued. While this does not offer an immediate solution, it also highlights that QKD is not well-suited for the banking and finance sector. Given the sensitive nature of this topic, we have maintained a low public profile and only discussed it in a few private meetings, as broader exposure could pose significant risks for those of us developing these technologies.

Post-quantum safety can only be achieved through a complex combination of deep security and processing measures. The level of defence required depends on the sensitivity, importance, and value of the services or data involved; attackers will expend disproportionately greater resources to overcome your defences.

The following introduction is neither complete nor does it provide a full solution. We are builders of ultimately secure cyber systems, not educators. If your IT team has exceptional industrial experience, consultancy may help — but development and implementation could take too long to protect your most critical data and transactions. Be aware that you are racing uphill, with a finish line set sometime in 2029. You may be critical of quantum computing — and that is fair, it may not suit many use cases — but the ability to break encryption is deeply embedded in its DNA, and there is no doubt it will eventually break RSA encryption.

## Key guidelines:

- Set up a complex, multi-layered architecture impenetrable to quantum computers. Ensure the architecture neutralizes every attack vector. Consider AI-supported hacks, quantum-AI-supported attacks, hybrid attacks (three types: external, internal, and physical manipulation at any part of your network), social engineering, identity theft, backdoor exploits, failures of critical infrastructure services, certificate authority vulnerabilities, operating system and software attack vectors, hardware vulnerabilities, and hidden malware.
- Neutralizing attack vectors is necessary but not sufficient. You need robust engineering, real-time deep and intelligent monitoring, flawless maintenance, security redundancy, and much more.
- The capabilities of AI-assisted cyberattacks are discussed in greater detail in Section 19.5.
- Stop relying solely on red teams. Their approach is outdated against sophisticated AI attacks. Sharing defenses is akin to handing a vault blueprint to attackers. No red team has the resources, training, or backup of state-sponsored actors.
- IT monocultures lead to breakdowns, each breakdown creating opportunities for hidden malware.
- As discussed later in Section 19.7, cybersecurity governance must operate at the highest level of institutional leadership.

Another reason to be concerned about post-quantum resilience — and why PQC alone will not provide the security your company needs — is the fact that NIST invited the NSA to participate in the PQC competition, due to their recognized expertise in encryption. In fact, the NSA could have participated regardless of invitation — by law, because of its backdoors for data retrieval. Several questions immediately come to mind:

- Will these backdoors remain secure against advances in AI? Attackers know they exist, and traffic only needs to be identified, isolated, and targeted.
- Will these backdoors withstand attacks from powerful quantum computers?
- How resilient are they to complex hybrid attacks, combining the full power of quantum computers with one or more high-performance classical systems?
- If your company is not U.S.-based, is it truly in your interest for sensitive data to be transferred to the NSA?
- Can we assume the current U.S. political landscape is benevolent — either in handling your data through U.S. companies or in counteracting your business or R&D efforts?

*Note: The above list is intentionally incomplete.*

What can you expect from NIST? The National Institute of Standards and Technology is a U.S. Department of Commerce agency whose mission is to promote American innovation and industrial competitiveness. Do you want to base the safety of your data and business on an institute whose mandate is to advance American interests?

The brief introduction to AI's current capabilities was intended to give a glimpse into why PQC will likely fail to protect your data, assets, and operations — especially high-value targets such as cryptocurrencies and stablecoins. Importantly, there was no need to invoke quantum hacking to exploit PQC or the additional countermeasures implemented by major banks and financial institutions, which are meant to safeguard assets. Publicizing security measures is essentially sabotaging your own interests.

AI-assisted attacks can succeed even against hardened systems, and this capability will only increase with the arrival of commercially accessible, powerful quantum computers. Hardware resources can already be operated by less technically trained attackers thanks to advanced quantum operating systems, such as Qrisp, which organize sparse hardware resources automatically and integrate AI tools. Smarter algorithms, developed or discovered by mathematicians and AI/quantum AI, will also improve classical supercomputer performance, potentially breaking first-generation PQC within four to five years rather than ten or fifteen.

This doomsday scenario, however, can be mitigated with today's engineering. Many of the mathematical and physical impacts of quantum technologies can already be accurately predicted, allowing countermeasures to protect critical assets over the coming decades.

Safer and less expensive methods exist to ensure post-quantum immunity and protect data than relying solely on PQC. Advances in mathematics alone will force the re-encryption of all sensitive data every few years — which is inefficient, given PQC's limitations in protecting high-value targets. However, these solutions are complex and require deep understanding of the entire system: main hardware, operating system, peripheral hardware, and the full network architecture.

Most decision-makers, including those in banks and financial institutions, are reluctant to develop a comprehensive quantum security roadmap. They intuitively know that neither PQC nor QKD (requiescat in pace) is sufficient — and they are right. Companies transitioning to quantum-resistant security often sense that the solutions they are implementing will not fully protect their data and operations. PQC can be used safely as a temporary layer of defence, but only if its limitations are fully acknowledged, and the system architecture is adjusted accordingly. Until PQC matures and is proven under real-world attacks, spending heavily on it may be unnecessary when safer, more effective options exist.

Consider the cost of replacing PQC solutions every two to five years across all long-term sensitive data. Each rollout carries risks: human error, internal sabotage, and implementation mistakes can easily outweigh the security benefits of the new encryption.

We do not condemn PQC. We view it as a medium-strength supplementary security measure — a thin defensive layer — but if your core defenses cannot withstand cyberattacks without PQC, they will not endure any longer with it.

## 19.2 The Quantum Edge: The Real Disruption

Quantum Technologies (QT) are genuinely disruptive, perhaps more so than the first Industrial Revolution, and developments will occur much faster. Quantum computers are inherently destructive because of the principles of qubits and superposition. You may have read that quantum computers can evaluate all possible computational paths simultaneously, collapsing their wave function to produce the desired solution — a complex but manageable task. Yet encryption breaking is a default capability of quantum computers. This explains the extraordinary investment in quantum technologies, even as much of industry remains sceptical. Encryption underpins the cyberworld; its failure can result in massive financial losses and potentially human harm.

While traditional cyberattacks have targeted critical infrastructure for decades, the arrival of large, powerful quantum computers will transform the landscape. Exact capabilities remain uncertain, but educated predictions are feasible. Financial sectors, particularly cryptocurrencies and stablecoins, should adopt a worst-case-scenario mindset — this is not science fiction but practical caution rooted in decades of IT experience.

Implementing new security technologies in finance takes, on average, seven years — pushing full adoption to around 2032. Many cryptographers predict that by then, all asymmetrical encryption could be broken; IBM forecasts this could happen as early as 2029. Immediate action is therefore required, with a focus on security architectures resilient to both post-quantum and hybrid attacks.

With the evolution of quantum AI, systems will become increasingly opaque, exceeding the comprehension of most humans, including physicists, mathematicians, and IT engineers.

## 19.3 Categorizing Quantum Technologies

To approach this strategically, QT should be understood in four categories:

1. **Quantum Computers:** Highly destructive by default, with potential long-term benefits for humanity.
2. **Quantum Sensors:** Extremely precise, capable of embedding cybersecurity.
3. **Quantum Networks:** The term “quantum network” is often used in the literature and by industry to refer simply to the connection of quantum computers—linking multiple quantum processors to share quantum states, perform distributed computations, or simulate quantum systems collaboratively. While this concept is important for research and computational purposes, it does not provide security or resilience against quantum or hybrid cyberattacks. In other words, connecting quantum computers alone does not make your systems immune to threats; it is purely a computational network, not a protective one.

## Our solution:

In contrast, we define a Quantum Network as an industrial-grade, post-quantum resilient network designed specifically to protect critical infrastructure, financial systems, and sensitive data from quantum-powered attacks. Key features of our network include:

- Layered post-quantum encryption across all seven OSI layers — ensuring no single point of failure.
- Embedded tamper-proof quantum sensors for real-time intrusion detection and environmental monitoring.
- AI-integrated monitoring and anomaly detection, capable of correlating classical and quantum network data to detect hybrid attack patterns.
- Seamless integration with existing IT infrastructure, providing high throughput, reliability, and scalability — unlike traditional QKD deployments.
- Stealth and jamming resilience, ensuring secure communications for financial transactions, machine-to-machine operations, and critical command links.
- Dynamic post-quantum adaptability, enabling the network to update cryptographic protocols and operational parameters as new threats emerge.

In short, unlike the academic notion of quantum networks, which focuses on connecting quantum computers, *our Quantum Network is a defensive architecture: it prevents, detects, and mitigates attacks in real time, and is designed to secure the financial sector, cryptocurrencies, stablecoins, and other critical digital assets for decades to come.*

## 19.4 Quantum Communication: Industrial-Grade, Post-Quantum Resilient

The term “quantum communication” is often used in research to describe protocols such as QKD or the transfer of quantum states between nodes. While these approaches are interesting academically, they are fragile, slow, difficult to scale, and generally unsuitable for industrial use — especially for critical sectors like finance and banking. They do not provide end-to-end resilience against hybrid attacks, AI-assisted intrusions, or large-scale cyber threats.

### Our solution:

We implement industrial-grade quantum communication that is:

- Post-quantum resilient: Secure against future advances in quantum computing and mathematical attacks.
- Industrial-grade and reliable: Engineered for continuous high-speed operation in real-world networks, not delicate lab setups.
- Immune to side-channel attacks: Designed to prevent exploitation through hardware or protocol vulnerabilities.

- High-throughput: Capable of handling large-scale financial transactions, real-time data, and complex machine-to-machine communications.
- Stealth and jamming resilient: Maintains operational integrity even under targeted interference or hostile environments.
- AI-integrated: Enables intelligent traffic monitoring, anomaly detection, and automated threat mitigation while supporting machine-to-machine communication (e.g., drone-to-drone, computer-to-drone, satellite-to-missile).

In short, unlike conventional “quantum communication” approaches, *our system is a robust, scalable, and secure industrial communication network*, engineered to protect sensitive data, financial systems, and critical infrastructure for decades — resilient to both current and future quantum threats.

Hacking does not follow academic rules. Many security experts—and even new academics entering quantum research — focus on DOS, DDOS, and a few other attack types, but lack awareness of the full spectrum. We cannot publish all defensive strategies publicly, but we are capable of defending even the most exposed systems against highly complex, AI-assisted attacks.

## 19.5 The Current Cyber Threat Landscape: AI as a Force Multiplier

The cybersecurity environment facing financial institutions is undergoing a profound transformation. While the disruptive potential of large-scale quantum computing remains a central concern for the coming decade, the threat landscape has already shifted dramatically due to rapid advances in artificial intelligence. These technologies are reshaping the scale, speed, and sophistication of cyber operations, enabling attackers to automate complex tasks and conduct highly targeted campaigns with unprecedented efficiency.

Historically, sophisticated cyberattacks required significant technical expertise, extended preparation, and substantial financial resources. Today, many of these capabilities are increasingly accessible through widely available AI tools and automated frameworks. As a result, activities that once required highly specialized attackers can now be executed by a far broader spectrum of threat actors. Artificial intelligence has effectively become a force multiplier in cyber conflict.

One emerging category of attacks involves the manipulation of AI systems themselves. Prompt injection techniques can exploit language models and automated decision systems by crafting inputs designed to bypass safeguards or trigger unintended behaviour. In corporate environments where AI tools are integrated with internal knowledge bases, document repositories, or operational systems, such attacks can reveal fragments of training data, internal procedures, or sensitive contextual information. Several demonstrations since 2023 have shown that enterprise AI platforms can inadvertently disclose confidential data when protective guardrails are insufficient.

Closely related are reverse-engineering attacks directed at AI models, APIs, and automated services. By analysing system responses over time, attackers may infer

aspects of underlying architectures, including database structures or decision logic. Once the internal behaviour of a system becomes partially understood, adversaries can design inputs that bypass fraud detection systems, manipulate automated workflows, or exploit hidden vulnerabilities within digital platforms.

Generative AI has also significantly enhanced the effectiveness of social engineering. Threat actors can now produce highly convincing phishing messages by analysing publicly available information, communication styles, and professional networks. Emails, messages, and even voice communications can be crafted to closely resemble legitimate correspondence. In recent incidents, attackers have used AI-generated messages and deepfake voice technologies to impersonate senior executives and authorize fraudulent financial transfers. These techniques dramatically increase the credibility of malicious communications while reducing the effectiveness of traditional detection mechanisms.

Beyond social manipulation, AI is increasingly used to automate technical phases of cyberattacks. Automated tools can rapidly scan networks, identify vulnerabilities, and prioritize exploitable weaknesses far more efficiently than manual methods. Once entry points are identified, AI-assisted frameworks can deploy malware — including ransomware — with a level of precision that significantly increases the likelihood of successful compromise.

Ransomware itself has evolved into a mature criminal ecosystem. The emergence of Ransomware-as-a-Service (RaaS) platforms allows individuals with limited technical expertise to launch sophisticated attacks using tools developed by professional cybercriminal groups. When combined with freely available AI systems, these platforms dramatically lower the barrier to entry for large-scale cybercrime.

Another important development is the increasing use of AI to analyse behavioural patterns within organizations. By examining communication styles, workflows, and operational structures, automated systems can generate messages or instructions that appear consistent with legitimate internal activity. In parallel, AI systems can evaluate network activity and operational schedules to determine when defensive monitoring is likely to be weakest, allowing attackers to time intrusions during periods of reduced oversight.

Fileless malware represents an additional challenge for modern cybersecurity systems. Unlike traditional malicious software that resides on disk, fileless attacks execute directly within system memory and therefore leave fewer detectable signatures. AI-enabled malware can dynamically adapt its behaviour by observing how defensive systems respond to specific actions. This adaptive capability allows malicious code to evade detection tools that rely primarily on known behavioural patterns.

Taken together, these developments illustrate a broader structural shift in cyber conflict. Artificial intelligence is not merely introducing new tools for attackers; it is fundamentally altering the economics of cyber operations. Activities that once required extensive expertise, preparation, and resources can now be conducted more rapidly, at lower cost, and at significantly greater scale.

For financial institutions—whose operations depend on complex and highly interconnected digital infrastructures — this transformation presents a major strategic

challenge. The attack surface continues to expand while adversarial capabilities evolve rapidly. Traditional cybersecurity models, designed for slower and more predictable threat environments, are increasingly strained by the speed and automation enabled by AI-driven attacks.

Yet artificial intelligence represents only one dimension of the emerging threat landscape. The development of large-scale quantum computing systems introduces an additional layer of disruption. While most contemporary cyberattacks rely on classical computing resources, the eventual availability of powerful quantum computers may dramatically expand the capabilities available to adversaries.

Breaking encryption has historically required either mathematical breakthroughs or vast computational resources. Quantum computers, however, operate according to fundamentally different physical principles. Algorithms such as Shor's algorithm demonstrate that sufficiently powerful quantum systems could theoretically break widely used asymmetric encryption schemes, including RSA and elliptic-curve cryptography. Since much of the global financial infrastructure relies on these cryptographic foundations, the implications for long-term data security are profound.

The most concerning scenario involves hybrid attacks that combine classical computing, artificial intelligence, and quantum capabilities. In such models, classical AI systems may conduct reconnaissance, identify vulnerabilities, and coordinate intrusion strategies, while quantum processors perform highly specialized tasks such as cryptographic attacks or advanced optimization calculations. This combination could allow attackers to exploit weaknesses across multiple layers of digital infrastructure simultaneously.

In parallel with these technological developments, cyber conflict is increasingly shaped by geopolitical dynamics. State-sponsored actors often possess access to extensive financial resources, technical expertise, and specialized infrastructure. Their objectives may extend far beyond financial gain to include strategic disruption, intelligence gathering, or economic coercion. As a result, the threat environment facing financial institutions is not limited to conventional cybercrime but includes sophisticated operations conducted by well-funded and highly organized adversaries.

In this environment, the resilience of financial infrastructure depends not only on the strength of individual security tools but also on the architecture of the entire system. Attackers frequently exploit indirect pathways rather than attempting to penetrate the most heavily protected components directly. Weaknesses in software supply chains, cloud services, telecommunications networks, or third-party providers can all serve as entry points into otherwise well-protected environments.

For this reason, modern cybersecurity must be approached as a systemic challenge rather than a purely technical one. Defensive strategies must account for the interaction between technological vulnerabilities, organizational processes, and human behaviour. As artificial intelligence continues to accelerate cyber operations and quantum technologies advance toward practical deployment, the need for resilient, multi-layered security architectures becomes increasingly urgent.

Understanding this evolving threat landscape is therefore essential for financial institutions seeking to protect digital assets, maintain operational continuity, and

preserve trust in an increasingly complex technological environment.

## 19.6 Quantum Escalation and the Future Threat Landscape

Breaking encryption is rarely about direct brute-force attacks. Quantum computers will, however, soon render most current encryption obsolete. Hybrid attacks capable of bypassing PQC are imminent. Two main drivers accelerate encryption-breaking capabilities:

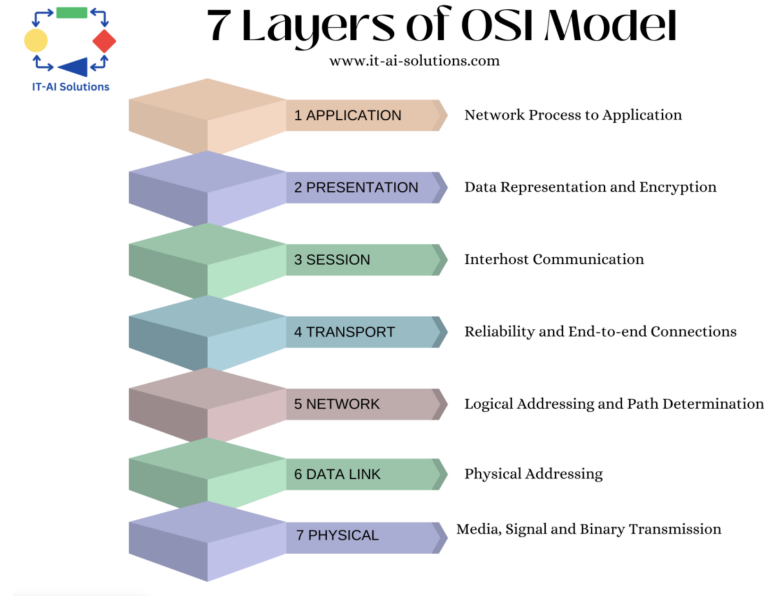
1. Quantum hardware is improving faster than anything else in human history, opening new avenues for algorithms and mathematical breakthroughs.
2. Quantum AI will rapidly surpass human comprehension, improving itself in synergy with classical AI and advanced quantum operating languages such as Qrisp.

**Quintessence:** Systems not already compromised by AI-assisted cyberattacks will likely be broken by 2033, or even by 2029 per IBM’s estimates. The financial sector, especially cryptocurrencies and stablecoins, is the primary target. Larger critical infrastructure attacks, exploiting cascade effects, could disable multiple sectors simultaneously, allowing attackers to extract assets without leaving traces. Disruption could last 10–20 days — far beyond government estimates of three days.

A cyberwar is already underway. Though the precise start date is a matter for historians, internal testing confirms that new classes of military-grade malware, previously unseen by commercial scanners, have been active for years. Our dual-use in-house solutions, designed since December 2022, have proven effective, though more complex and slightly costlier than conventional approaches. Given the value of the assets at risk, these additional costs are negligible. Unfortunately, executives often fail to prioritize cybersecurity at the board level, leaving organizations vulnerable to the best-funded opponents — disaster is the inevitable result.

The only way to avoid chaos, destruction, and loss is to ensure that all seven layers of the OSI model are protected beyond today’s military-grade standards. There is no quick or improvised solution for achieving post-quantum hacking resilience. Current approaches — including PQC — are unable to defend even against advanced classical AI-driven cyberattacks. Data centres protected with so-called military-grade solutions have already been compromised and infiltrated with malware. Contemporary IT defences are no match for deeply trained, highly resourced attackers — whose teams are, in many cases, part of specialized military units.

Securing all seven layers of the OSI model is far from trivial. This was demonstrated by the compromise of the first Chinese quantum satellite, *Mozi* (*Micius*), which was used to deploy the world’s largest QKD network across thousands of kilometres. In 2025, *Mozi/Micius* was compromised because only one OSI layer was protected by QKD, rendering the entire network vulnerable. China has since managed to secure three of the seven layers, significantly increasing the difficulty of penetration — but not eliminating it. There are already credible evaluations demonstrating how QKD



could be defeated even if applied across all seven layers. In other words, *QKD is not a viable solution* — not for defense and certainly not for banking and finance.

Executives must understand that their businesses, assets, transactions, supply chains, intellectual property, and strategic data are already exposed to state-sponsored attackers, even before the arrival of large-scale quantum computers and quantum AI. The disruptive potential of hybrid attacks — where a large, fault-tolerant quantum computer operates in concert with multiple high-performance classical systems — cannot be overstated. Our team assessed these threats in what we call paranoia mode: not optimizing for cost, but for absolute integrity, security, and resilience of every system component and all interconnections. This approach ensures long-term protection measured in decades, not years.

Paradoxically, this strategy proves more cost-effective in the long run than today’s cost-minimization approaches to cybersecurity, which merely postpone failure. Ransomware incidents often represent the least damaging outcomes compared to what is already technically feasible.

Crypto agility is not a solution either. It is slow, expensive, and largely ineffective. Even NIST does not fully trust the three selected PQC algorithms, which is why it recommends classical encryption first, followed by PQC — a resource-intensive and inefficient process. Sector-wide implementation of PQC requires enormous financial and operational effort, and projections indicate that post-quantum keys would need replacement every three to five years. This model is not sustainable.

The good news is that *it is already possible to secure all seven layers of the OSI model impenetrably using technologies available today*, while involving only specific, carefully selected quantum technologies. Our solution is fully compatible with existing internet and network architectures, as it was designed — both in hardware and software — to integrate seamlessly with current infrastructure. It is scalable, industrially robust, and designed for dual use. Moreover, it can be deployed rapidly; implementation speed is limited primarily by administrative and regulatory constraints rather than

technical feasibility.

However, traditional regulatory and administrative approaches must be abandoned. These limitations can be discussed in direct, confidential conversations — but not in publications.

For security reasons, this section remains intentionally incomplete. Our objective is not visibility, but resilience. Publishing sensitive architectural, mathematical, or physical details effectively hands blueprints to attackers. No system — no matter how sophisticated — can survive if its full defensive design is publicly disclosed. Such transparency is not openness; it is self-destruction.

It is critical to understand that *official countermeasures are insufficient and that every sector remains exposed*. Do not rely on the assumption that attackers will focus on short-term objectives such as cryptocurrency theft. These are marginal gains. Strategic objectives come later — and by then, the outcome is already decided.

One final remark on quantum hacking:

One day, you may go to sleep under your country's flag — and wake up under another. That is what is at stake.

Invest, plan, and act accordingly.

## 19.7 Industry Guidance and Security Frameworks

Security solutions must remain compatible with established industry standards in order to ensure interoperability, regulatory compliance, and operational continuity across the financial ecosystem. At the same time, institutions should avoid excessive dependence on a narrow set of dominant technology providers. Platforms that become widely adopted across industries inevitably attract intensive scrutiny from adversaries. Over time, attackers invest significant resources into analysing these systems, reverse-engineering their behaviour, and developing scalable exploitation techniques. What appears to be market dominance can therefore translate into systemic vulnerability.

By the mid-2020s, cybersecurity and institutional resilience have become matters of corporate governance at the highest level. Security can no longer be treated as a purely operational responsibility delegated to IT departments. Strategic accountability and funding authority must reside with executive leadership and the board. Within financial institutions in particular, cybersecurity expenditure should be understood as a form of institutional insurance rather than discretionary cost. Modern financial systems operate under constant pressure from adversarial ecosystems that collectively invest billions of dollars annually in offensive capabilities, frequently pooling resources, intelligence, and infrastructure across national borders.

Architectural design must also move beyond simplistic models of either full centralization or uncontrolled decentralization. Purely centralized infrastructures create high-value single points of failure, while excessive decentralization reduces visibility

and governance. Effective financial infrastructure therefore combines both approaches in a deliberate hybrid architecture. Centralized components enable monitoring, policy enforcement, and coordinated response, while decentralized elements limit the operational impact of a successful intrusion by containing the potential blast radius. When implemented correctly, this hybrid model improves resilience against post-quantum threats, AI-driven attacks, and complex hybrid intrusion scenarios.

Resilient architectures must also account for cascading failures across interconnected technological systems. Financial institutions depend on a wide ecosystem of infrastructure components, including telecommunications networks, cloud services, payment platforms, and external data providers. Disruptions in any one component can propagate rapidly across operational environments. Robust system design therefore requires mechanisms capable of isolating faults, containing disruptions, and maintaining operational continuity even when surrounding infrastructure experiences partial failure.

Regulatory complexity adds another critical dimension to cybersecurity planning. Financial institutions operating across multiple jurisdictions must comply with domestic legal frameworks while simultaneously navigating cross-border regulatory obligations. Data environments are often mirrored or distributed across regions to satisfy legal, operational, or resilience requirements, which can expose organizations to overlapping legal regimes. Effective cybersecurity planning must therefore incorporate legal jurisdiction as an explicit component of system architecture and risk management.

Technological sophistication alone does not guarantee security. Large enterprise platforms — regardless of origin — inevitably accumulate vulnerabilities due to their complexity, rapid development cycles, and extensive integration with third-party components. When critical infrastructure depends heavily on a small number of dominant vendors, the discovery of a single vulnerability can expose thousands of institutions simultaneously. Technology selection must therefore consider not only functionality and market reputation but also security track record, patch management discipline, supply-chain transparency, and long-term maintainability. Vendor dominance should never be mistaken for resilience.

Infrastructure design must also assume that individuals within the organization — including senior leadership — may leave unexpectedly. Systems that depend on concentrated knowledge, singular authority nodes, or exclusive access privileges create structural vulnerabilities. Resilient organizations distribute institutional knowledge, enforce separation of responsibilities, and eliminate single points of operational exposure.

Outsourcing decisions require similarly rigorous scrutiny. Brand recognition and historical success no longer provide reliable indicators of security robustness. Some of the most prominent technology providers have experienced significant breaches in recent years. Fragmented architectures assembled from multiple vendors without systemic integration frequently perform poorly under AI-driven attack conditions, and such weaknesses may become even more pronounced in a future environment that includes quantum-enabled compromise.

Operational experience also remains a critical component of effective defence. Academic expertise provides valuable theoretical insight, but it does not necessarily expose practitioners to real-time cyber conflict characterized by global scale, rapid adversarial adaptation, and coordinated attacks across multiple infrastructures. Responding effectively to such threats requires experience gained through practical engagement with complex operational environments.

Ultimately, the challenge facing financial institutions is not a shortage of technical specialists but the strategic deployment of expertise. Security depends less on the absolute number of professionals employed than on how effectively their knowledge is integrated into governance structures, architectural design, and long-term resilience planning. In an environment increasingly defined by AI-accelerated cyber threats and emerging quantum capabilities, resilient financial infrastructure cannot be achieved through isolated technological measures. It requires coordinated governance, disciplined engineering, and a sustained institutional commitment to security embedded across the entire organization.

## 19.8 Conclusion: The Quantum Imperative for Finance

The financial world is entering an era of unprecedented technological disruption. Quantum technologies, AI-driven attacks, and hybrid threats are no longer theoretical, they are actively reshaping the cybersecurity landscape. Your assets, transactions, and critical infrastructure face risks that are more sophisticated, faster, and more persistent than anything traditional cybersecurity can address.

From a business perspective, cybersecurity is not just a technical issue- it is a strategic, board-level imperative. Investment in post-quantum resilience is an investment in:

- **Asset protection:** Safeguarding digital wealth, cryptocurrencies, and high-value transactions.
- **Operational continuity:** Preventing systemic disruption, supply chain breakdowns, or cascade failures.
- **Reputation and trust:** Ensuring stakeholders and clients have confidence in your ability to manage extreme technological risk.
- **Competitive advantage:** Turning security into a differentiator rather than a cost center.

And here is the key insight: cybersecurity cannot rest in the corner of your organization. It cannot be a “cool back-office function” that you glance at once in a while. It must sit at the very centre of decision-making. If finance wants to continue business safely in the quantum era, cybersecurity is no longer optional — it is the defining driver of operational resilience, trust, and competitive survival.

We are not here to sell fear, we are here to provide clarity, guidance, and actionable solutions. As builders and strategists, we translate complex science into practical, deployable systems that protect your business today and for decades to come.

The financial sector cannot wait for standards to catch up. *The time to act is now.* The investments you make today in post-quantum resilience will determine whether your assets, operations, and reputation survive or whether they become the next headlines in a world reshaped by quantum-powered threats.

*Quantum resilience and central, strategic cybersecurity is no longer optional. It is the defining safeguard of the next generation of finance.*

# Chapter 20

## An Accidental Quantum Love Story

### *A translator's journey into the quantum era*

Orlagh Neary

Language was my first love.

I studied scientific and economic translation and interpreting in college and immersed myself in the mechanics of meaning: how ideas move between cultures, how tone shifts context, how a word in one language rarely maps cleanly onto another. However, in the early 1990s, translating was incredibly tedious. Technical terminology required hours of cross-referencing, and specialized dictionaries weren't that accessible, I think we had a handful of them in our college library. I remember creating my own glossaries, lugging them around and feeling a growing frustration with how much time I spent looking up words instead of thinking about meaning. Somewhere in that frustration, I began to suspect there had to be a better way.

As a teenager, I had taught myself BASIC programming on my beloved ZX Spectrum. I had seen how code could automate repetitive tasks. So, when I was studying on my Erasmus scholarship in Barcelona in my third year at college and discovered computer-assisted translation technologies (early forms of what we would now call applied AI), I fell in love with applied technology. Perhaps it was laziness, or perhaps it was curiosity, but I became obsessed with leveraging technology to free up my time to focus on the intelligent work of translating meaning versus just words.

After graduating college, I carved a path out for myself as a machine translation technology specialist in Ireland's localization industry. I eventually found myself in Redmond, Washington, at Microsoft where I started as a localization engineering group manager. Later, after over a decade serving as Chief of Staff to developer, data and AI teams, my love for language and applied technologies led me to fall in love with the art of storytelling and helping non-tech audiences embrace new technologies.

---

This chapter is a featured contribution to the monograph *A Portrait of Quantum Technologies in Finance*, edited by Oswaldo Zapata, PhD, and published by The Quantum Finance Boardroom (2026). All rights reserved by the chapter's author.

That began my career in marketing and communications.

I was lucky enough to work with Microsoft Research and was involved in the early days of Microsoft AI, which ultimately led to my team being reorganized in 2022 to what was then the newest division in Microsoft, a division focused on the future of the cloud, with technologies ranging from AI for scientific discovery to space and quantum. Call me fickle, but I fell in love yet again, and my love affair with quantum is still going strong.

For someone like me (someone who came from the world of community building and ecosystem engagement, someone whose formal training was in translation, not in physics), it was disorienting in the best possible way. I was suddenly in rooms with people who thought in timescales I'd never considered, who talked about problems that had been unsolvable for decades and were now, slowly, cautiously, beginning to crack open.

With AI, I had fallen in love with the idea of machines helping humans focus on meaning rather than mechanics. With quantum, I found myself captivated by something even more profound: a new way of thinking about complexity itself.

## 20.1 My Love at First Sight Moment

Of all the conversations I sat with in those early days, one has never left me. A scientist began describing the scenarios where quantum computing would have its most profound near-term impact. Drug discovery came up early. The explanation, as I recall, went something like this: nature operates at the quantum level, molecules behave according to quantum mechanical rules. Classical computers are fundamentally limited in their ability to simulate quantum systems accurately, but quantum computers are not. It meant that, for the first time, we may be able to model complex molecules with a level of precision that could accelerate the discovery of new medicines, improve the design of therapies, and potentially change outcomes for people living with diseases that have resisted treatment for generations.

I'm not going to pretend I absorbed every scientific nuance in that room, but I absorbed the meaning of it. When I heard scientists talking about quantum's potential in drug discovery, in climate modeling, in materials science that could change how we power our world, something settled into place for me. I wasn't just learning about a new field. I was learning about a field that could matter to people who would never set foot in a quantum lab and might never hear the word "qubit," but whose lives might be quietly and radically transformed by what happens in those labs.

That is all worth understanding, and it's worth explaining so others can also fall in love with this wonderful world, bring their unique perspectives to it, and accelerate the impact quantum can have.

## 20.2 What AI Taught Us, and What the World of Quantum is Still Learning

When artificial intelligence began its ascent from academic curiosity to world-changing technology, the early years were dominated by a relatively small group of researchers, mathematicians, and computer scientists. But as AI moved from theory into application, from labs into boardrooms, it became clear that the field had a problem. The people building it didn't always reflect the people it would affect. Some of the most important applications took longer to surface, and some of the risks went unexamined for too long.

The lesson was hard-won: advancing transformative technology requires a diverse set of skills, perspectives, and lived experiences, not just technical ones. The people who asked “but what does this mean for patients?” were as important as the people who designed the algorithm. The operators were critical to bring together with the builders.

Quantum computing is at an earlier and more critical juncture. The decisions being made right now, about which applications to prioritize, about how to train the next generation of talent, about what “responsible quantum” even means, will shape the field for decades. And if we make the same mistake of treating it as a purely scientific and technical discipline, accessible only to those with the right academic credentials, we will slow it down and narrow it unnecessarily.

I've spent four years in the quantum ecosystem, and some of the most impactful conversations I've witnessed happen when someone with serious technical depth is in dialogue with someone who brings a different kind of intelligence: business acumen, human empathy, storytelling, systems thinking. The ability to look at a breakthrough and ask not just “how does this work?” but “who and what does this serve, what value does this create or extract?” These are questions that help shape the path of innovation from the lab to adoption at scale.

## 20.3 An Invitation

If you are reading this and you are not a physicist, if your background is in finance, or healthcare, or law, or communications, or policy, or community development, or any of the thousands of other disciplines that make a complex world function, I want you to hear this clearly. Quantum computing needs you. The field is at the stage where the bridges between science, business and society need to be built deliberately and well. That work requires people who understand how industries operate, how communities are shaped, how audiences engage, how trust is established and broken.

You don't need to understand the physics to contribute meaningfully. You need to be curious, willing to learn, and genuinely motivated by the possibility that this technology, handled well, developed thoughtfully, communicated honestly, could be one of the most powerful forces for good that our generation gets to be part of.

I fell in love with this field on Valentine’s Day, four years ago, in a room full of scientists who were thinking about things I didn’t yet have words for. You don’t need to wait for a reorganization to find your way in. You just need to decide that you belong in the conversation. Because you do!

## 20.4 So Where Do You Start?

The beautiful thing about quantum being early is that the paths in are still being carved. You need a little bit of courage, a lot of curiosity, and a willingness to sit in rooms where you don’t yet speak the language.

If you’re in business or finance, start asking questions about quantum’s relevance to your industry. Attend industry conferences that are beginning to include quantum tracks. Organizations like the Quantum Economic Development Consortium (QED-C) are actively building bridges between quantum researchers and industry applications. Your ability to articulate return on investment, de-risk early adoption, or identify strategic timing matters as much as the hardware itself.

If you’re in policy or government, countries around the world are making significant investments in quantum research through initiatives like the U.S. National Quantum Initiative and the EU Quantum Flagship. These programs need people who understand regulatory frameworks, workforce development, international cooperation, and responsible innovation. The decisions being made now about standards, export controls, education infrastructure, and public-private partnerships will shape the field for decades.

If you’re a communicator, writer, or educator, the quantum field desperately needs people who can explain without oversimplifying and inspire without hyping. Write about it. Make videos. Develop curricula. Start a newsletter. Create content that meets people where they are. I wrote a series in the summer to help explain quantum using golf metaphors – Quantum Links. Come play a round! The scientists are brilliant at the physics and many of them would welcome partners who are brilliant at translation.

If you’re in healthcare, climate, or social impact, learn about the applications that align with your mission. Quantum’s potential in drug discovery, materials science, climate modeling, and optimization could transform these fields. Your domain expertise is essential for identifying which problems quantum should be trying to solve and how to evaluate whether the solutions are actually serving people. Join interdisciplinary working groups. Push for collaborations between quantum researchers and the communities who would benefit most.

If you’re simply curious, take a course. MIT offers a quantum strategy course for business leaders, and there are a growing number of free courses online. Join communities focused on quantum literacy and ecosystem development. Follow quantum companies, research labs, and thought leaders. Read, ask questions.

Most importantly, don’t wait to feel fully qualified. The field is moving fast, and the people shaping its direction right now are the ones who showed up before they felt

ready. Quantum doesn't need more gatekeeping. It needs more translators, more connectors, more people willing to bridge the gap between what's possible in the lab and what's needed in the world.

I came into this space as a translator who loved language, fell in love with applied AI, and discovered that quantum needed exactly that kind of love too. Your path will be different but the door is open, all you have to do is open your heart and walk through the door.



# Chapter 21

## Where Quantum Computing Meets Optimization

*Promise, practice, and perspective*

Pascal Halfmann

### 21.1 Introduction

*Quantum Computing* has spent much of the last decade oscillating between two extremes in public perception. On the one hand, it is presented as a revolutionary technology that will upend entire industries and render today's computers obsolete. On the other hand, it is dismissed as an academic curiosity – impressive in principle, but too fragile, too small-scale, or too far from real-world relevance to matter in practice. Both perspectives miss something essential. At its core, quantum computing is neither magic nor myth. It is a new computational paradigm, shaped by the laws of quantum mechanics, that offers fundamentally different ways of representing and exploring complex decision problems. Whether this difference translates into practical value depends not on the novelty of the physics, but on the structure of the problems we ask it to solve – and on the expectations we bring to it.

In this essay, I approach quantum computing from the perspective of optimization. Optimization problems arise whenever limited resources must be allocated under competing objectives and constraints – problems that are ubiquitous in industry and society, and notoriously hard. Decades of research in operations research and

---

This chapter is a featured contribution to the monograph *A Portrait of Quantum Technologies in Finance*, edited by Oswaldo Zapata, PhD, and published by The Quantum Finance Boardroom (2026). All rights reserved by the chapter's author.

mathematical optimization have produced powerful classical methods, but also a clear understanding of their limits. Quantum computing in optimization, here referred to as Quantum Optimization, promises to offer alternative ways to search, sample, and reason about large and highly structured decision spaces. Rather than assuming that quantum optimization will outperform classical methods, I focus on how it changes the way complex decisions can be modeled and explored. This emphasizes structure, trade-offs, and practical constraints over raw computational power. The experiences described here are drawn from concrete projects in which quantum methods were embedded in real optimization workflows. They reflect neither uncritical enthusiasm nor outright skepticism, but the perspective of a researcher seeking to understand where quantum computing already adds value, where it does not, and what would need to change for its impact to grow. This pragmatic viewpoint inevitably leads to nuanced conclusions: some encouraging, some sobering, and many conditional.

My academic background is rooted in mathematical optimization. I studied Business Mathematics at RPTU Kaiserslautern, Germany, where I early on focused on optimization methods. I continued along this path during my doctoral studies at RPTU, which I completed with *summa cum laude*, Halfmann, 2021. My PhD research focused on multiobjective optimization, which studies decision problems involving several competing goals that cannot be reduced to a single notion of optimality. Such problems are particularly challenging because improving one objective often comes at the expense of another, leading not to a single optimal solution but to entire sets of trade-offs that must be understood and navigated. It is fair to say that I am a proven expert in mathematical optimization, which is underlined by numerous scientific publications, regular contributions to international conferences, and invited talks<sup>1</sup>. At Fraunhofer ITWM, I deliberately expanded this academic expertise not only towards data driven approaches in AI and machine learning but in particular towards industrial application. Much of this applied work has been situated in the domains of finance and energy, where I also acquired substantial domain knowledge. Together, these experiences have shaped my perspective as an optimization expert who evaluates new computational approaches not only by their theoretical promise, but by their ability to function in practice.

When I first encountered quantum computing after joining Fraunhofer ITWM in 2021, I was – by my own admission – a complete outsider to the field. Apart from two introductory semesters of physics during my studies and a general interest in physics, I had no prior exposure to quantum mechanics or even quantum computer. Entering the field therefore did not feel like a natural extension of my previous work, but rather like stepping into a new intellectual territory with its own language, assumptions, and ways of thinking. This initial distance, however, turned out to be an advantage. Approaching quantum computing without a background in physics or computer architecture, I engaged with it primarily through the lens of optimization: What kinds of problems does it naturally represent? Where do classical methods struggle? And under what conditions could quantum approaches offer something genuinely different? This outsider perspective shaped both my learning process and my focus and it has been appreciated at Fraunhofer ITWM: after several positions as project manager of research and industry project, I became the research coordinator

---

<sup>1</sup>See my Google Scholar profile at <https://scholar.google.de/citations?user=TXEBfIUAAAAJ>.

for our quantum computing activities in the institute. I never saw myself proving the next quantum phenomenon or building novel hardware. Instead, I became interested in a narrower but ultimately more practical question: when, and for which problem types, might quantum computing become useful at all? To address this, my work focuses on identifying and translating real-world problems that are suitable for quantum algorithms, as well as on developing new quantum algorithmic approaches. Figuratively speaking, if quantum computing is a gold rush, I am not digging for gold – I see myself as a goldsmith, turning raw potential into real-world value.

To build fluency in quantum computing, I combined structured learning with hands-on experimentation. My foundation was established through participation in the IBM Qiskit Summer School, which provided a systematic introduction to the fundamentals of the field, standard quantum algorithms, and quantum hardware through lectures and guided coding exercises. I complemented this coursework with intensive reading of core textbooks, most notably *Quantum Computation and Quantum Information: 10th Anniversary Edition* by Nielsen and Chuang (2012).

At the time, however, structured learning opportunities were either scarce or not easily accessible to me. As a result, much of my further immersion was self-directed and practice-oriented: studying survey and technical papers on arXiv – primarily in the area of quantum optimization – and implementing and testing quantum algorithms using Qiskit. Based on these experiences, I recommend the following resources for readers seeking an entry point into the field:

- the seminal textbook *Quantum Computation and Quantum Information: 10th Anniversary Edition* by Nielsen and Chuang (2012),
- the book *Quantum Computing for Everyone* by Bernhardt (2019),
- the lecture notes by John Preskill at Caltech: <https://www.preskill.caltech.edu/ph229/>,
- courses such as the IBM Qiskit Summer School and other IBM Qiskit resources, as well as offerings from Coursera<sup>2</sup>, Udemy<sup>3</sup>, or universities such as MIT,
- and dedicated courses offered by Fraunhofer ITWM<sup>4</sup>.

In addition, I am currently contributing to the book “Combinatorial Optimization Using Quantum Computing – A Gentle Introduction”, which aims to lower the entry barrier for researchers and practitioners with a background in optimization.

The remainder of this essay is organized as follows. I first present two concrete success stories in quantum optimization drawn from my own work and close collaboration with industrial and academic partners. The first focuses on optimization problems in the energy industry, while the second addresses multiobjective optimization in the context of portfolio management. In each case, I describe the problem setting,

<sup>2</sup><https://www.coursera.org/search?query=quantum%20computing>

<sup>3</sup><https://www.udemy.com/courses/search/?src=ukw&q=quantum%20computing>

<sup>4</sup>[https://www.itwm.fraunhofer.de/en/fields-of-application/quantum-computing/quantencomputingschulungen\\_ueberblick-en.html](https://www.itwm.fraunhofer.de/en/fields-of-application/quantum-computing/quantencomputingschulungen_ueberblick-en.html)

the motivation for exploring quantum approaches, the methods employed, and the insights gained. I then conclude with a broader assessment of the current state of quantum computing from an optimization perspective, outlining open challenges, realistic expectations, and directions for future development.

## 21.2 Success Story: Quantum Optimization in the Energy Industry

### 21.2.1 Problem context and classical limits

The transformation of the energy system is one of the central technological and societal challenges of our time. The transition toward renewable energy sources such as wind and solar power is essential for reducing greenhouse gas emissions and mitigating climate change. At the same time, it also reduces dependence on fossil fuels and on geopolitical actors controlling their supply. While these developments are broadly desirable, they fundamentally change how energy systems must be planned and operated. Unlike conventional power plants, many renewable energy sources are inherently volatile. Their availability depends strongly on weather conditions and time of day, while energy demand – particularly for electricity – must be satisfied continuously and reliably. As a result, modern energy systems must cope with fluctuating supply, uncertain forecasts, and increasing decentralization. This requires not only new physical infrastructure, such as upgraded grids and energy storage systems, but also more sophisticated methods for steering, planning, and coordinating energy production and consumption.

At the heart of many of these challenges lies a class of large-scale optimization problems that govern operational decision-making in power systems. One of the most prominent examples is the unit commitment problem (UCP). In simplified terms, unit commitment determines which power generation units should be switched on or off over a given planning horizon, and at what times, in order to meet predicted electricity demand at minimal cost while respecting a wide range of technical and regulatory constraints, see Figure 1. These constraints include minimum up- and down-times of power plants (time a power plant has to stay on/off when turned on/off), ramping limits (maximal change in energy production in a certain time window), maintenance schedules, power grid restrictions, and increasingly the integration of renewable generation and storage. Decisions made at this level have a cascading effect on system stability, costs, and emissions and are politically sensitive: blackouts would lead to societal instabilities.

From an optimization perspective, the unit commitment problem is particularly challenging. It combines discrete decisions – such as whether a power plant is running or not – with continuous variables describing power output, and it must be solved repeatedly under uncertainty and tight time constraints. Classical optimization methods have been highly successful in addressing such problems and are an integral part of today’s energy markets. Small- to medium-sized problem instances with simplified constraints can be solved in under five minutes using state-of-the-art

commercial solvers such as Gurobi. However, as system size grows, renewable penetration increases, and additional objectives and constraints such as robustness and sustainability are introduced, these methods approach practical limits in terms of scalability, flexibility, and solution diversity. It is in this context that alternative computational approaches, including quantum optimization, have attracted interest – not because classical methods have failed, but because the evolving structure of modern energy systems places growing demands on existing optimization methods: there is a pressure to provide better solutions to more complex problems faster. This raises the question of whether new approaches could complement established tools in practice.




Figure 21.1: Example production schedule obtained from solving a unit commitment problem. Based on a forecast of electricity demand over a 24-hour horizon, generation is allocated across different types of power plants while respecting production costs and technical constraints. The schedule combines renewable sources (run-of-river hydropower, wind, and photovoltaic solar), baseload generation (nuclear), and dispatchable thermal units (lignite, coal, and gas-fired combined-cycle (CC) and gas-turbine (GT) plants). Such schedules illustrate the complexity of modern energy systems, where economic efficiency, technical feasibility, and variability of renewable generation must be balanced simultaneously.

### 21.2.2 Quantum approach and results

The work presented in this section is based on results obtained within the EnerQuant project, funded by the German Federal Ministry for Economic Affairs and Climate Action. The project investigated the applicability of quantum optimization methods to decision problems in the energy sector, with a focus on embedding quantum approaches into realistic optimization workflows motivated by unit commitment-type

problems<sup>5</sup>. A detailed project report is available in German, see Halfmann, Lenk, et al., 2024.

At an early stage of the project, a deliberate decision was made to focus on quantum annealing rather than gate-based quantum computers. At the time (2021), only a small number of gate-based processors with a few dozen qubits were publicly accessible, limiting experiments to toy-sized instances and yielding results that were difficult to interpret due to hardware noise and qubit variability. Quantum annealing, by contrast, offered a more mature and immediately usable platform for combinatorial optimization, as exemplified by commercially available systems from D-Wave Systems Johnson et al., 2011. Conceptually, quantum annealing can be viewed as a quantum analogue of simulated annealing, using quantum effects to explore complex energy landscapes. All experiments were conducted using the standalone quantum annealer rather than vendor-provided hybrid solvers, as the latter obscure the relative contribution of quantum and classical computation, whereas focusing on pure annealing allowed for a clearer assessment of the strengths and limitations of quantum optimization. In practice, the fixed optimization mechanism of quantum annealers shifted the project’s focus away from low level algorithm design toward modeling, hybrid integration, and result analysis – areas closely aligned with my expertise in mathematical optimization.

**Finding a suitable model:** Classical optimization problems are typically formulated using an objective function, decision variables of various types, and a set of constraints capturing technical and operational requirements. Quantum annealers impose a much stricter interface: problems must be expressed as a Quadratic Unconstrained Binary Optimization (QUBO) problem. A QUBO represents the objective as a quadratic function of binary variables only, with constraints incorporated indirectly through penalty terms added to the objective.

Generic transformation techniques exist to convert constrained mixed-integer formulations into a QUBO. These methods binarize non-binary variables, reduce higher-order terms to quadratic form using auxiliary variables, and encode constraints via penalty terms. While systematic, this approach comes at a significant cost: it increases the number of variables – and thus required qubits – and typically produces dense interactions between variables. Because current annealers do not provide all-to-all connectivity, missing couplings must be realized through minor embedding, which introduces additional auxiliary qubits and often leads to precision issues once penalties are applied.

In the EnerQuant project, this motivated a different modeling philosophy: if the final formulation must be quadratic, then the starting point should already be as close to quadratic as possible. Rather than formulating the unit commitment problem conventionally and applying a generic reduction, we developed compact quadratic constrained formulations that preserve problem structure and minimize auxiliary variables. In particular, we introduced novel formulations for constraints such as minimum up-time and minimum down-time requirements, see Halfmann et al., 2023. In the resulting model, auxiliary variables were required only for the discretization of

---

<sup>5</sup><https://www.itwm.fraunhofer.de/en/departments/analytics-computing-en/quantum-computing/enerquant-quantencomputing-power-industry.html>

generation output, while the commitment logic and operational constraints remained compact. Compared to generic transformations, this yielded a significantly sparser QUBO with roughly 80% fewer couplings, enabling larger problem instances to be mapped to the annealer.

Despite these modeling improvements, solution quality on quantum hardware remained a major challenge. Even for small instances – such as 5 generating units over 12 time steps – classical solvers like Gurobi were orders of magnitude faster and returned provably optimal, feasible solutions. In contrast, raw annealing runs often produced solutions with inferior objective values and constraint violations requiring post-processing. This gap becomes even more apparent when contrasted with realistic operational scales, where unit commitment problems are solved over 24-hour horizons with up to 96 time steps and involve dozens to hundreds of generation units.

**Hybrid workflows:** These observations led to a shift in focus from encoding the problem to improving solution quality. In particular, attention centered on penalty factors used to enforce constraints in QUBO formulations. If penalty values are chosen too small, infeasible solutions dominate; if chosen too large, feasibility is enforced at the expense of meaningful objective optimization. Identifying an appropriate balance is therefore critical, yet difficult to determine a priori – even for compact formulations.

To address this issue, we developed an adaptive penalty calibration strategy that operates during the solution process. Instead of fixing all penalties upfront, the method iteratively constructs the full QUBO by introducing constraint classes step by step. Penalty values are tuned empirically based on observed feasibility and solution quality, with hard technical constraints enforced first and softer requirements, such as demand satisfaction with allowable deviations, incorporated later.

Using this approach, experiments were rerun on more mature annealing hardware provided by DWave. The improvements were substantial: the number of violated constraints dropped from dozens per solution to only two or three in typical runs, and achieved objective values were within a factor of approximately 1.5–2 of the classical optimum, placing the approach in the range of competitive classical heuristics.

In terms of runtime, classical solvers such as Gurobi continued to outperform quantum annealing by orders of magnitude. However, as problem size increased, the performance gap narrowed compared to the deterministic baseline, and first indications of more favorable scaling behavior for annealing were observed. Achieving a practical break-even point would nevertheless require substantially larger problem instances, and thus more qubits, than are available on current hardware.

**Classical postprocessing to incorporate uncertainties in the energy demand:** For deterministic quadratic or linear mixed-integer problems such as unit commitment, these results confirm that classical solvers set an exceptionally high benchmark. However, deterministic formulations do not fully reflect operational reality. Electricity demand and supply are inherently uncertain due to factors such as renewable generation variability and unexpected outages, motivating the use of robust optimization. In robust optimization, solutions are evaluated based on their performance across multiple scenarios, with robustness defined as maintaining

feasibility or requiring only minor adaptation under uncertainty. Classical robust formulations typically increase computational complexity significantly due to scenario expansion.

In this context, we explored whether the inherent stochasticity of quantum annealing could be exploited constructively, see Halfmann, Trebing, and Lenk, 2024. Rather than returning a single solution, the annealer produces a distribution over many candidate solutions across repeated runs. We used this property to generate diverse near-optimal solutions for a deterministic model (e.g., based on mean demand) and evaluated them post hoc with respect to a robust objective. Each solution was assessed across multiple demand and supply scenarios and aggregated using a robustness criterion such as worst-case cost.

The resulting solutions exhibited good robustness properties, with objective values within a factor of roughly 1.5 of the classical robust optimum on tested instances. While Gurobi remained faster overall, the runtime gap was noticeably smaller than in the deterministic setting. Importantly, these results did not constitute a quantum advantage in the strict sense; the observed benefits arose from combining quantum sampling with classical evaluation rather than from raw quantum speedup.

### 21.2.3 Lessons learned and future impact

The EnerQuant project marked my first in-depth engagement with quantum computing and served as my entry point into the field. Coming from a background in classical optimization, the learning curve was steep, but also highly instructive. Working through the full pipeline – from modeling and hardware constraints to algorithmic behavior and result interpretation – made it possible to assess quantum optimization not as an abstract concept, but as a practical computational tool. Several overarching lessons emerged from this experience:

**Using quantum algorithms effectively is non-trivial.** Successful application requires careful attention to many interacting components, including problem formulation, parameter choices, and algorithmic tuning. In the case of quantum annealing, penalty factors critically influence feasibility and solution quality; for gate-based approaches, variational parameters play a similarly decisive role. These aspects cannot be treated as technical details, but are central to performance.

**Quantum methods are most promising in hybrid settings.** Across the project, the most meaningful improvements arose when quantum algorithms were embedded into classical workflows. Combining quantum sampling with classical preprocessing, post-processing, and evaluation allowed the strengths of both paradigms to be leveraged. The adaptive penalty calibration strategy developed in EnerQuant is a concrete example of such a hybrid approach.

**For standard optimization problems, there is currently no quantum advantage.** For well-structured, deterministic mixed-integer problems such as unit commitment, classical solvers like Gurobi set an extremely high benchmark. Our experiments confirm observations reported widely in the literature: under these conditions, quantum optimization does not yet outperform classical state-of-the-art

methods in terms of runtime or solution guarantees.

**Robust optimization may offer a more promising entry point.** Problems involving uncertainty and robustness are not only of high practical relevance in energy systems, but are also significantly more challenging for classical solvers than their deterministic counterparts. Here, the inherent stochasticity and solution diversity produced by quantum algorithms may become an asset rather than a limitation, suggesting a potential “can-opener” for future quantum advantage.

## 21.3 Success Story: Quantum Multiobjective Optimization for Portfolio Optimization

### 21.3.1 Problem context and classical limits

In the previous section, we have seen that for well-structured, deterministic optimization problems, state-of-the-art classical solvers are difficult to outperform. Even for large and complex instances, commercial solvers such as Gurobi routinely deliver provably optimal solutions within short runtimes. This sets a high benchmark for quantum optimization approaches that aim to compete directly on such problems. However, this picture changes when optimization problems move beyond a single, well-defined objective. In many practical settings, decision-makers are faced with competing goals that cannot be meaningfully reduced to a single performance measure. This is true in your decisions (choosing a car that is fast and cheap) and even more so in companies’ decisions: instead of optimizing one objective, several conflicting objectives must be considered simultaneously.

For such a *Multiobjective Optimization Problem*, there is typically no unique optimal solution. Instead, one seeks a set of Pareto-optimal solutions, each representing a different trade-off between the objectives. Improving one objective inevitably comes at the expense of at least one other. From a computational perspective, the task is therefore not only to find a single solution, but to explore and characterize an entire solution set. While classical methods for multiobjective optimization exist, exact methods often rely on repeated scalarization, combining objectives into a single objective and solving the resulting problem multiple times, which becomes increasingly cumbersome to steer and manage. Heuristics are often evolutionary-based with fixed populations. In total, a comprehensive solver as for single objective problems does not exist yet, leaving room for novel approaches and technologies.

These characteristics suggest a potential role for quantum optimization. Rather than excelling at producing one optimal solution, quantum algorithms naturally generate distributions over many candidate solutions they sample from. This makes them conceptually well suited for exploratory tasks in which understanding trade-offs and alternative compromises is more important than identifying a single optimum.

Portfolio optimization is a canonical example of a multiobjective decision problem. Investors must balance expected return against risk, while often also accounting for additional criteria such as liquidity, transaction costs, or sustainability considerations.

The resulting trade-offs are central to real-world decision-making and align naturally with the structure of multiobjective optimization problems. For this reason, portfolio optimization provides a natural testbed for exploring whether quantum optimization can offer value in navigating complex trade-off spaces.

### 21.3.2 Quantum approach and results

In this section, I turn to quantum approaches for multiobjective optimization, focusing on portfolio optimization as a representative and practically relevant case. The work presented here draws on results from several sources, including the QuSAA project, funded by the German Federal Ministry of Education and Research<sup>6</sup>. Within this project, we investigated how strategic asset allocation problems can be addressed using quantum optimization methods.




Figure 21.2: Multiobjective solution set for a portfolio optimization problem. Each point represents a feasible portfolio evaluated with respect to two competing objectives (e.g., risk and return). In contrast to single-objective optimization, optimality here does not refer to one best solution, but to Pareto optimality: a portfolio is considered optimal if no other portfolio improves one objective without worsening at least one other. The upper-left envelope of the cloud of points resembles the Pareto front and illustrates the trade-offs available to a decision-maker. Additionally, the color scale encodes the Solvency Capital Requirement (SCR), a regulatory risk measure, highlighting how additional constraints and objectives can be analyzed simultaneously within the same solution set. A reference portfolio is shown for comparison.

**Modeling hard portfolio optimization problems:** The classical starting point for portfolio optimization is the Markowitz mean–variance model, which seeks to balance expected return against risk by minimizing portfolio variance subject to a

<sup>6</sup><https://www.itwm.fraunhofer.de/en/departments/analytics-computing-en/quantum-computing/quantum-algorithms-strategic-asset-allocation.html>

budget constraint. Because the objective function is quadratic and the constraints are simple, this problem has become a popular testbed for quantum optimization algorithms. Under additional simplifying assumptions – such as allowing each asset to be selected at most once – the model can be transformed relatively easily into a QUBO formulation suitable for quantum annealing. What is often overlooked, however, is that the Markowitz problem is convex and can be solved efficiently and reliably using classical optimization methods. As such, it offers little room for quantum methods to demonstrate a meaningful advantage.

In practical decision-making, however, the Markowitz model is far too simplistic. Real-world portfolio optimization must account for additional objectives and constraints, such as transaction costs, regulatory requirements, and risk measures mandated by financial authorities. One particularly challenging example is the Solvency Capital Requirement (SCR), which plays a central role in insurance and financial regulation. The SCR does not admit a simple explicit formula, and even commonly used approximations are nonlinear and non-quadratic, making them difficult to integrate into QUBO-based formulations. To address this, we adopted a data-driven approach: using real-world data, we trained a machine learning model to approximate the SCR with a quadratic function. This enabled us to construct an enhanced Markowitz-type model that could be expressed as a QUBO and solved using a quantum annealer, see Turkalj et al., 2024. As in the energy use case, classical solvers outperformed the annealer in terms of runtime and solution optimality. Nevertheless, the quantum approach exhibited a higher diversity of near-optimal solutions, which becomes relevant when the problem is viewed from a multiobjective perspective rather than as a single-objective optimization task.

**Hybrid quantum algorithm for multiobjective optimization:** Building on this observation, we next considered portfolio optimization explicitly as a multiobjective problem, in which several objectives are optimized simultaneously rather than combined into a single weighted sum. In this context, researchers have proposed quantum algorithms based on the Quantum Approximate Optimization Algorithm (QAOA). QAOA is a hybrid quantum–classical algorithm that alternates between two types of operations: one that encodes the objective function and one that mixes between feasible solutions. These operations are parameterized, and a classical optimization loop is used to tune the parameters so as to optimize the desired objective.

In the multiobjective setting, this idea is extended by Ekstrom et al. (2025) via introducing separate objective-encoding components for each objective function within the quantum circuit. We followed up on this line of work and enhanced the proposed approach in several ways, including improvements to the circuit ansatz, the choice of mixing operations, and – most importantly – the classical optimization layer, see Turkalj et al., 2025. The original approach relied on a scalar aggregation of objectives within the classical optimizer, which proved computationally expensive and limited solution quality. We replaced this component with a genetic algorithm that directly operates on the multiobjective values of the sampled solutions. By doing so, the classical optimizer no longer required an explicit aggregation of objectives, but instead evolved parameter settings based on Pareto dominance and diversity. In our experiments, this led to a noticeable improvement in the quality and spread of

the resulting solution sets.

**Using a quantum circuit as a sampling device:** Despite these improvements, no quantum advantage over classical multiobjective optimization methods was observed. However, recent work points toward a promising alternative perspective. Instead of viewing quantum algorithms as solvers that must outperform classical methods on a single optimal solution, they can be interpreted as sampling devices for structured solution spaces. In this approach, a standard QAOA circuit is applied to a scalarized version of the multiobjective problem, but with randomly chosen objective weights. By repeatedly executing the circuit with different weights, the quantum algorithm generates a diverse set of solutions that approximate the Pareto front.

his idea closely mirrors our earlier use of quantum annealing to explore robust solution spaces. In the referenced work by Kotil et al. (2025), the approach was evaluated on the MaxCut problem – a quadratic combinatorial optimization problem structurally similar to portfolio optimization – and the results were encouraging. The quantum method produced Pareto-optimal solutions with solution quality and runtimes close to those of classical exact algorithms. While this does not yet constitute quantum advantage, it reinforces the view that quantum optimization may be particularly well suited for exploratory tasks, where generating diverse, high-quality candidate solutions is more important than identifying a single provably optimal one.

### 21.3.3 Lessons learned and future impact

The results of this success story reaffirm a central conclusion from the previous section: for standard optimization problems – especially simplified or artificially constructed benchmarks – quantum optimization is unlikely to demonstrate a clear advantage in the near future. Classical solvers are highly optimized for such settings, and transforming these problems into quantum-compatible formulations does not change their fundamental computational structure. The insights gained from portfolio optimization therefore point less toward near-term speedups and more toward identifying problem classes where quantum methods may offer complementary value.

**Artificial or overly simplified optimization problems are poor indicators of quantum utility.** Canonical models such as the classical Markowitz portfolio problem are attractive as benchmarks due to their mathematical structure, but they are convex and easily solvable using classical methods. Success on such problems provides limited insight into the practical potential of quantum optimization.

**Multiobjective optimization exhibits a problem structure well suited to quantum approaches.** Similar to robust optimization, multiobjective problems shift the focus from finding a single optimal solution to exploring trade-offs among competing objectives. This exploratory nature aligns naturally with quantum algorithms, which tend to generate distributions of near-optimal solutions rather than isolated optima.

**Hybrid workflows that use quantum computing as a sampling device are particularly promising.** Treating quantum algorithms as generators of diverse candidate solutions – embedded within classical preprocessing, evaluation, and

decision-support pipelines – proved to be a more effective paradigm than viewing them as stand-alone solvers. In this role, quantum methods can enhance the exploration of Pareto fronts and support decision-making under uncertainty.

Taken together, these findings suggest that the value of quantum optimization in portfolio management is unlikely to arise from outperforming classical solvers on traditional benchmarks. Instead, its potential lies in supporting complex, multiobjective decision processes where diversity, structure, and robustness are more important than exact optimality.

## 21.4 State of Quantum Computing & and a Look Ahead

### 21.4.1 Do we already have quantum advantage?

In short: no.

Looking at quantum computing as a whole, there is currently no convincing and broadly accepted demonstration of quantum advantage for problems of practical relevance. While individual experiments have claimed forms of advantage under carefully controlled conditions, these results have either relied on highly specialized problem constructions or have not translated into utility for real-world applications.

A recent example illustrates this well. In 2024–2025, D-Wave Systems reported the observation of quantum states and algorithmic behavior that were claimed to be hard to simulate classically, King et al. (2025). While scientifically interesting, these claims were rapidly scrutinized and subsequently challenged by independent groups, who demonstrated classical simulation methods that closed much of the purported gap, e.g., by Mauron and Carleo (2025). This episode reflects a broader pattern in the field: apparent breakthroughs often trigger rapid methodological advances on the classical side, underscoring how difficult it is to establish a robust and uncontested notion of quantum advantage. At the same time, it would be misleading to interpret the absence of quantum advantage as stagnation. On the contrary, progress in quantum computing has been substantial. Hardware platforms have evolved steadily, with increases in qubit counts, improvements in coherence times, and better control and calibration techniques. In parallel, algorithmic research has matured, particularly in the development of hybrid quantum–classical methods, error mitigation strategies, and application-oriented workflows tailored to near-term devices. The overall pace of development remains impressive.

Nevertheless, from an applied perspective, these advances have not yet translated into clear quantum utility. For the types of problems encountered in industrial optimization, finance, or energy systems, classical computing continues to dominate decisively. Even where quantum devices can be used meaningfully, they typically do so as components within hybrid pipelines, without delivering standalone performance benefits that would justify replacing classical methods.

This assessment becomes even clearer when focusing specifically on optimization. As the case studies in this essay have shown, state-of-the-art classical solvers set an extraordinarily high bar for deterministic optimization problems. Quantum optimization methods have so far not matched classical approaches in terms of speed, scalability, or solution guarantees. Where they do exhibit distinctive behavior – such as generating diverse candidate solutions or supporting exploratory analyses – this behavior complements rather than surpasses classical optimization.

In this sense, the current state of quantum optimization is best understood not as a failure to achieve advantage, but as a clarification of where such advantage is unlikely to arise. Competing head-on with classical solvers on standard optimization benchmarks is improbable in the near future. Any meaningful impact of quantum computing will instead have to come from problem classes and decision contexts that emphasize uncertainty, robustness, and exploration over exact optimality.

## 21.4.2 Future challenges and necessary developments

If quantum computing is to move beyond experimental demonstrations and become a meaningful component of applied computation, several challenges must be addressed in parallel. These challenges span hardware, algorithms, software engineering, system integration, and broader strategic considerations.

**Hardware maturity beyond the NISQ era:** Most current quantum devices belong to the so-called *Noisy Intermediate-Scale Quantum (NISQ)* era. NISQ systems are characterized by tens to hundreds of physical qubits, limited coherence times, and the absence of full error correction. While they enable valuable experimentation, their noise levels and scale fundamentally restrict achievable performance. Current roadmaps across multiple hardware platforms are now targeting the transition toward *Fault-Tolerant Quantum Computing (FTQC)*, in which logical qubits are protected by quantum error correction. Early demonstrations of small numbers of logical qubits mark an important milestone, but to challenge classical hardware in practice, substantially more progress is required. In particular, scaling the number of logical qubits must go hand in hand with long coherence times, high-fidelity operations, sufficient qubit connectivity, and – often overlooked but crucial for practical workloads – fast gate execution times. Without advances across all of these dimensions, quantum devices will struggle to support large-scale, time-critical applications.

**Hybrid QC–HPC integration as a first-class challenge:** Quantum computers alone will not replace classical high-performance computing. Many computational tasks will remain more efficiently solvable on classical architectures, and the most realistic near- and mid-term vision is one of hybrid computation. In such settings, quantum devices act as accelerators embedded within classical workflows. Achieving this requires deep integration between quantum systems and classical HPC environments, including scheduling, data movement, and workflow orchestration. Initial efforts in this direction are already underway at large supercomputing centers such as Leibniz Supercomputing Centre, but the challenge extends beyond national facilities. Industrial users operating private HPC clusters may wish to integrate smaller or specialized quantum devices into their existing infrastructure. Developing

scalable, robust, and portable integration solutions will therefore be essential. A related bottleneck arises in data-intensive applications such as machine learning: many algorithms require large training datasets, yet transferring classical data into quantum representations is costly and nontrivial. Addressing this data interface problem could become a key enabler for quantum machine learning.

**Algorithmic development beyond NISQ:** On the algorithmic side, much of today’s research understandably focuses on NISQ-compatible methods. However, if fault-tolerant devices become available, entirely different algorithmic regimes will open up. Designing algorithms explicitly for FTQC – rather than retrofitting NISQ-era techniques – will be critical. First approaches in this direction have begun to emerge, but a comprehensive algorithmic toolbox is still lacking. Across optimization, simulation, and machine learning, a major open challenge lies in classical–quantum co-design: developing algorithms in which classical and quantum components are conceived jointly, rather than treated as loosely coupled modules. For quantum optimization in particular, this means moving beyond isolated quantum subroutines toward hybrid strategies that exploit classical strengths in modeling and feasibility while using quantum resources for structured exploration or sampling.

**Usability and software engineering:** Beyond hardware and algorithms, usability represents a metalevel challenge that is often underestimated. For quantum computing to be adopted widely, quantum devices must integrate seamlessly with existing software ecosystems and computational workflows. This requires standardized interfaces, robust middleware, and tooling that allows quantum hardware to interact reliably with other devices and with human users. Importantly, most future users of quantum computing will not be quantum experts. Software abstractions, development environments, and debugging tools must therefore evolve to hide unnecessary complexity while preserving performance and flexibility. From this perspective, progress in quantum software engineering is just as critical as advances in hardware.

**Technology sovereignty and strategic considerations:** Finally, quantum computing must be viewed within a broader geopolitical and economic context. As with artificial intelligence and large language models, questions of technology sovereignty are becoming increasingly important. Many quantum hardware and software platforms are currently dominated by actors in the United States or China, creating dependencies that may pose strategic risks for other regions. For Europe and other parts of the world, fostering independent quantum ecosystems – spanning research, industry, and infrastructure – will be essential to ensure long-term autonomy, resilience, and competitiveness. This includes not only hardware development, but also open software stacks, talent cultivation, and application-driven innovation.

### 21.4.3 Personal outlook

Reflecting on the projects and experiences described in this essay, my perspective on quantum computing has become both more grounded and more focused. Entering the field from a background in classical optimization, I was initially confronted with a steep learning curve and with expectations that quantum computing might deliver rapid breakthroughs. What emerged instead was a more nuanced understanding:

quantum computing is neither a near-term replacement for classical methods nor a distant curiosity, but a developing computational paradigm whose value will depend critically on how and where it is applied.

Looking ahead, I expect the first clear and sustained impact of quantum computing to arise outside of optimization. In particular, the simulation of molecules and quantum systems in chemistry appears to be the most natural near-term application. Here, the underlying problem structure aligns closely with quantum physics itself, and progress has direct implications for areas such as materials science and drug discovery, where substantial investments are already being made. In comparison, optimization is likely to be a runner-up: while challenging, it offers structured problem classes – such as robust and multiobjective optimization – where quantum methods may eventually complement classical techniques. Quantum machine learning, by contrast, currently appears overhyped. Fundamental challenges, especially around data encoding and the integration of large classical datasets into quantum workflows, remain largely unresolved. Beyond computing, other quantum technologies, including communication and sensing, may well deliver practical utility earlier than quantum algorithms running on general-purpose processors.

Within optimization, my own focus will remain firmly on hybrid quantum–classical approaches. The experiences from energy systems and portfolio optimization suggest that the most promising role for quantum algorithms lies in supporting exploration, robustness, and trade-off analysis rather than in competing with classical solvers on deterministic benchmarks. In particular, I intend to continue working on quantum and hybrid algorithms for robust and multiobjective optimization, where uncertainty, competing objectives, and the need for diverse solution sets are intrinsic features of the problem rather than complications to be eliminated.

More broadly, these experiences have reinforced my belief that progress in quantum computing will be incremental and interdisciplinary. Meaningful advances will not come from isolated algorithmic breakthroughs alone, but from sustained collaboration between hardware developers, algorithm designers, and application experts who understand real decision-making contexts. For practitioners in optimization, this means engaging with quantum computing not as a disruptive alternative, but as a complementary tool whose strengths must be carefully matched to the structure of the problem at hand.

In this sense, my outlook on quantum computing is one of curious but cautious commitment. I do not expect rapid or universal quantum advantage (I might be mistaken), but I do see medium- to long-term potential – provided we remain honest about current limitations and disciplined in aligning quantum methods with problems where they can genuinely add value.

Quantum utility or advantage will not come from doing the same things  
faster, but from doing different things better.

## References

Bernhardt, C. (2019). *Quantum computing for everyone*. The MIT Press. <https://doi.org/10.7551/mitpress/11860.001.0001>

Ekstrom, L., Wang, H., & Schmitt, S. (2025). Variational quantum multiobjective optimization. *Physical Review Research*, 7(2), 023141. <https://doi.org/10.1103/physrevresearch.7.023141>

Halfmann, P. (2021). *Advances in multiobjective optimisation: Scalarisation, approximation, and complexity* [PhD Thesis]. Dr. Hut Verlag. <https://doi.org/10.13140/RG.2.2.26352.94726>

Halfmann, P., Holzer, P., Plociennik, K., & Trebing, M. (2023). A quantum computing approach for the unit commitment problem. In *Operations Research Proceedings 2022* (pp. 113–120). Springer. [https://doi.org/10.1007/978-3-031-24907-5\\_14](https://doi.org/10.1007/978-3-031-24907-5_14)

Halfmann, P., Lenk, S., Hegemann, N., & Oberthaler, M. (2024). *Enerquant – quantum optimization for energy systems: Final report*. Fraunhofer ITWM. <https://www.itwm.fraunhofer.de/content/dam/itwm/de/documents/anwendungsfelder/EnerQuant-Final-Report.pdf>

Halfmann, P., Trebing, M., & Lenk, S. (2024). Harnessing inferior solutions for superior outcomes: Obtaining robust solutions from quantum algorithms. *Proceedings of the Genetic and Evolutionary Computation Conference Companion*, 1950–1953. <https://doi.org/10.1145/3638530.3664160>

Johnson, M. W., et al. (2011). Quantum annealing with manufactured spins. *Nature*, 473(7346), 194–198. <https://doi.org/10.1038/nature10012>

King, A. D., et al. (2025). Beyond-classical computation in quantum simulation. *Science*, 388(6743), 199–204. <https://doi.org/10.1126/science.ado6285>

Kotil, A., et al. (2025). Quantum approximate multi-objective optimization. *Nature Computational Science*, 5(12), 1168–1177. <https://doi.org/10.1038/s43588-025-00873-y>

Mauron, L., & Carleo, G. (2025). Challenging the quantum advantage frontier with large-scale classical simulations of annealing dynamics. <https://doi.org/10.48550/ARXIV.2503.08247>

Nielsen, M. A., & Chuang, I. L. (2012). *Quantum computation and quantum information: 10th anniversary edition*. Cambridge University Press. <https://doi.org/10.1017/cbo9780511976667>

Turkalj, I., et al. (2024). Quadratic unconstrained binary optimization approach for incorporating solvency capital into portfolio optimization. *Risks*, 12(2), 23. <https://doi.org/10.3390/risks12020023>

Turkalj, I., et al. (2025). Enhancing variational quantum algorithms for multicriteria optimization. <https://doi.org/10.48550/ARXIV.2506.22159>



# Chapter 22

## From Bits to Qubits

### *A banker's journey to quantum Monte Carlo*

Rafał Pracht

#### 22.1 Opening Perspective

When I began my studies in computer science at the Military University of Technology in 2002, quantum computing existed only within physics laboratories and theoretical discourse. Like most engineers of my generation, I was trained within a classical paradigm of computation: bits, logic gates, deterministic algorithms, and hardware architectures optimized to execute Boolean operations with ever-increasing efficiency. At that time, the perceived limits of computing were primarily technological. Progress was expected to come through faster processors, larger memory capacity, and more powerful parallel systems.

Years later, while working in global financial institutions such as Moody's Analytics, Deloitte, KBC Group, SAS Institute, and PZU, I encountered a fundamentally different type of limitation. In finance, the primary constraint is not hardware performance, but computational complexity itself. Modern financial markets rely heavily on simulation-based methods, particularly Monte Carlo techniques, to price derivatives and measure risk exposures. These methods require the generation of millions, and often billions, of possible future scenarios. Despite decades of advances in high-performance computing, such simulations remain slow, costly, and constrained by inherent mathematical scaling laws.

This realization led me toward an unexpected domain: quantum computing. Initially, I viewed the field with skepticism regarding its practical applicability. However,

---

This chapter is a featured contribution to the monograph *A Portrait of Quantum Technologies in Finance*, edited by Oswaldo Zapata, PhD, and published by The Quantum Finance Boardroom (2026). All rights reserved by the chapter's author.

in 2018, I began studying quantum computing during a professional certification program at MIT. What I discovered was not merely a faster computational platform, but an entirely new computational paradigm. Quantum computing does not simply accelerate classical algorithms; it fundamentally transforms how complex problems can be represented, processed, and ultimately solved.

Motivated by this insight, I began conducting research on the application of quantum computing to financial mathematics. As a Quantum Researcher at BNP Paribas, I developed algorithms for derivative pricing using quantum computational frameworks. At the same time, I commenced doctoral studies under the supervision of Professor Dariusz Gałtarek, co-author of the LIBOR Market Model (BGM). This experience provided a deep understanding of interest rate modeling and clarified the specific areas in financial computation where quantum methods may deliver meaningful advantages.

Following the MIT program, I also initiated the development of an open-source Julia library designed for quantum computing applications in finance. This work culminated in a presentation at JuliaCon 2023 and ultimately led to the founding of FinQbit in September 2023 together with my co-founder Tomasz. The library became the technological foundation of the company's efforts to develop practical quantum solutions for the financial industry.

This chapter presents an accessible introduction to Quantum Monte Carlo, one of the most promising applications of quantum computing in finance. Rather than emphasizing technical complexity, the discussion is structured as a narrative journey: from the foundations of classical computation, through the computational challenges faced by modern financial institutions, to the emerging quantum techniques that may fundamentally transform how banks price complex derivatives and manage risk.

## 22.2 From Classical Logic to Computational Limits

### 22.2.1 The classical computing paradigm

To understand why quantum computing represents such a profound shift, it is essential to first examine the foundations of classical computation. Modern digital computers, regardless of their scale or sophistication, are built upon a remarkably simple conceptual framework: the manipulation of binary information through logical operations.

At the most fundamental level, classical computation relies on *bits*, units of information that can take only two possible values, conventionally denoted as 0 and 1. These binary states are physically realized in electronic circuits as voltage levels, but conceptually they correspond to logical truth values: *false* and *true*. Complex computations emerge from the systematic transformation of these bits through logical rules governed by Boolean algebra.

**Boolean algebra and logic gates.** Boolean algebra provides the mathematical

language of classical computing. It defines how binary variables can be combined using basic logical operations such as AND, OR, and NOT. These operations are implemented physically as *logic gates*, which form the elementary building blocks of digital circuits.

For example, the logical AND operation can be expressed as:

$$x \wedge y = \begin{cases} 1 & \text{if } x = 1 \text{ and } y = 1, \\ 0 & \text{otherwise.} \end{cases}$$

Similarly, the NOT operation simply reverses the value of a bit:  $\neg x = 1 - x$ . Although these operations appear simple, their combinations allow classical computers to perform arbitrarily complex tasks. Any algorithm, no matter how sophisticated, ultimately reduces to sequences of such elementary logical transformations.

**Deterministic computation and irreversibility.** A defining characteristic of classical computation is its *deterministic* nature. Given an initial input and a sequence of operations, the outcome is uniquely determined. Each computational step transforms the system from one well-defined state to another.

Importantly, most classical logic operations are *irreversible*. For instance, if we consider the AND gate, knowing the output alone does not allow us to reconstruct the original inputs. If the result equals 0, multiple input combinations could have produced it. This loss of information is not merely conceptual; it has physical consequences.

According to Landauer's principle, erasing one bit of information necessarily dissipates a minimum amount of energy:

$$E \geq k_B T \ln 2,$$

where  $k_B$  is the Boltzmann constant and  $T$  is temperature. This principle establishes a fundamental link between computation, information, and thermodynamics, implying that classical computation inevitably involves energy dissipation due to irreversibility.

**Arithmetic operations as digital circuits.** Higher-level numerical operations, such as addition, multiplication, and division, are constructed from networks of logic gates. For example, a simple binary addition is implemented using *half-adders* and *full-adders*, circuits that combine logical operations to compute both the sum and the carry bit.

In essence, all arithmetic computations reduce to Boolean transformations of binary representations. Even complex financial models, large-scale simulations, and machine learning algorithms ultimately rely on vast sequences of these basic logical operations executed at extraordinary speed.

**Scaling through parallelism and hardware improvements.** For decades, improvements in computational performance were driven primarily by hardware scaling. The exponential growth described by Moore's Law enabled continuous increases in processing speed, memory capacity, and circuit density.

When physical miniaturization approached fundamental limits, progress shifted toward *parallelism*. Modern computing architectures distribute tasks across multiple cores, processors, and distributed systems. High-performance computing clusters and cloud infrastructures now enable massive simulations that would have been inconceivable just a few decades ago.

However, this scaling strategy encounters an unavoidable constraint. While hardware improvements can reduce constant factors in computation time, they cannot alter the underlying mathematical complexity of problems. If a task requires a number of operations that grows exponentially with system size, no amount of hardware parallelization can fully overcome this limitation.

**Intrinsic efficiency limits of classical computation.** This distinction between technological and mathematical limits is critical. Classical computers excel at tasks whose computational complexity scales polynomially with problem size. Yet many problems central to science, optimization, and financial risk modeling exhibit exponential growth in computational requirements.

Monte Carlo simulation provides a clear example. To reduce statistical error by a factor of two, one must increase the number of simulated scenarios by a factor of four. This square-root convergence rate is not a hardware limitation but a fundamental mathematical property of random sampling methods.

Thus, even with unlimited engineering progress, classical computation remains bound by intrinsic efficiency limits imposed by algorithmic scaling laws. These limits define the boundary beyond which new computational paradigms, rather than faster hardware, become necessary.

### 22.2.2 Reversible computation and the bridge to quantum logic

The transition from classical to quantum computing does not begin with exotic physics, but with a subtle yet fundamental shift in how computations are performed. This shift revolves around a single concept: *reversibility*. Understanding reversibility provides the most intuitive bridge between familiar classical logic and the principles of quantum computation.

**Why quantum operations must be reversible.** In classical computing, most logical operations are irreversible. As discussed earlier, gates such as AND or OR lose information about their inputs. Once the computation is performed, it is generally impossible to reconstruct the original state from the output alone.

Quantum computing operates under a fundamentally different rule. The evolution of a quantum system must obey the laws of quantum mechanics, which are governed by deterministic and reversible transformations. Mathematically, these transformations are described by *unitary operations*. A unitary transformation preserves the total probability of the system and, crucially, can always be inverted.

This requirement means that every quantum computation must be reversible: given the final state of the system, it must be possible, at least in principle, to reconstruct the initial state by applying the inverse transformation.

**Unitary transformations as quantum gates.** In quantum computing, logical operations are implemented through unitary matrices acting on quantum states. If we denote the state of a quantum system by a vector  $|\psi\rangle$ , then a quantum operation is represented as:

$$|\psi_{\text{out}}\rangle = U|\psi_{\text{in}}\rangle,$$

where  $U$  is a unitary matrix satisfying:

$$U^\dagger U = I.$$

This condition ensures that the transformation is reversible and that no information is destroyed during computation. Unlike classical gates, which map multiple input states to a single output, quantum gates preserve the full structure of the information encoded in the system.

**Classical reversible logic: the Toffoli gate.** Interestingly, reversibility is not unique to quantum computing. It is possible to construct classical logic circuits that are fully reversible. The most important example is the *Toffoli gate*, also known as the controlled-controlled-NOT gate.

The Toffoli gate operates on three bits:

$$(x, y, z) \mapsto (x, y, z \oplus (x \wedge y)),$$

where  $\oplus$  denotes addition modulo two. In simple terms, the third bit flips only if the first two bits are both equal to one.

This gate is reversible because the original inputs can always be recovered from the outputs. Remarkably, it has been proven that the Toffoli gate alone is sufficient to implement any classical computation in a reversible manner. This makes it a crucial conceptual link between classical digital logic and quantum circuits.

**Embedding classical logic in quantum circuits.** Because classical reversible gates can be expressed as unitary transformations, classical computations can be directly embedded within quantum circuits. A classical bit corresponds to a quantum bit (qubit) restricted to the states  $|0\rangle$  and  $|1\rangle$ . Classical reversible gates then become special cases of quantum gates acting on these basis states.

This observation has profound implications. It means that quantum computers do not replace classical computation; rather, they generalize it. Any classical algorithm can, in principle, be implemented on a quantum computer using reversible logic.

However, quantum computation extends far beyond this classical subset. Once we allow quantum states to exist in superpositions of  $|0\rangle$  and  $|1\rangle$ , entirely new computational possibilities emerge, including the ability to process many potential

outcomes simultaneously within a single coherent state.

**Preparing for the quantum paradigm.** Reversibility thus serves as the conceptual bridge between classical and quantum computation. It transforms the familiar language of digital logic into a framework compatible with quantum mechanics, while preserving intuitive connections to classical circuit design.

### 22.2.3 From reversibility to superposition: entering the quantum world

While reversible computation forms the bridge between classical and quantum logic, it does not yet capture what makes quantum computing truly powerful. The defining features of quantum computation arise from three uniquely quantum phenomena: superposition, interference, and measurement. Together, these concepts transform computation from a deterministic sequence of operations into a fundamentally probabilistic yet highly structured process.

**From bits to qubits.** In classical computing, the basic unit of information is the bit, which can take one of two possible values: 0 or 1. Every computation ultimately reduces to manipulating these discrete states through logical operations.

Quantum computing introduces a new unit of information known as the *quantum bit*, or *qubit*. Unlike a classical bit, a qubit can exist in a combination of both states simultaneously. Mathematically, this is represented as:

$$|\psi\rangle = \alpha|0\rangle + \beta|1\rangle,$$

where  $\alpha$  and  $\beta$  are complex numbers satisfying:  $|\alpha|^2 + |\beta|^2 = 1$ .

This expression means that a qubit is not simply “between” 0 and 1, but rather exists in a weighted combination of both possibilities at once. The coefficients determine the probabilities of observing each outcome when the qubit is measured.

**Superposition: parallelism at the fundamental level.** The ability of qubits to exist in superposition is the first key advantage of quantum computation. When multiple qubits are combined, their joint state can represent an exponentially large number of classical configurations simultaneously.

For example, while three classical bits can represent only one of eight possible states at any given moment, three qubits can represent all eight states at the same time within a single quantum state:

$$|\psi\rangle = \sum_{i=0}^7 \alpha_i |i\rangle.$$

This property is often described as *quantum parallelism*. It allows a quantum computer to process many potential outcomes simultaneously, rather than sequentially

as in classical computation.

**Interference: selecting the correct answer.** Superposition alone does not guarantee a computational advantage. The true power of quantum computing emerges from the phenomenon of *interference*. Quantum algorithms are designed so that the amplitudes corresponding to correct solutions reinforce each other, while those corresponding to incorrect solutions cancel out.

This process is analogous to wave interference in physics, where overlapping waves can either amplify or diminish one another depending on their phase relationships. In quantum computation, carefully constructed sequences of unitary operations manipulate these amplitudes to steer the system toward the desired outcome.

**Measurement: the fundamental constraint.** Despite the richness of quantum superposition, a fundamental limitation remains: when a quantum state is measured, it collapses to a single classical outcome. The result of the measurement is probabilistic, determined by the squared magnitudes of the amplitudes in the quantum state.

This measurement constraint has profound implications for algorithm design. Unlike classical algorithms, which can directly access intermediate results, quantum algorithms must encode the final answer in such a way that it can be extracted efficiently from a single measurement or a small number of repeated measurements.

**A new computational paradigm.** Superposition, interference, and measurement together define a radically different computational paradigm. Quantum computers do not simply perform calculations faster; they manipulate probability distributions in ways that are fundamentally inaccessible to classical machines.

For decision-makers in finance, the key takeaway is not the physics itself, but the consequence: quantum computation provides a new method for estimating complex probabilistic quantities. Since many financial problems, such as derivative pricing and risk measurement, are inherently probabilistic, this paradigm aligns naturally with the core computational challenges faced by modern financial institutions.

## 22.3 Financial Simulation: The Computational Heart of Modern Banking

### 22.3.1 Why finance depends on simulation

Modern financial institutions operate in an environment defined by uncertainty. Prices of assets evolve unpredictably, interest rates fluctuate continuously, and market conditions can shift rapidly in response to economic, political, and technological forces. As a result, decision-making in finance is fundamentally probabilistic rather than deterministic.

To manage this uncertainty, banks rely heavily on computational models that quantify potential future outcomes. These models form the backbone of many critical financial processes, including:

- **Derivative pricing.** Financial derivatives, such as options and structured products, derive their value from uncertain future market movements. Determining their fair value requires estimating expected payoffs under probabilistic scenarios.
- **Market risk measurement.** Banks must evaluate how their portfolios may perform under a wide range of possible market conditions. Risk metrics such as Value-at-Risk (VaR) and Expected Shortfall depend on large-scale scenario generation.
- **Capital requirement calculations.** Regulatory frameworks such as Basel III and Basel IV require banks to quantify potential losses under stressed market environments. These calculations involve massive scenario-based simulations.
- **Stress testing.** Financial institutions routinely conduct forward-looking stress tests to assess resilience under extreme but plausible economic conditions.

At the mathematical core of all these tasks lies a common computational problem: estimating expectations under uncertainty.

**The central role of expected values.** Many financial quantities can be expressed in a simple conceptual form:

$$\text{Financial Value} = \mathbb{E}[\text{Future Outcome under Uncertainty}].$$

For example:

- The price of a derivative equals the expected discounted payoff under a risk-neutral probability measure.
- A risk metric equals the expected loss under adverse scenarios.
- A capital requirement reflects the expected exposure under stressed market conditions.

Although this formulation appears straightforward, the difficulty arises from the complexity of the underlying uncertainty. Financial markets involve many interacting variables, evolving over time in ways that cannot be captured by simple formulas.

**The curse of dimensionality in finance.** In realistic settings, a financial system may involve:

- Multiple risk factors (interest rates, equity prices, credit spreads, exchange rates)
- Complex stochastic dynamics
- Long time horizons
- Path-dependent contractual features

The number of possible future scenarios grows exponentially with each additional risk factor and time step. This phenomenon, known as the *curse of dimensionality*, makes analytical solutions infeasible for most practical problems in modern finance.

As a result, closed-form formulas exist only for a small subset of simplified models. In real-world applications, financial institutions must rely on numerical methods capable of approximating expectations in very high-dimensional spaces.

**Simulation as the universal solution.** The most widely used approach to overcoming these challenges is simulation. Instead of attempting to solve complex probabilistic models analytically, banks generate large numbers of hypothetical future scenarios and estimate expectations by averaging the resulting outcomes.

This simulation-based paradigm has become the computational engine of modern finance. It is used daily across trading desks, risk departments, asset management firms, and regulatory frameworks. However, while simulation provides remarkable flexibility, it comes with a fundamental limitation: computational cost. Accurate estimation typically requires an enormous number of simulated scenarios, often reaching millions or even billions of paths. As financial systems become more complex and regulatory requirements more stringent, the computational burden continues to grow.

### 22.3.2 Monte Carlo methods in finance

The most widely used computational tool for dealing with uncertainty in finance is the Monte Carlo method. Its popularity stems from a simple yet powerful idea: when analytical solutions are unavailable, one can approximate expected values by simulating many possible future scenarios and averaging the outcomes.

**Simulating possible futures.** At its core, the Monte Carlo approach mirrors how one might reason intuitively about uncertainty. Instead of trying to predict a single future path for a financial variable, the method generates a large number of plausible scenarios, each representing one possible evolution of the market.

For example, to price an option on a stock, the simulation process typically involves:

- Generating many possible future price paths for the underlying asset.
- Computing the payoff of the option in each simulated scenario.
- Averaging these payoffs to estimate the expected value.

Mathematically, the price of a derivative can be expressed as an expected value:

$$V = \mathbb{E}[\text{Payoff under future uncertainty}].$$

Because the true probability distribution is too complex to evaluate directly, Monte Carlo methods approximate this expectation through random sampling.

**Averaging as a universal strategy.** The strength of the Monte Carlo framework lies in its generality. It does not require simplifying assumptions about the structure of the problem and can handle:

- High-dimensional risk factors
- Complex payoff structures
- Nonlinear market dynamics
- Long time horizons

For this reason, Monte Carlo simulation has become the universal computational engine for pricing derivatives, computing risk exposures, and performing regulatory stress testing across the global financial system.

**The fundamental limitation: slow convergence.** Despite its flexibility, Monte Carlo simulation suffers from a fundamental mathematical limitation: its slow rate of convergence.

If  $N$  independent scenarios are simulated, the statistical error of the estimate decreases according to the well-known law:

$$\text{Error} \sim \frac{1}{\sqrt{N}}.$$

This relationship has profound practical implications. To reduce the estimation error by a factor of ten, one must increase the number of simulations by a factor of one hundred. Achieving very high precision, therefore, requires an enormous computational effort.

In real banking applications, this translates into:

- Millions to billions of simulated scenarios
- Long computation times even on large computing clusters
- High infrastructure and energy costs

As financial models become more complex and regulatory expectations more stringent, this computational burden continues to grow.

**A fundamental computational bottleneck.** Importantly, this slow convergence is not a limitation of hardware or software engineering. It is a fundamental property of classical statistics. No matter how fast computers become, the  $\frac{1}{\sqrt{N}}$  scaling law cannot be overcome within the classical Monte Carlo framework. This scaling law imposes a fundamental computational constraint. For large financial institutions, reducing pricing or risk estimation error from, for example, 1% to 0.1% requires one hundred times more computational effort. When models involve thousands of risk factors and millions of instruments, such scaling quickly becomes prohibitively expensive.

The second challenge arises from memory requirements rather than computational speed alone. For instance, when pricing American-style options using the

Longstaff–Schwarz Least-Squares Monte Carlo method, it is necessary to store all simulated paths in memory in order to implement the backward dynamic programming procedure. In high-dimensional settings, such as basket options involving multiple underlying assets, this requirement becomes prohibitive. Even if the desired statistical accuracy could, in principle, be achieved under the classical convergence law, the sheer volume of data that must be retained may exceed the memory capacity of even the most advanced supercomputing systems. As a result, many practically relevant high-dimensional optimal stopping problems remain computationally intractable within classical architectures.

This observation naturally raises a fundamental question:

*Is it possible to estimate expectations more efficiently by changing not the hardware speed, but the very principles of computation?*

## 22.4 From Classical to Quantum: A New Paradigm for Estimation

As discussed in the previous section, the primary limitation of classical Monte Carlo methods is not technological but mathematical. Their convergence speed is governed by a fundamental statistical law, which dictates that improving accuracy requires a disproportionately large increase in the number of simulated scenarios. This constraint has long shaped the computational economics of financial modeling, determining the time, cost, and feasibility of large-scale risk calculations.

The emergence of quantum computing introduces a fundamentally different approach to estimation. Rather than relying solely on repeated sampling in the classical sense, quantum algorithms exploit superposition and interference to extract statistical information more efficiently. This shift does not merely accelerate existing techniques; it changes the underlying relationship between computational effort and precision. As a result, problems that currently demand enormous computational resources could, in principle, become significantly more tractable.

One of the most important developments in this direction is a quantum technique designed specifically to enhance sampling efficiency. This method, known as Quantum Amplitude Estimation, represents a cornerstone of quantum-enhanced simulation and provides the conceptual bridge between classical Monte Carlo methods and their quantum counterparts.

### 22.4.1 Quantum amplitude estimation: a breakthrough in sampling

The most important result for financial simulation is the existence of the *Quantum Amplitude Estimation* (QAE) algorithm. QAE provides a quantum analogue of classical Monte Carlo estimation. Its key achievement is an improved convergence

rate:

$$\text{Quantum Error} \sim \frac{1}{N}.$$

This represents a quadratic improvement over the classical scaling:

$$\frac{1}{\sqrt{N}} \longrightarrow \frac{1}{N}.$$

The implications of this improvement are profound:

- To achieve the same precision, a quantum algorithm requires dramatically fewer simulations.
- Computational costs for large-scale risk calculations could be reduced by orders of magnitude.
- Problems currently considered intractable may become feasible.

From a financial perspective, this improvement directly translates into faster pricing, more accurate risk measurement, and the ability to perform more detailed scenario analyses.

## 22.4.2 From theory to practice

While current quantum hardware remains limited, the theoretical foundations of quantum sampling are well established. Quantum Monte Carlo methods demonstrate that quantum computing is not merely a faster version of classical computation, but rather a fundamentally different approach to handling uncertainty and expectation.

Instead of repeatedly sampling independent scenarios, quantum algorithms encode entire probability distributions into quantum states and extract statistical properties through amplitude manipulation. This paradigm shift forms the foundation of quantum computational finance.

## 22.5 Quantum Monte Carlo in Finance: Intuition Without Physics

At its core, most problems in quantitative finance can be reduced to a simple mathematical structure: the computation of an expected value under uncertainty. Quantum Monte Carlo takes a fundamentally different approach. Instead of simulating scenarios one by one, it encodes the entire probability distribution into a quantum state.

### 22.5.1 Encoding probability into a quantum state

The key idea behind quantum Monte Carlo is surprisingly intuitive. Rather than generating individual samples, a quantum computer prepares a superposition that represents all possible scenarios simultaneously.

Formally, a probability distribution  $\{p_i\}$  over possible outcomes  $\{x_i\}$  can be encoded into a quantum register as:

$$|\psi\rangle = \sum_i \sqrt{p_i} |x_i\rangle.$$

Each basis state  $|x_i\rangle$  corresponds to a possible market scenario, while the amplitude  $\sqrt{p_i}$  encodes its probability. This representation is powerful because it stores an entire distribution within a single quantum object. Instead of running millions of simulations sequentially, the quantum system holds all possibilities at once. Once the distribution is loaded, the payoff function can be applied through reversible quantum arithmetic. This transforms the quantum state so that the expected payoff is encoded in the amplitude of a specific measurement outcome. Quantum Amplitude Estimation can then extract this expectation with quadratic efficiency.

### 22.5.2 Why state preparation is the real bottleneck

While the theoretical speed-up of quantum Monte Carlo is well established, practical implementation faces a major challenge: how to efficiently load realistic probability distributions into quantum hardware. In finance, probability distributions rarely have simple closed forms. They typically arise from stochastic differential equations that describe the evolution of asset prices, interest rates, or volatility. Preparing such distributions directly on a quantum computer is difficult because:

- Financial models involve continuous dynamics.
- Distributions evolve over time.
- Dependencies between variables can be complex.

If state preparation is inefficient, the potential quantum speed-up disappears. In fact, this problem has long been recognized as the central obstacle to practical quantum Monte Carlo in finance.

## 22.6 The Quantum Binomial Tree: Loading Financial Markets into Qubits

### 22.6.1 The classical binomial tree: a familiar foundation

One of the most widely used tools in financial engineering is the binomial tree model. Introduced as an intuitive discretization of stochastic processes, it remains a cornerstone of derivative pricing, particularly for options and interest rate models.

The idea is simple. Instead of modeling continuous market evolution directly, time is divided into discrete steps. At each step, the underlying asset can move up or down with certain probabilities. Over time, this produces a branching structure that represents all possible future price scenarios. After  $T$  steps, the number of possible paths equals:  $2^T$ . This exponential growth is both the strength and the weakness of

the binomial approach. On one hand, it provides a flexible and intuitive representation of stochastic dynamics. On the other, it quickly becomes computationally infeasible.

For example, a tree with only 50 time steps already contains more than one quadrillion possible paths. Classical computers cannot explicitly store or process such structures. In practice, financial engineers use approximations, pruning techniques, or Monte Carlo simulations to manage this explosion. Nevertheless, the fundamental scaling limitation remains.

### 22.6.2 A natural fit for quantum representation

From the perspective of quantum computing, the exponential growth of binomial trees is not a problem. In fact, it is precisely the type of structure that quantum systems can represent naturally. A quantum register of  $n$  qubits can store a superposition of  $2^n$  states simultaneously.

This means that a binomial tree with exponentially many paths can be encoded compactly within a linear number of qubits. Instead of storing each path explicitly, the quantum computer represents all possible paths as components of a single quantum state:

$$|\psi\rangle = \sum_{\text{paths}} \sqrt{p(\text{path})} |\text{path}\rangle.$$

Here, each basis state corresponds to one possible trajectory of the market, while its amplitude encodes the probability of that trajectory.

Thus, what appears as an intractable combinatorial explosion in classical computing becomes a natural and compact representation in the quantum domain.

### 22.6.3 From classical trees to quantum circuits

The key innovation of the Quantum Binomial Tree lies in translating the classical branching structure into a sequence of quantum operations. Each time step of the classical model corresponds to applying a controlled rotation on a qubit that represents the direction of movement at that step. Intuitively:

- Each qubit represents one time increment.
- The state  $|0\rangle$  encodes a downward movement.
- The state  $|1\rangle$  encodes an upward movement.
- Controlled rotations embed the correct transition probabilities.

By sequentially applying these operations, the algorithm constructs a quantum state that encodes the entire stochastic evolution of the underlying asset.

This process requires only  $O(T)$  quantum operations for a model with  $T$  time steps, even though the classical tree contains  $2^T$  paths. This exponential compression is the central insight behind the Quantum Binomial Tree.

### 22.6.4 Preserving financial model flexibility

A critical requirement for any practical financial algorithm is flexibility. Real markets cannot be described by simplistic assumptions such as constant volatility. The Quantum Binomial Tree is designed to address this need.

Unlike earlier state-preparation methods, it can efficiently encode multi-dimensional stochastic differential equations of the form. Regard an economy model with  $d$  assets, whose prices  $S(t) = [S_1(t), S_2(t), \dots, S_d(t)]^\top$ , are described by a system of Stochastic Differential Equations:

$$dS(t) = \mu(S(t), t) dt + \Sigma(S(t), t) dW_t \quad (22.1)$$

with  $W_t$  a  $d$ -dimensional Brownian motion,  $\Sigma(S(t), t)\Sigma(S(t), t)^\top = \mathbf{Cov}(S(t), t)$  is deterministic covariance function to  $\mathbb{R}^{d \times d}$  and  $\mu$  is deterministic function of the current states  $S(t)$  and time  $t$  to  $\mathbb{R}^d$ . This model can be effectively modeled using Quantum Binomial Tree.

This capability allows the algorithm to support a broad class of market models, including:

- Black–Scholes dynamics with term-structure volatility
- Local volatility models
- Interest rate lattice models
- Multi-period risk-neutral valuation frameworks

By retaining the intuitive structure of classical binomial trees while leveraging quantum superposition, the method bridges financial engineering practice and quantum computation.

### 22.6.5 Enabling quantum Monte Carlo speed-up

Once the stochastic dynamics are loaded via the Quantum Binomial Tree, the system is ready for expectation estimation using Quantum Amplitude Estimation. This combination yields a two-layer advantage:

- Quadratic acceleration in expectation estimation
- Exponential compression of scenario representation
- Elimination of classical memory bottlenecks in American option pricing

Together, these features allow quantum Monte Carlo to evaluate financial expectations using dramatically fewer computational resources than classical methods.

### 22.6.6 A new computational perspective for finance

The Quantum Binomial Tree represents more than a technical improvement. It reflects a deeper shift in how financial computation can be conceptualized. Classical

finance treats uncertainty as a collection of independent scenarios that must be simulated one at a time. Quantum finance instead treats uncertainty as a structured superposition that can be manipulated holistically.

In this sense, the Quantum Binomial Tree does not merely accelerate existing algorithms. It redefines how stochastic processes themselves are represented within computational systems. This paradigm shift opens the door to solving classes of financial problems that have long been considered computationally prohibitive.

## 22.7 Looking Ahead: Why Quantum Monte Carlo Matters for the Future of Finance

When reflecting on pivotal moments in the evolution of quantitative finance, it is natural to recall landmark intellectual breakthroughs such as the Black–Scholes–Merton model, which fundamentally transformed how markets understand and price risk. Yet alongside these theoretical achievements, equally profound progress has come from advances in computation. The history of quantitative finance is, in many ways, a history of computational revolutions that enabled practitioners to translate mathematical insight into practical decision-making tools.

One of the earliest examples of this transformation was the introduction of the binomial option pricing model by Cox, Ross, and Rubinstein. Its importance did not lie solely in its elegance, but in its computational practicality. By discretizing market uncertainty into a tractable numerical structure, the binomial model allowed practitioners to solve pricing problems that had previously been inaccessible. For the first time, financial engineers could harness digital computers to explore complex derivative payoffs and dynamic hedging strategies. In doing so, the model marked a decisive shift: finance was no longer constrained by analytic formulas alone, but could evolve alongside computational capabilities.

The subsequent rise of Monte Carlo simulation represented another decisive leap. As computing power expanded, financial institutions increasingly relied on large-scale stochastic simulations to value complex instruments, measure portfolio risk, and determine regulatory capital requirements. Monte Carlo methods became the computational backbone of modern banking, supporting applications ranging from derivative pricing to enterprise-wide stress testing. Yet despite enormous investments in high-performance computing infrastructure, a fundamental limitation persists: classical Monte Carlo converges slowly. Achieving higher accuracy requires exponentially greater computational effort, imposing significant costs in time, energy, and hardware resources.

Today, the financial industry stands at the threshold of another computational transition. Quantum computing offers not merely incremental speed improvements, but a fundamentally new way to represent and process uncertainty. Quantum Monte Carlo, in particular, illustrates this paradigm shift. By encoding entire probability distributions within quantum states and leveraging quantum interference, it promises to reduce the computational burden of expectation estimation dramatically. Problems

that currently require millions of simulations could, in principle, be solved with orders of magnitude fewer computational steps.

For financial institutions, the implications are far-reaching. Faster simulation capabilities could transform real-time risk management, enabling banks to compute exposures dynamically as market conditions evolve. More efficient valuation could support the pricing of complex structured products that are currently too costly to model with high precision. Perhaps most importantly, quantum acceleration could enable entirely new classes of financial analysis that remain computationally infeasible today.

History suggests that the most influential breakthroughs in quantitative finance emerge when theoretical insight converges with computational innovation. Just as early computers enabled the binomial model and later powered the Monte Carlo revolution, quantum computing may provide the next foundational infrastructure for financial modeling.

As Wayne Gretzky famously observed, one should “skate to where the puck is going, not where it has been.” For the next generation of quantitative professionals, this insight carries particular relevance. The future landscape of finance will increasingly be shaped by quantum technologies. Developing familiarity with quantum principles, computational thinking, and emerging quantum algorithms will position practitioners to lead rather than follow this transformation.

Quantum Monte Carlo does not represent the end of this journey, but rather its beginning. It signals a shift from a world in which computational limitations constrain financial imagination, to one in which new computational paradigms expand the very scope of problems that finance can address.



# Chapter 23

## Quantum Finance and Venture Capital

### *Why investors should pay attention to quantum computing in financial markets*

Reinaldo Coelho

#### 23.1 Quantum Finance Startups as an Investment Opportunity

Across financial history, durable competitive advantage has often been rooted in superior information processing. From ticker-tape arbitrage in the nineteenth century to high-frequency trading in the twenty-first, financial markets have rewarded those capable of modelling risk, uncertainty, and optimisation more effectively than their peers.

In my years operating at the intersection of financial markets and alternative investments, I have rarely encountered a technological domain that combines scientific depth, strategic optionality, and economic leverage as powerfully as quantum computing. In particular, quantum computing applied to financial services represents one of the most compelling frontier opportunities for venture capital over the next three to seven years.

Financial markets are, at their core, computational engines. Portfolio optimisation, risk aggregation, derivative pricing, scenario simulation, capital allocation, and market microstructure modelling all rely on solving mathematically complex problems at scale. As regulatory demands increase and market complexity intensifies, classical computing approaches, while powerful, are increasingly stretched in high-dimensional

---

This chapter is a featured contribution to the monograph *A Portrait of Quantum Technologies in Finance*, edited by Oswaldo Zapata, PhD, and published by The Quantum Finance Boardroom (2026). All rights reserved by the chapter's author.

optimisation and simulation contexts.

Quantum computing offers the prospect of meaningful advantages in selected computational classes, particularly combinatorial optimisation and Monte Carlo estimation. The sector has reached what some analysts describe as an inflection point: private investment accelerated sharply in 2024 and 2025, enterprise experimentation expanded, and maturing software layers brought commercial use cases closer to viability. In 2024 alone, venture capital backed quantum startups raised approximately \$1.9 billion across 62 rounds – a 138 per cent increase over the prior year—and the pace continued to accelerate through the first half of 2025.

From a venture capital perspective, the attractiveness is threefold. First, finance converts even marginal computational improvement into significant economic value. Second, the customer base made of global banks, hedge funds, and asset managers is capable of paying enterprise-grade pricing for a performance edge. Third, even incremental hybrid quantum advantages can create asymmetric upside in market positioning.

The financial impact would not merely be operational, it would be structural. Institutions that can simulate faster, optimise more effectively, or price complex products more accurately will deploy capital more efficiently and potentially generate superior risk-adjusted returns. For venture capital funds, backing the enabling infrastructure of such a transformation represents a rare alignment between frontier science and financial scalability.

## 23.2 High-Value Domains in the Quantum Finance Frontier

The first step in evaluating the quantum finance landscape from an investment point of view, is to acquire knowledge about the areas where new investable opportunities might arise. The most promising near-term applications of quantum computing in finance do not represent a broad-spectrum replacement of classical IT infrastructure. Instead, they represent a surgical enhancement of specific and high-value computational domains: optimisation, simulation, pricing, and cryptography. Each of these areas addresses a bottleneck where classical brute-force scaling has reached diminishing returns.

### 23.2.1 Portfolio optimisation and combinatorial complexity

At its heart, portfolio construction under real-world conditions is a constrained optimisation problem of immense proportions. Investors seek to maximise expected return while minimising risk, subject to transaction costs, liquidity constraints, cardinality limits, and increasingly complex regulatory requirements. These formulations frequently reduce to quadratic unconstrained binary optimisation (QUBO) problems or related NP-hard structures. Classical solvers, even when running on large server farms, must rely on heuristics and approximations as dimensionality increases, often

settling for local rather than global optima.

Recent research has demonstrated that hybrid quantum-classical algorithms may improve performance in portfolio selection under realistic constraints. Quantum variational algorithms, in particular, have been applied to constrained risk-return trade-off frameworks that were previously computationally prohibitive. While current quantum devices remain noisy and limited in qubit count, the so-called NISQ (Noisy Intermediate-Scale Quantum) era, hybrid approaches represent a pragmatic middle ground. By combining classical preprocessing with quantum subroutines, firms can begin to capture practical advantages in selected problem regimes today.

From a venture capital perspective, optimisation represents a commercially tractable entry point. It is directly monetisable, integrates into existing institutional workflows, and offers a path for incremental improvement as hardware matures.

### 23.2.2 Risk modelling and Monte Carlo simulation

Monte Carlo simulation underpins the edifice of modern risk management. Institutions simulate thousands or millions of market paths to estimate tail risk and capital adequacy metrics such as Value at Risk (VaR) and expected shortfall. However, the precision of these simulations is fundamentally limited by the inverse square root of the number of samples, which is a constraint that makes high-precision estimation computationally expensive.

Quantum amplitude estimation has been shown, under idealised conditions, to offer a quadratic speedup in Monte Carlo sampling, potentially transforming hours of computation into minutes. These theoretical advantages have been formalised in the academic literature, providing a roadmap for how financial institutions might eventually bypass the sampling bottleneck. Subsequent research has extended such methods to derivative pricing contexts, demonstrating potential efficiency gains for path-dependent option valuation.

The economic relevance of this acceleration is substantial. Reducing computational time for VaR or expected shortfall calculations would allow institutions to rebalance portfolios more frequently, adjust hedging dynamically in response to sudden volatility, and optimise capital allocation on an intraday basis. Although fault-tolerant quantum computers capable of large-scale implementation remain a horizon goal, the algorithmic groundwork is established and is being tested in pilot programmes at several major financial institutions.

### 23.2.3 Derivative pricing

Complex derivative pricing involves high-dimensional integrals and conditional path structures that are notoriously difficult for classical processors to handle in real time. Quantum amplitude estimation frameworks have been proposed as potential accelerators for these specific calculations. Given that structured product desks often operate on thin margins with significant computational overhead, pricing efficiency improvements could materially alter cost structures and enable the creation of new,

more granularly priced financial instruments.

### 23.2.4 Cryptography and post-quantum security

Quantum computing also introduces systemic implications that cannot be ignored by the investment community. The potential for future quantum systems to break widely used public-key cryptography schemes, a scenario sometimes referred to as Q-Day, has prompted an urgent global effort to develop post-quantum cryptographic (PQC) standards. Financial institutions, which depend on public-key encryption for everything from interbank transfers to client authentication, face a mandatory transition window that regulators are beginning to formalise.

For venture capitalists, startups positioned in quantum-safe cryptography for financial institutions represent a parallel and complementary opportunity within the broader quantum-finance thesis. These firms are not merely offering computational speed, they are selling resilience. Collectively, the domains of optimisation, simulation, pricing, and cryptography form the intellectual and commercial foundation upon which the first generation of quantum finance startups is being built.

## 23.3 The Startup Landscape in Quantum Finance

The next step in preparation to invest in a particular thesis, is to evaluate some benchmarks. The quantum finance startup ecosystem has matured considerably over the past five years, evolving from academic curiosity toward enterprise-grade software solutions. The most finance-relevant startups have increasingly focused on the software layer and hybrid workflows. They adopt a hardware-agnostic posture that allows them to capture value regardless of which hardware architecture (superconducting circuits, trapped ions, or neutral atoms) eventually prevails. What follows is an analytical overview of some of the most prominent companies operating in this space.

### 23.3.1 Multiverse Computing

Founded in 2019 and headquartered in San Sebastián, Spain, Multiverse Computing has emerged as one of the most prominent finance-focused quantum startups. The firm develops quantum and quantum-inspired algorithms tailored to portfolio optimisation, risk management, and financial machine learning. Its early trajectory was defined by a €25 million Series A round in March 2024, led by Columbus Venture Partners with participation from Quantonation, CDP Venture Capital, and others.

Multiverse subsequently expanded internationally and established partnerships with financial institutions including BBVA, Bank of Canada, Crédit Agricole, and Moody's Analytics. In June 2025, the company closed a \$215 million Series B led by Bullhound Capital, with participation from HP Tech Ventures, Forgepoint Capital International, Toshiba, and Santander Climate VC, among others. The round reportedly valued the

company at over \$500 million, a roughly fivefold increase from its Series A valuation.

Multiverse’s product strategy rests on two pillars. The first is Singularity, a software development kit that allows financial professionals to access quantum algorithms through familiar interfaces. The platform offers solutions for parametric risk, foreign exchange trading, and dynamic portfolio optimisation. The second is CompactifAI, a compression technology rooted in tensor networks (a mathematical framework drawn from quantum physics) that reduces the size of large language models by up to 95% while maintaining near-original accuracy. This product addresses the escalating infrastructure costs associated with deploying large-scale AI, allowing institutions to run advanced models on-premise or on edge devices at a fraction of the usual energy and compute cost.

As of early 2026, Multiverse was reportedly seeking a €500 million funding round at a target valuation of €1.5 billion, signalling the company’s ambition to move from a finance-focused quantum startup into a broader AI efficiency platform. With over 160 patents, more than 100 enterprise clients, and recognition by CB Insights among the top 100 most promising AI companies, Multiverse represents one of the more commercially advanced players in the quantum-finance space.

### 23.3.2 SandboxAQ

SandboxAQ emerged from Alphabet in 2022 as an independent company focused on the intersection of artificial intelligence and quantum techniques. By early 2026, it had established itself as one of the most highly capitalised companies in the sector, having raised over \$950 million in total funding across multiple rounds. Its April 2025 Series E brought in more than \$450 million from investors including Google, NVIDIA, BNP Paribas, Ray Dalio, and Horizon Kinetics, and valued the company at approximately \$5.75 billion.

The company’s core technology centres on what it calls Large Quantitative Models (LQMs), which is a proprietary class of AI that, unlike text-based large language models, is grounded in the mathematical structures of physics and economics. LQMs are designed to simulate complex, high-dimensional systems with greater fidelity than purely statistical approaches, making them relevant to financial applications such as market simulation, tail-risk estimation, and capital adequacy modelling.

SandboxAQ’s product portfolio spans multiple domains. AQtive Guard is a cryptographic management platform that helps financial institutions inventory, monitor, and migrate their encryption infrastructure to post-quantum standards, which is a compliance-driven market that is likely to grow as fault-tolerant quantum computers approach viability. The platform has been deployed by partners including Accenture and has been approved (at least partially) for governmental use in the USA (achieved FedRAMP Ready status). In the financial services vertical, SandboxAQ’s LQMs are being applied to real-time VaR calculation, credit scoring, and fraud detection, while its quantum-sensing navigation technology (AQNav) is being explored for ultra-precise timestamping in high-frequency trading environments.

From an investor’s standpoint, SandboxAQ is notable for the breadth of its commercial

relationships and the dual nature of its quantum thesis: it offers both computational upside through LQMs and defensive value through post-quantum cryptography, positioning it to benefit regardless of the precise timeline for fault-tolerant quantum hardware.

### 23.3.3 QC Ware

QC Ware, headquartered in Palo Alto, California, with an additional office in Paris, occupies what might be described as the middle layer of the quantum software stack. Founded in 2014, the firm develops enterprise platforms that enable hybrid quantum-classical workflows for large institutional clients. Its 2021 Series B round of \$25 million attracted strategic investors including Goldman Sachs, Samsung, and Koch Disruptive Technologies, which was a notable signal that major financial institutions were moving from observation to active participation in the quantum software ecosystem.

QC Ware's Forge platform abstracts the complexity of quantum computing for financial analysts and data scientists, providing a cloud-based interface for developing and testing quantum algorithms without deep quantum physics expertise. More recently, the firm has expanded into quantum chemistry through its Promethium platform, targeting pharmaceutical and materials discovery applications.

### 23.3.4 Classiq

Based in Tel Aviv, Classiq Technologies has positioned itself as a provider of high-level quantum development tools, often described as a compiler and operating system for quantum computers. Founded in 2020, the company's platform allows banks and other enterprises to design complex quantum circuits without low-level quantum programming expertise, and it supports hardware-agnostic deployment across major cloud providers including AWS Braket, Microsoft Azure Quantum, and Google Cloud.

In May 2025, Classiq raised \$110 million in a Series C round led by Entrée Capital, which it described as the largest single funding round for a quantum software company at the time. The round attracted participation from HSBC, Samsung Next, NightDragon, and others. The company subsequently extended the round with investments from SoftBank Vision Fund 2, AMD Ventures, and Qualcomm Ventures, bringing total Series C proceeds to over \$200 million. Classiq's enterprise clients have include BMW, Citi, Deloitte, Rolls-Royce, and Mizuho, and its platform has been integrated into the curriculum at several leading universities.

### 23.3.5 Pasqal

Pasqal, headquartered near Paris and co-founded by Nobel laureate Alain Aspect, occupies a somewhat different position in the quantum-finance ecosystem: it is primarily a hardware company, building quantum processors based on neutral-atom technology. The company raised €100 million in a Series B round in January 2023,

led by Temasek, with participation from the European Innovation Council Fund, Bpifrance, and Quantonation. Its total funding has since exceeded \$300 million.

Pasqal’s relevance to quantum finance lies in the nature of its neutral-atom architecture, which is particularly well suited to modelling interconnected systems, which is a property that maps naturally onto problems in credit-risk analysis, portfolio correlation, and network-based financial modelling. The company has worked with Cr dit Agricole CIB on financial optimisation problems and counts over 25 commercial customers and partners. In March 2026, Pasqal announced plans to go public through a SPAC merger with Bleichroeder Acquisition Corp II, at a pre-money valuation of \$2 billion, accompanied by \$200 million in committed convertible financing—a signal of the growing institutional appetite for quantum-hardware exposure.

## 23.4 The Investor Landscape in Quantum Finance

While it is important to understand which startups have been standing out in the market, it is also important understand which investors have been most active in this scenario and how they are positioned. Currently, the venture capital landscape for quantum finance is in general split between deep-tech specialists and strategic corporate investors. While hardware has historically absorbed the majority of private quantum funding, roughly two-thirds of the approximately \$11 billion in disclosed private quantum investment over the past four years, software-layer startups are increasingly attracting attention due to shorter commercial timelines and asset-light business models. The following profiles illustrate the range of investor types active in the space.

### 23.4.1 Quantonation

Headquartered in Paris, Quantonation was among the first venture funds dedicated entirely to quantum technologies. The firm focuses on seed and Series A rounds, providing the patient, thesis-driven capital that deep-physics companies require in their earliest stages. Quantonation was an early backer of both Pasqal and Multiverse Computing, giving it meaningful exposure to what might be described as the European quantum stack. Its investment thesis focuses on the conviction that the transition to quantum computing represents a foundational shift in the economy’s computational infrastructure.

### 23.4.2 Temasek

As Singapore’s state-owned investment company, Temasek operates at a different scale and with a different mandate. The firm has provided large growth-stage rounds, most notably leading Pasqal’s €100 million Series B, and its investment thesis is rooted in national strategic interest. By backing quantum-hardware companies at scale, Temasek has helped position Singapore as a credible hub in the global quantum ecosystem, balancing commercial return expectations with sovereign technology

objectives.

### **23.4.3 QED**

QED is a premier fintech VC that recognizes quantum as a critical “infrastructure” play for the next generation of banking. Their thesis focus on the application-layer startups that can immediately. Their solutions can improve credit scoring, fraud detection, and capital efficiency, among other things. Their investment in Quantinuum signals a move toward backing firms that have a clear path to commercial revenue in the mid-term.

### **23.4.4 Deep-tech generalists and late-stage investors**

The investor base has broadened considerably beyond quantum specialists. NVIDIA backed three quantum startups in a single week in September 2025 (Quantinuum, PsiQuantum, and QuEra Computing) signalling the chipmaker’s view that quantum hardware will eventually integrate with its GPU ecosystem. SoftBank Vision Fund 2 entered the quantum-software space through Classiq. Bullhound Capital led Multiverse Computing’s \$215 million Series B. The entrance of these generalist and late-stage investors reflects a maturing capital structure in which quantum is no longer the exclusive province of specialist deep-tech funds, but is instead becoming a recognised asset class within broader technology portfolios.

### **23.4.5 Corporate venture capital**

Besides financial investors, like venture capital firms, strategic investors have also showing interest. Perhaps the strongest indicator of the sector’s maturation is the growing activity of bank-affiliated corporate venture arms. Institutions such as JPMorgan, Goldman Sachs, HSBC, and Citi are not merely prospective customers of quantum-finance software, they have become direct investors. For instance, Goldman Sachs participated in QC Ware’s Series B, HSBC backed Classiq’s Series C, and JPMorgan has built one of the most active in-house quantum research teams in the financial industry.

The logic behind these investments is both defensive and strategic. For large financial institutions, being excluded from proprietary quantum algorithms that may one day price complex financial products or optimise capital allocation represents a form of systemic competitive risk. By investing early, these institutions seek not only financial returns but also strategic access to the intellectual property and technical talent that will define the next generation of financial infrastructure.

## 23.5 The Geography of Quantum Finance Innovation

Finally, investors interested in quantum finance should also understand the dynamics that drive innovation in this space and where these entrepreneurial initiatives are coming from. Unlike the SaaS boom, which was heavily concentrated in Silicon Valley, quantum finance have been globally dispersed. Innovation clusters have emerged in regions where three conditions overlap: elite quantum research institutions, mature venture capital infrastructure, and concentrated financial-services demand.

Europe has established a notable lead in quantum-finance software. Paris anchors one end of the continental innovation corridor, providing deep physics talent through institutions such as the CNRS and École Polytechnique, which has given rise to companies like Pasqal and the fund Quantonation. On the other hand, San Sebastián and Madrid contribute financial-application expertise, most prominently through Multiverse Computing. Government support has also been a meaningful accelerant. Spain's PERTE Chip programme and France's quantum strategy under the France 2030 investment plan have provided both direct funding and regulatory frameworks that encourage private co-investment.

The United States East Coast remains the centre of gravity for financial-market proximity. The MIT-Harvard research hub in Boston has produced multiple quantum startups, and New York's concentration of banks and hedge funds provides a natural customer base. Silicon Valley, meanwhile, continues to lead in hardware scaling through the research divisions of Google, IBM, and NVIDIA. The Illinois Quantum and Microelectronics Park, anchored in the Chicago area, has recently attracted companies including Pasqal, PsiQuantum, and IBM, adding a Midwestern node to the US quantum geography.

Additionally, the United Kingdom, Singapore, and Israel also have been showing notable developments. The United Kingdom's combination of the City of London's financial depth with Cambridge's quantum-hardware breakthroughs, home to Quantinuum, among others, makes it a natural anchor in Europe. Singapore, through Temasek's strategic investments and a proactive regulatory environment, has positioned itself as a quantum-finance bridge for Asia. Israel has emerged as a quantum-software power, with Classiq's success demonstrating the country's capacity to scale deep-tech companies beyond seed and Series A. Other active clusters include Zurich and the Toronto–Waterloo corridor.

For venture investors, this geographic distribution carries practical implications. Deal sourcing in quantum finance requires a global network. The most promising companies may emerge from European physics laboratories, Israeli software incubators, or North American university spin-outs. Understanding the interplay between local research ecosystems, government industrial policy, and the proximity of financial end-users is essential to constructing a well-informed quantum-finance investment thesis.

## 23.6 In Conclusion, a Call to Action

As a practitioner operating at the intersection of finance and technology, I view quantum computing for financial services not as speculative futurism but as structured optionality. The opportunity rests on three pillars: finance’s inherent ability to monetise even marginal computational advantages; the direct alignment between quantum subroutines and core financial workflows in optimisation, simulation, and pricing; and the emergence of hybrid approaches that can deliver enterprise value today, well before fault-tolerant hardware arrives.

The companies profiled in this chapter, Multiverse Computing, SandboxAQ, QC Ware, Classiq, and Pasqal, illustrate the range of viable commercial strategies, from pure quantum-software platforms to AI-compression engines rooted in quantum mathematics to full-stack hardware providers with financial-sector applications. Execution, capital discipline, and the ability to navigate the long timeline between scientific potential and paying customers remain decisive factors.

The investor landscape is evolving in tandem. Specialist quantum funds such as Quantonation have been joined by sovereign wealth vehicles, corporate venture arms of major banks, and generalist technology investors. The entrance of NVIDIA, SoftBank, Goldman Sachs, and HSBC into the quantum-software space signals that the sector has crossed a threshold of institutional credibility.

The disruption horizon is not decades away. Over the next three to seven years, hybrid quantum workflows are likely to become embedded in specific, mission-critical functions within financial institutions, initially in optimisation and cryptographic migration, and progressively in risk simulation and derivative pricing as hardware improves. For venture capital funds, particularly those focused on fintech or deep tech, the mandate is clear: invest early in companies that solve current computational pain points while positioning themselves for a future in which quantum advantage is real and monetisable.

This journey requires patience, rigorous technical due diligence, and a genuine appetite for frontier technology. It also requires honesty about the risks: hardware timelines remain uncertain, commercialisation paths are long, and the history of deep-tech investing is littered with companies that were right about the science but wrong about the business. For those who approach the opportunity with clear eyes and sufficient conviction, however, the potential payoff is a transformative stake in the future of financial infrastructure. The time to build, partner, and invest is now.

# Chapter 24

## Post Quantum Cryptography

### *Securing today's data against tomorrow's quantum threats*

Santanu Ganguly

#### 24.1 Introduction

The author of this article started off his journey in quantum computing as a student when Caltech open-sourced John Preskill's notes on quantum computing circa 1999. Fast forward to 2019, and he was involved in early field testing of Post Quantum Crypto solution by a major vendor with the largest telecom group in the UK as a partner. Following that, his first paper on a cloud based roll out of hybrid classical-quantum crypto was published in January 2020 [14], to be followed by a book on quantum machine learning (QML), several patents on quantum tech, and research papers.

The author worked for many years in security, networking, automation and machine learning across Europe, with an academic background in Physics and Mathematics. This triggered a natural interest in security aspects and vulnerability of existing encryption systems surrounding the rise of quantum computers.

#### 24.2 Background

Encryption is used to protect the confidentiality of data. It takes the data as an input and in return outputs code or ciphertext that looks as though it has no meaning.

---

This chapter is a featured contribution to the monograph *A Portrait of Quantum Technologies in Finance*, edited by Oswaldo Zapata, PhD, and published by The Quantum Finance Boardroom (2026). All rights reserved by the chapter's author.

Only someone with a special key can decipher that code to obtain the original data, and only the intended recipient should have access to that special key. There are two forms of encryption algorithms: Symmetric Key and Asymmetric Key. To understand the difference, imagine the encryption of data to be equivalent to placing the data into a box and locking it, before sending it to the recipient. Therefore, decrypting data would be equivalent to unlocking the box and accessing the data inside.

Symmetric Key encryption is when the same key is used to lock and unlock that box. This is cost-effective (in terms of computation) but carries a risk. The key (called a Pre-shared Key) needs to be sent to the recipient so they can unlock the box, but we must ensure that it is not intercepted during transit. Any hacker who manages to obtain this key would be able to decipher the message and access confidential data. So, the key must be transmitted securely, just as securely as the locked box itself.

This is where Asymmetric Key Encryption is used. The symmetric key itself is encrypted, but this time, the key used to lock this box (called a Public Key) is different to the key used to unlock the box (called a Private Key). The public and private keys are associated to one another through a complex mathematical computation and only the public key is shared.

Today's most commonly used asymmetric key encryption algorithms are based on how difficult it is to break down large numbers into their prime factors. For example, we can easily verify that the multiplication of 3, 5, 11 and 17 is equal to 2805, but it is much harder to start with the number 2805 and break it down into its prime components. As the number gets larger, the problem becomes exponentially harder to solve, even for a computer. In fact, if the number is made up of at least 2048 bits, scientists believe it will take a standard computer so long to solve that it is practically impossible and so this problem is considered to be 'intractable'. Asymmetric key encryption is computationally expensive. It requires a lot of time, memory and processing power and so is not used to encrypt voluminous data itself, rather just the symmetric key that needs to be sent to the recipient.

However, the study of mathematics has evolved and new algorithms that can quickly reverse the factorising problem are being researched and developed. These algorithms would still require a greater amount of processing power than what is available today, but technology is quickly catching up and it is thought that a classical computer with enough power to run these new algorithms will exist in the near future.

Furthermore, a quantum computer would be able to solve this in just a matter of hours, thanks to the parallel processing capability discussed earlier. This would mean the symmetric key could be compromised, so all data that was encrypted using that key would also be compromised. Fortunately, a quantum computer with this level of processing power has not yet been built, and it is not expected to be around for at least another decade.

So, why are we worried? For the everyday user who makes online transactions (which use 2048-bit encryption), it is not a large reason for concern. If a transaction made today is decoded in twenty years' time, there is little chance of much damage being caused as a result. However, if we are sending and receiving confidential data that will still need to remain secure twenty years into the future, then it would definitely

be worth considering the use of post-quantum cryptography. And the trend of HNDL (harvest now, decrypt later) of encrypted data has been viably established via evidence across various geographical locations for the last six to seven years: there are entities who are actually storing encrypted data today, with hopes of decrypting them in future when computers become powerful enough to crack them. Keeping this in mind, organizations such as NIST (National Institute of Standards and Technology), NCSC (National Cyber Security Centre) et al., scientists and researchers have already placed resources into developing post-quantum encryption algorithms and exploring Quantum Key Distribution to deliver the symmetric key.

### 24.3 So, What Is This PQC Anyway?

**Post-Quantum Cryptography (PQC)** refers to a new generation of cryptographic algorithms designed to remain secure even against attacks from quantum computers. While today's widely used public-key systems, such as RSA and elliptic curve cryptography, are considered secure against classical computers, they are expected to become vulnerable once large-scale quantum computing becomes practical. PQC addresses this risk by introducing alternative algorithms for encryption, key exchange, and digital signatures that can withstand both classical and quantum attacks.

For organisations, PQC is not only a technical upgrade but also a strategic risk management priority. Sensitive data with long retention periods, such as financial records, personal information, government data, and intellectual property, may already be exposed to “harvest now, decrypt later” (HNDL) threats, where encrypted information is collected today and decrypted in the future when quantum capabilities mature. As a result, many organisations are beginning to assess their cryptographic landscape, build crypto inventories, and introduce crypto-agile architectures that allow algorithms to be replaced without major system redesign.

### 24.4 Is PQC the Holy Grail of Security Against All Quantum Attacks?

In reality PQC algorithms are purely classical, based on complex mathematical models that are supposed to be “quantum proof” due to their mathematical complexity. However, in reality, the attacks from a quantum computer, if and when that comes, will come from the domain of physics. Since, no known quantum attacks have been detected yet, it is hard to predict whether or not these classical PQC algorithms will actually be fruitful against such attacks.

As an example, several specific NIST post-quantum cryptography (PQC) candidates, including finalists like SIKE and Rainbow, were broken by researchers using conventional, non-quantum, present day ordinary computers. SIKE (Supersingular Isogeny Key Encapsulation) was broken in 2022 by researchers using a single-core PC in roughly 62 minutes, despite it being a top-round candidate for NIST. Rainbow signature scheme broken in February 2022 by Ward Beullens, who demonstrated a

practical key recovery attack on a laptop. However, the primary algorithms selected for standardization, such as CRYSTALS-Kyber, are perceived to be secure against future quantum threats, with NIST releasing the first finalized standards in August 2024.

The other option for security against quantum attacks is **Quantum Key Distribution (QKD)**. Instead of relying on today's complexity of a mathematical problem which could be overcome in the near future, Quantum Key Distribution (QKD) relies on the unchanging principles of quantum mechanics. It is based on the capability of manipulating a photon particle in order to encode data. A unique key code is generated and then shared only between the sender and the recipient. One thing that is quite unintuitive about a quantum particle is that it behaves differently when someone tries to observe the current state and differently when it is not being observed. This idea of being unable to observe the state of the particle is essential in QKD. If a hacker tries to intercept a key being transmitted via QKD, the state of the photon would change, so the recipient would know immediately that it has been tampered with and could terminate all communication. There are however some limitations to QKD, including the fact that it can only operate via satellite or fibre optics (which is what caught the interest of many service providers, such as BT). Whilst using a satellite can allow QKD over huge distances, transmission over commercial fibre keeping the coherence of information at an acceptable level, is currently limited to around 50km-60km; however, this distance can be extended through the use of specific optical repeaters.

Note: Whilst QKD sounds very promising, it is still in the stage of early trials and research, are not quite standardized for usage at a large scale (even though smaller scale, campus deployments are up and running in real world), and therefore the recommended direction, given by the National Cybersecurity Center UK, is as follows:

- Do not endorse QKD for any government or military applications
- Advise against replacing any existing public key solutions with QKD for commercial applications<sup>1</sup>

Despite this, there has been strong progress into the creation and testing of QKD networks around the globe. In Vienna, a QKD network has been designed and implemented and what is more is that China and Vienna have shared a key that was distributed via QKD-satellite. The UK has also made huge investments in implementing and testing QKD networks. The author of this article published a defensive article in January 2020 named "System and Architecture of a Quantum Key Distribution (QKD) Service over SDN" [19] along with a few patents on quantum security with quantum tech as well as with hybrid-quantum classical approaches [see references].

---

<sup>1</sup>Source: <https://www.ncsc.gov.uk/whitepaper/quantum-key-distribution>

## 24.5 PQC: Some Technical Overview

Post-quantum cryptography (PQC) is the transition from quantum-vulnerable public-key cryptography (notably RSA/ECC-based key establishment and digital signatures) to algorithms designed to resist future quantum attacks. As mentioned before, NIST approved the first three PQC standards in August 2024: FIPS 203 (ML-KEM), FIPS 204 (ML-DSA), and FIPS 205 (SLH-DSA), providing a standards-based foundation for enterprise planning and implementation.

NIST has announced 3 different standards:

The new accepted PQC standards are designed for two essential tasks for which encryption is typically used: general encryption, used to protect information exchanged across a public network; and digital signatures, used for identity authentication. NIST announced its selection of four algorithms — CRYSTALS-Kyber, CRYSTALS-Dilithium, SpHincS+ and FALCON — slated for standardization in 2022.

- **For general encryption**, used when users access secure websites, NIST has selected the CRYSTALS-Kyber algorithm. Among its advantages are comparatively small encryption keys that two parties can exchange easily, as well as its speed of operation.
- **For digital signatures**, often used when we need to verify identities during a digital transaction or to sign a document remotely, NIST has selected the three algorithms CRYSTALS-Dilithium, FALCON and SPHINCS+ (read as “Sphincs plus”). Reviewers noted the high efficiency of the first two, and NIST recommends CRYSTALS-Dilithium as the primary algorithm, with FALCON for applications that need smaller signatures than Dilithium can provide. The third, SPHINCS+, is somewhat larger and slower than the other two, but it is valuable as a backup for one chief reason: It is based on a different math approach than all three of NIST’s other selections.

Three of the selected algorithms are based on a family of math problems called structured lattices, while SPHINCS+ uses hash functions. The additional four algorithms still under consideration are designed for general encryption and do not use structured lattices or hash functions in their approaches.

While there have been no substantive changes made to the standards since the draft versions, NIST has changed the algorithms’ names recently to specify the versions that appear in the three finalized standards, which are:

- Federal Information Processing Standard (FIPS) 203, intended as the primary standard for general encryption. Among its advantages are comparatively small encryption keys that two parties can exchange easily, as well as its speed of operation. The standard is based on the CRYSTALS-Kyber algorithm, which has been renamed ML-KEM, short for Module-Lattice-Based Key-Encapsulation Mechanism.
- FIPS 204, intended as the primary standard for protecting digital signatures. The standard uses the CRYSTALS-Dilithium algorithm, which has been renamed ML-DSA, short for Module-Lattice-Based Digital Signature Algorithm.

- FIPS 205, also designed for digital signatures. The standard employs the Sphincs+ algorithm, which has been renamed SLH-DSA, short for Stateless Hash-Based Digital Signature Algorithm. The standard is based on a different math approach than ML-DSA, and it is intended as a backup method in case ML-DSA proves vulnerable.

Similarly, when the draft FIPS 206 standard built around FALCON is released, the algorithm will be dubbed FN-DSA, short for FFT (fast-Fourier transform) over NTRU-Lattice-Based Digital Signature Algorithm.

While the standard is in development, *NIST encourages security experts to explore the new algorithms and consider how their applications will use them, but not to bake them into their systems yet, as the algorithms could change slightly before the standard is finalized.*

*One of the most important considerations for PQC deployment is migration.* NIST NCCoE guidance emphasizes that migration to PQC from our existing encryption standards requires organizations to identify where quantum-vulnerable public-key cryptography is used across hardware, software, and services, and to build prioritized migration roadmaps. This is not a one-time library upgrade; it is a multi-year program involving discovery, architecture updates, interoperability testing, phased deployment, supplier coordination, and operational monitoring.

This engineering-focused report provides a practical framework with migration phases, test matrices, rollout guardrails, runbooks, KPIs, and dual-lens guidance for executives and engineering teams.

**Callout.** The decision is not whether PQC matters - it is whether the organization will execute a controlled, evidence-based migration early enough to avoid operational disruption later. NIST standards are available; planning and pilots should begin now.

**Technical Acumen.** Treat PQC as a platform modernization program: crypto inventory (CBOM), protocol support, interoperability testing, telemetry, rollout automation, and rollback readiness are as important as algorithm selection.

## 24.6 Why PQC Matters Now

### 24.6.1 Migration lead time is the immediate challenge

Even if the exact date of cryptographically relevant quantum computers remains uncertain, cryptographic transitions are slow. Organizations need time to identify cryptography embedded across third-party products, devices, protocol stacks, PKI, code-signing workflows, and customer/partner integrations. To this end, NIST has rolled out a migration guideline.

The dominant short-term risk for many enterprises is migration unreadiness rather than immediate cryptanalytic break. Delayed planning increases the chance of rushed deployments, compatibility failures, and uncontrolled exception sprawl.

**Callout.** Migration unreadiness is a business resilience risk. Budgeting for discovery, testing, and supplier engagement early reduces future operational disruption.

## 24.6.2 Harvest Now, Decrypt Later (HNDL) exposure

Adversaries may collect encrypted traffic or archives today and attempt decryption later if the underlying public-key mechanisms become vulnerable. This is especially relevant for long-lived sensitive data, regulated records, intellectual property, and software trust chains.

## 24.6.3 What is most exposed first

PQC migration primarily targets public-key mechanisms used for key establishment and digital signatures. Typical high-priority domains include TLS/VPN handshakes, PKI and certificates, code signing, firmware signing, secure boot, and identity/authentication systems.

Although some symmetric cryptography such as RSA below 2048 and ECC are deemed vulnerable, the overall domain remains comparatively resilient in the near term when appropriately parameterized, but policy and architecture reviews and updates are still needed.

# 24.7 NIST PQC Standards and How They Work

## 24.7.1 Current NIST baseline

NIST approved FIPS 203 (ML-KEM), FIPS 204 (ML-DSA), and FIPS 205 (SLH-DSA) in 2024. FALCON remains selected for future standardization (FIPS 206 in development).

## 24.7.2 Naming translation for teams

- CRYSTALS-Kyber → ML-KEM
- CRYSTALS-Dilithium → ML-DS
- SPHINCS+ → SLH-DSA (standardized form)
- FALCON → future FIPS 206 (in development)

Organisations use the FIPS names consistently in architecture reviews, procurement, and audit artifacts, reducing confusion when teams reference older competition-era names. NSA's CNSA FAQ explicitly clarifies that ML-KEM and ML-DSA are the fully specified NIST standards corresponding to those earlier submission names.

### 24.7.3 ML-KEM (FIPS 203): conceptual behavior and deployment relevance

ML-KEM is a key encapsulation mechanism (KEM). At a high level, one party generates a KEM public/private key pair; the peer encapsulates a shared secret to the public key, producing a ciphertext and shared secret; the private key holder decapsulates and derives the same shared secret. The resulting shared secret is used to derive symmetric session keys.

- Deployment relevance: TLS/QUIC-like session establishment
- VPN key establishment
- Service-to-service secure channels
- Secure messaging/session setup

**Technical Acumen.** KEM integration changes more than algorithm names. It can affect handshake payload sizes, CPU cost, retry/timeout behavior, and observability requirements.

### 24.7.4 ML-DSA and SLH-DSA

ML-DSA (FIPS 204) and SLH-DSA (FIPS 205) are PQC digital signature standards used for integrity, authenticity, and non-repudiation use cases, subject to ecosystem support. Typical uses include code signing, signed metadata, firmware signing, and identity artifacts.

### 24.7.5 Standards and deployability

Production rollout depends on standards, protocol support, libraries, HSM/TPM support, certificate, client/server compatibility, and operations readiness. Fig. 1 below shows a protocol level conceptual schematic for KEM-based processes.

## 24.8 Engineering Migration Roadmap (3–5+ Year Program)

A practical transition to post-quantum security is typically phased. It begins with discovery and risk assessment, followed by pilot implementations using hybrid cryptography, where classical and post-quantum methods are used together, and then expands into broader enterprise adoption. In regulated sectors such as finance and government, PQC is increasingly viewed as a foundational component of long-term cyber resilience, Zero Trust architecture, and compliance readiness.

Any upgrade to PQC from existing crypto will require thorough preparation and diligent planning. The exact timeline depends on sector, legacy footprint, and vendor readiness. Public milestone models such as the UK NCSC's 2028/2031/2035 guidance

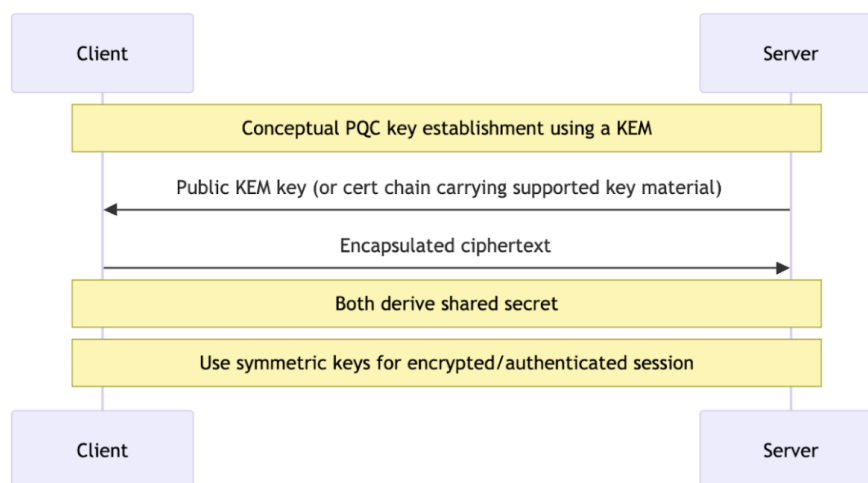


Figure 24.1: High-level conceptual flow for KEM-based key establishment. This is an explanatory schematic, not a protocol specification.

can be used as external planning anchors while preserving an enterprise-specific phased roadmap.

### 24.8.1 Crypto-agility

Prior to upgrading to PQC, it is strongly advised that the systems go through a thorough cryptographic inventory and define protocols for crypto-agility. **An crypto agile system** is one that automatically adjusts to any security threats or compliance lapse in the digital environment, as for example, may happen during mergers of two companies with two different security standards.

**Crypto-agility** is formally defined as the ability of a system to adapt its cryptographic mechanisms without requiring major redesign, disruption, or prolonged downtime. As cryptographic standards evolve and new threats emerge, organisations need architectures that can replace algorithms, keys, certificates, and protocols quickly and safely. The core design principle of crypto-agility is to avoid hard-coding specific algorithms into business applications. Instead, cryptographic functions should be abstracted through reusable services, APIs, or policy-driven platforms so that encryption, signing, and key exchange methods can be updated with minimal impact on the wider system. This kind of policy enforcement is of optimal importance in environments that contain sensitive data, such as Financial organisations or Federal government.

*Effective crypto-agility also depends on centralised governance and clear policy control.* Cryptographic choices should be defined by enterprise policy rather than individual developers or isolated teams. This includes approved algorithms, key sizes, certificate profiles, rotation schedules, and deprecation rules. Strong key and trust management is equally important, with lifecycle controls for key generation, storage, rotation, revocation, and destruction. Systems should also be designed for interoperability and phased migration, allowing classical and next-generation cryptography to coexist

during transition periods. This is especially important in post-quantum migration, where hybrid approaches help maintain compatibility while reducing future risk.

Finally, *crypto-agile systems must be observable, testable, and auditable*. Organisations need visibility into where cryptography is used, which algorithms are active, and which systems are exposed to legacy or have weakened controls. Logging, inventory, and continuous monitoring make it possible to assess risk and respond quickly when standards change. In this sense, crypto-agility is not only a technical design choice but a resilience strategy: it enables security teams to respond to new threats, regulatory requirements, and technology shifts without rebuilding the enterprise each time cryptography needs to evolve. A general phased roadmap for a PQC migration is shown in Fig. 2 below.

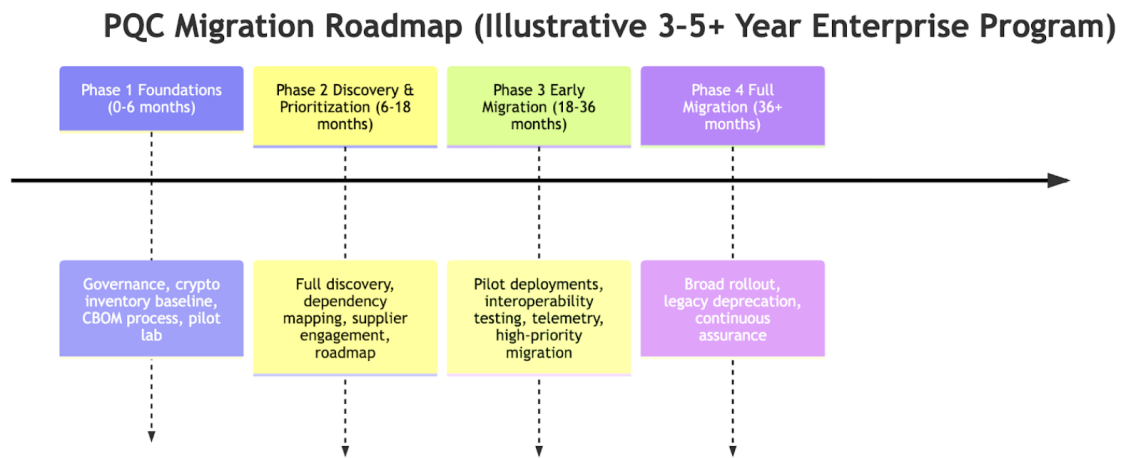


Figure 24.2: Example phased migration roadmap aligned to enterprise delivery cadence, dependent on timelines to asset complexity and supplier readiness.

Fig. 3 below depicts a broader view of phased PQC roll-out as per NIST guidance




Figure 24.3: Illustrative enterprise roadmap for phased PQC migration. Actual timelines depend on asset complexity, ecosystem readiness, and vendor support.

## 24.9 Machine Learning (ML) for PQC Deployment, Monitoring, and Capacity Control

ML is most valuable in operational support for PQC migration: discovery and inventory enrichment, telemetry anomaly detection, compatibility prediction, rollout prioritization, and capacity forecasting. This is focused on solving one major problem in today's networks: Crypto usage is scattered across code, configs, appliances, and managed services.

Machine learning can play an important supporting role in Post-Quantum Cryptography (PQC), not by replacing cryptographic standards, but by helping organisations manage the complexity of discovery, migration, and continuous monitoring. As enterprises prepare for the transition to quantum-resistant algorithms, one of the biggest challenges is identifying where current cryptography is used across applications, infrastructure, certificates, protocols, and data flows. Machine learning can assist by analysing code, configurations, network telemetry, and security logs to detect cryptographic usage patterns, classify algorithms, and build a more complete cryptographic inventory. This helps organisations understand which systems rely on vulnerable classical cryptography and where post-quantum migration should be prioritised.

Machine learning is also valuable in risk assessment and operational monitoring. By correlating factors such as data sensitivity, exposure level, retention period, and current algorithm strength, ML models can help score systems according to post-quantum migration urgency. During rollout, machine learning can support anomaly detection by identifying unusual handshake failures, compatibility issues, or unexpected changes in cryptographic behaviour across hybrid environments where classical and post-quantum methods coexist. However, its role should remain carefully bounded: machine learning should inform discovery, prioritisation, and monitoring, while final cryptographic decisions remain governed by approved policy, standards, and human review. In this way, machine learning strengthens PQC programmes by improving visibility, scalability, and resilience without replacing formal security controls. Figure 4 below shows ML-Enabled PQC Rollout Control Loop.

In summary, main ML use cases for PQC involve the following:

- Repository/config classification for crypto indicators
- Protocol/cipher pattern detection in telemetry
- Certificate metadata clustering
- NLP on tickets/docs for hidden dependencies
- Improved CBOM completeness
- Reduced manual discovery time
- Audit evidence generation
- Capacity planning
- Engineering

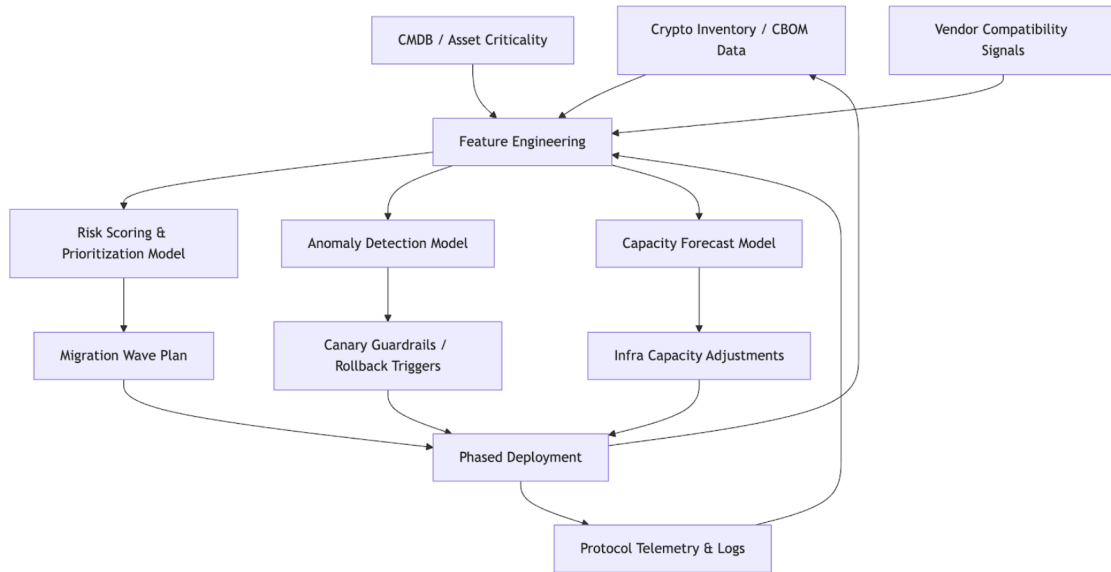


Figure 24.4: Example control loop showing how classical ML can support PQC migration planning and safe phased rollout operations.

### 24.9.1 Performance and implementation considerations (including GPU acceleration/cuPQC)

#### Why performance matters

PQC migration is not only a security project. It can affect handshake CPU costs, connection setup latency, message sizes, memory usage, hardware acceleration assumptions just to name a few. This is where NVIDIA’s PQC stack becomes absolutely critical.

NVIDIA’s cuPQC SDK documentation describes it as a high-performance SDK for cryptographic applications on GPUs, and its feature docs position cuPQC-PK as GPU-accelerated implementations of NIST-standardized PQC algorithms with batched operations and device-side integration. Following defines the way ecosystem is adapting to performance concerns:

- high-throughput services,
- cloud-scale termination workloads,
- telecom / infra-heavy deployments,
- research benchmarking.

GPU acceleration (e.g., NVIDIA cuPQC) is one emerging option for organizations that need high-throughput PQC operations or batched cryptographic workloads, but it should be treated as an optimization path rather than a prerequisite for migration.

## 24.10 QML for PQC: Opportunities and Limits

PQC migration today is driven by standards, software engineering, interoperability, and operations. Even though officially quantum machine learning (QML) is treated as an R&D track, not a prerequisite for enterprise PQC migration, research in the area has picked up considerably in recent times. There has been some works on quantum-enhanced intrusion detection and PQC, albeit, existing studies remain largely performance-focused, with limited analysis of geometric interpretability, stability, and operational robustness.

The author of this article published a book in 2021 called “Quantum Machine Learning: An Applied Approach”, addressing general applied QML as well as security aspects that can be addressed via such methods. Since then research has picked up and QML for cybersecurity and PQC is a popular area for research. An example is a recent focus of the author where a hybrid classical-quantum autoencoder (HCQAE) architecture was investigated to analyse CIC DDoS data from University of New Brunswick.<sup>2</sup>

Anomaly detection [15, 16, 17, 18] is central to cybersecurity and cyber-physical systems, where rare events must be identified without labelled attack data. Classical autoencoders detect anomalies via reconstruction error and quantum variants have been explored through variational classifiers and optimization-based methods [16, 17, 18].

Quantum anomaly detection remains comparatively underexplored. Existing work includes quantum kernels, variational classifiers, and generative quantum models for detecting distributional shifts. Yet it is largely unresolved whether anomalies appear as structured displacement in quantum state space or whether apparent gains reflect unstable dynamics rather than robust representations. These are exciting areas where QML can be very useful in recent future given NVIDIA’s focus on making quantum AI simulation reachable to industry via HPC boost.

### Plausible applications of QML for PQC:

1. Optimization experiments for rollout sequencing
2. Anomaly detection research on specialized telemetry representations
3. Benchmarking hybrid classical/quantum workflows

**Callout.** QML should be funded as innovation if it aligns with R&D strategy, but do not gate PQC migration readiness on it.

## 24.11 Risks, Pitfalls, and Recommendations

In every technology roll-out, more so for new technologies, there are always risks and caveats to be aware of.

<sup>2</sup><https://www.unb.ca/cic/datasets/ddos-2019.html>

### 24.11.1 Common mistakes to avoid

1. Treating PQC as a simple library upgrade
2. Waiting for “all standards to be final everywhere”: NCSC and NSA both indicate the transition is multi-year and requires early planning, not last-minute action.
3. Starting pilots without a crypto inventory: No crypto inventory means no CBOM. To prioritize migration, identification of where vulnerable cryptography lives is of paramount importance.
4. Skipping interoperability and telemetry testing
5. Using hybrid as a permanent architecture by default: Hybrid approach can help transition, but NSA notes the complexity and extra long-term transition burden it can introduce.
6. Ignoring supplier and partner dependencies.
7. Lacking rollback criteria and rehearsed runbooks.

## 24.12 Conclusion

PQC migration is now a practical engineering and governance program, not a speculative future exercise. NIST’s approval of FIPS 203/204/205 provides a standards baseline for planning and pilots, while NIST NCCoE and NCSC guidance make clear that early discovery, interoperability testing, and phased execution are essential.

For executives, success means funding and governing a crypto-agility transformation. For engineering teams, success means delivering inventory, testability, telemetry, staged rollout control, and reliable rollback. ML can materially improve migration operations; QML remains exploratory.

### References

1. NIST CSRC, ‘Post-Quantum Cryptography FIPS Approved’ (FIPS 203/204/205 approval, Aug. 2024).
2. NIST CSRC, ‘Post-Quantum Cryptography Standardization’ project page (includes FIPS 206 status for FALCON).
3. NIST FIPS 203, Module-Lattice-Based Key-Encapsulation Mechanism Standard (ML-KEM), 2024.
4. NIST NCCoE, ‘Migration to Post-Quantum Cryptography’ project guidance and migration considerations.
5. UK NCSC, ‘Timelines for migration to post-quantum cryptography’ (public milestones and planning guidance).
6. NSA, ‘Commercial National Security Algorithm Suite 2.0 (CNSA 2.0) FAQ’ (transition and hybrid considerations).

7. NIST CSRC, Digital Signatures project page (digital signature roles and PQC signature standards references).
8. NVIDIA, cuPQC SDK documentation (example acceleration ecosystem tooling for PQC-related workloads).
9. “CUDA-Q Assisted Hybrid Classical-Quantum Autoencoders for Anomaly Detection in multi-class Colour and Greyscale Medical Images”, Ganguly, S., Liang, X., & Makris, D. (2026). To be Published.
10. “Bosonic Interference Reveals Spectral Structure in Graph-Coupled Quantum Learning”, Ganguly, S., Liang, X., & Makris, D. (2026). To be Published.
11. “Hybrid Classical-Quantum Generative Algorithms for Financial Modelling and Prediction”, Ganguly, S., Liang, X., & Makris, D. (2024). IEEE International conference on “Intelligent Control, Computing and Communication (ICCC-2025)”
12. “Design of a Quantum Machine Learning Course for a Computer Science Program”, Sathish Kumar, Temitope Adeniyi, Ahmad Alomari, and Santanu Ganguly, IEEE QCEE 2023.
13. “Classification of NEQR Processed Classical Images using Quantum Neural Networks (QNN)”, 2022, arXiv:2204.02797 [quant-ph].
14. “System and Architecture of a Quantum Key Distribution (QKD) Service over SDN”, Santanu Ganguly, January 2020
15. A. Gong et al., “Memorizing Normality to Detect Anomaly: Memory-augmented Deep Autoencoder for Unsupervised Anomaly Detection,” arXiv:1904.02639 (2019),  
<https://doi.org/10.48550/arXiv.1904.02639>
16. Z. Chen, C. Yeo, B. Lee, and C. T. Lau, “Autoencoder-based network anomaly detection,” in 2018 Wireless Telecommunications Symposium (WTS), 1–5 (2018),  
<https://doi.org/10.1109/WTS.2018.8363930>
17. V. Bergholm et al., “PennyLane: Automatic differentiation of hybrid quantum-classical computations,” arXiv:1811.04968 (2018),  
<https://doi.org/10.48550/arXiv.1811.04968>
18. J. Basit, D. Hanif, and M. Arshad, “Quantum Variational Autoencoders for Predictive Analytics in High Frequency Trading Enhancing Market Anomaly Detection,” *Int. J. Emerg. Multidisciplinaries: Comput. Sci. Artif. Intell.* 3(1), 21 (2024),  
<https://doi.org/10.54938/ijemdc sai.2024.03.1.319>
19. “Quantum Machine Learning: An Applied Approach”, Springer Nature, August 2021

## Patents

1. US Patent 12362921: “Systems and methods for providing user authentication for quantum-entangled communications in a cloud environment”, Santanu Ganguly and Brice Achkir
2. US Patent 12192344: “Systems and methods for providing dynamic quantum cloud security through entangled particle distribution”, Santanu Ganguly and Brice Achkir
3. US Patent 11818257: “Systems and Methods for Providing User Authentication for Quantum-Entangled Communications in a Cloud Environment”, Santanu Ganguly and Brice Achkir
4. US Patent 11716151: “Routing Methods for Quantum Communication Paths across a Mesh Quantum Network”, Luca Della Chiesa, Sam Samuel, Paul Polakos, Scott Fluhrer, Santanu Ganguly
5. “Controlling Quantum Communication via Quantum Memory Management”, Louis Gwyn Samuel, Santanu Ganguly, Maria Gragera Garces and Luca Della Chiesa

## Chapter 25

# Quantum Readiness for Financial Adoption

*From systematic alpha to quantum probability: a personal and technical roadmap for the modern CTA*

Sebastian Torres

The global asset management industry stands at a critical juncture. For the past two decades, the pursuit of “Alpha” has been largely defined by the speed of information acquisition and the refinement of classical statistical models. We have optimized the latency of our execution and the precision of our linear regressions to their absolute asymptotic limits. However, the fundamental mathematical challenges of finance, specifically combinatorial portfolio optimization, stochastic volatility forecasting, and systemic risk simulation, remain computationally irreducible under the current silicon- based paradigm.

Drawing upon my personal journey from analyzing emerging market volatility in Colombia to developing advanced quantitative models in North America, this paper posits that the financial industry has hit a “Classical Ceiling” The heuristics we rely on to manage portfolios are essentially approximations of a reality that is too complex for classical bits to model perfectly.

Using my operational experience in developing the PRO Systematic Capital Management program as a primary case study, this White Paper outlines a pragmatic framework for the adoption of Quantum Computing (QC) in institutional finance. It moves beyond the hype of “quantum supremacy” to focus on “quantum utility”, the specific and narrow applications where quantum algorithms can solve optimization and simulation problems that are currently intractable. I argue that the transition to

---

This chapter is a featured contribution to the monograph *A Portrait of Quantum Technologies in Finance*, edited by Oswaldo Zapata, PhD, and published by The Quantum Finance Boardroom (2026). All rights reserved by the chapter’s author.

quantum readiness is not merely a hardware upgrade, but a fundamental reimagining of risk from a deterministic variable to a probabilistic state.

## 25.1 Origins: The Deterministic Illusion in a Stochastic World

My personal fascination with quantum computing did not originate in a theoretical physics laboratory. It was born out of profound, late-night frustration on the trading floor.

Growing up and studying in Bogotá, Colombia, I was surrounded by the realities of an emerging market. During my undergraduate studies in Accounting, Finance, and Business Administration at the University of the Andes, I was taught the classical, elegant theories of modern finance. I learned the Markowitz Mean-Variance optimization, the Black-Scholes options pricing model, and the Capital Asset Pricing Model. On the whiteboard, these equations were perfect.

However, when I began my career as a Risk Management Consultant and later as a Quantitative Risk Analyst at StoneX, the illusion of deterministic finance shattered. My daily reality involved analyzing the Colombian Peso (USDCOP), agricultural commodities like Palm Oil, and forecasting volatility using GARCH models. I quickly learned that the elegant models I studied in university often failed to capture the chaotic, fat-tailed reality of the actual markets. In emerging markets, correlations do not behave normally. A sudden geopolitical shift or a liquidity crisis can cause previously uncorrelated assets to move in perfect, destructive lockstep.

This realization, that traditional financial math was insufficient for the real world, was the catalyst for my transition into data science. I knew that if the classical financial formulas were failing, I needed better computational tools. This drive pushed me to continuously upskill. I moved beyond Excel and Stata, diving deeply into Python, R, and SQL. I pursued the Applied Data Science Program at MIT and the Machine Learning Specialization at Stanford University online, spending countless hours pushing code to my GitHub repository and learning how to build neural networks that could handle market noise.

For a time, these advanced machine learning tools felt like the ultimate answer. I could process multi-formatted client data through automated ETL pipelines. I could utilize Python libraries like Pandas, NumPy, and Scikit-learn to build predictive models that far outperformed traditional regressions. But as my responsibilities grew, so did the size of the datasets and the complexity of the mandates. Eventually, I realized that even with advanced machine learning, I was simply running headfirst into the physical limits of classical computing. I was no longer limited by my understanding of the market, I was limited by the processing power of silicon.

## 25.2 The Classical Apex: A Case Study of PRO Systematic CM

To truly understand my personal pivot toward quantum mechanics, we must analyze the specific environment where I hit the “Optimization Wall.” This occurred during my tenure developing and managing the PRO Systematic Capital Management program, particularly our quantitative strategies focused on capital preservation and long-term equity growth.

As Chief Investment Officer, my mandate was to deliver performance that could weather market downturns while capturing upside volatility. The PRO Systematic CM program was designed to operate as a highly diversified Commodity Trading Advisor strategy. We utilized rigorous, data-driven quantitative models to capture trends across a vast universe of equities, FX, and commodities.

The success we achieved was significant. By leveraging advanced classical statistics, we achieved an 18% increase in portfolio growth and maintained risk within approved Value at Risk levels 95% of the time. We employed a multi-frequency approach, trading across different time horizons to smooth out the equity curve. By utilizing complex correlation matrices, we ensured that the portfolio remained diversified, achieving a 30% improvement in diversification metrics through strategic, algorithmic rebalancing.

Yet, behind the strong performance metrics, I was fighting a constant, invisible battle against computational latency.

While the PRO Systematic strategy was robust, it was strictly bound by the limitations of linear algebra and classical optimization solvers. For example, finding the optimal set of parameters for our trend-following algorithms required a massive “brute force” computational approach. We had to test millions of combinations of look-back periods, volatility filters, and entry thresholds. This process was incredibly computationally expensive. I would set Python scripts to run overnight, hoping that by the morning, the classical solver had found a viable parameter set that wasn’t completely overfitted to the historical data.

Furthermore, the portfolio rebalancing logic, while highly sophisticated, fundamentally relied on historical correlation data. As I had learned years earlier in Colombia, in moments of extreme market stress, these correlations have a dangerous tendency to converge to one. The classical algorithms we used for execution had to make mathematical compromises to run in real-time, simplifying non-linear market impacts into linear approximations just so the computer could process the trade before the price moved.

The PRO Systematic CM program stands as a testament to the absolute power of data-driven decision making. It represents the pinnacle of what a CTA can achieve with classical beta. However, it also serves as my personal control group for the quantum experiment. Sitting in front of those models, watching the CPU usage spike to 100% while the algorithm struggled to find the global minimum of a risk function, I realized that the future of finance required a completely new architecture. This is

where my personal passion for quantum algorithms was ignited.

## 25.3 The Combinatorial Wall: Rethinking Portfolio Construction

The central mandate of any portfolio manager is the efficient allocation of capital. On the surface, this appears to be a straightforward, linear problem, buy the assets with the highest expected return and the lowest correlation to the existing book.

In practice, portfolio construction is a “Combinatorial Optimization” problem of staggering, mind-bending complexity. When managing a portfolio of over 100 equities, the number of possible portfolio combinations exceeds the number of atoms in the visible universe.

As we introduce necessary real-world constraints, the mathematical difficulty explodes exponentially. Consider the standard institutional mandates we managed,

1. Cardinality Constraints, the fund must hold a specific, limited number of stocks to minimize operational drag.
2. Sector Neutrality, no single sector can exceed a certain percentage of the total weight to prevent concentration risk.
3. Transaction Costs, turnover must be minimized to preserve net asset value.
4. Liquidity Thresholds, the fund cannot hold positions that are too large relative to the average daily trading volume of the asset.

Mathematically, adding these constraints transforms a simple quadratic problem into a Mixed-Integer Quadratic Programming problem. In computer science terminology, this is classified as “NP-Hard.” Classical computers simply cannot solve this perfectly. They cannot check every valid combination to find the absolute global minimum of risk. Instead, they must make compromises, utilizing heuristic algorithms like Simulated Annealing or Genetic Algorithms to find a “local minimum”, a solution that looks good compared to its immediate neighbors, but might be vastly inferior to the true, undiscovered optimal solution.

This specific bottleneck is the primary reason I began studying quantum circuits. It represents the ultimate “Killer App” for quantum finance, the Quantum Approximate Optimization Algorithm (QAOA).

QAOA is specifically designed to tackle these exact types of combinatorial problems. By mapping the complex portfolio constraints to a Hamiltonian, which is essentially an energy function representing the cost or risk of the portfolio, a quantum processor can explore the entire vast solution space simultaneously. Unlike classical solvers that get stuck in the valleys of a local minimum, quantum tunneling allows the algorithm to pass directly through energy barriers to locate the true global minimum.

For a program like PRO Systematic, the future implementation of QAOA would mean the ability to rebalance a massive portfolio dynamically with absolute mathematical certainty. It would completely eliminate the “optimization error” that currently plagues the entire asset management industry, potentially adding significant, risk-free basis points to the annual return simply by operating at maximum mathematical efficiency.

## 25.4 The Latency of Risk: Accelerating Monte Carlo Simulations

If portfolio optimization is the offensive engine of asset management, risk management is the defensive shield. Throughout my career, whether I was consulting in Miami, working remotely for StoneX, or executing trades based in Vancouver, the primary tool for assessing risk in complex derivatives and foreign exchange books has been the Monte Carlo simulation.

During my work managing a massive \$170 million USD annual hedging policy for Uniban SA, the largest agricultural company in Colombia, the stakes were incredibly high. We were dealing with real-world, physical supply chains heavily exposed to currency fluctuations. To calculate the Value at Risk and the Expected Shortfall of the currency book, we had to run tens of thousands of random walk simulations. We would simulate the potential future path of the Peso or the Euro 10,000 times to see exactly how often the portfolio would lose money under extreme conditions.

The fatal limitation of classical Monte Carlo simulations is their painfully slow convergence rate. To improve the accuracy of a simulation by a factor of 10, you must increase the number of simulations by a factor of 100. This is known as the square root law.

In a high-frequency trading environment, or during a highly volatile market event, this latency is completely unacceptable. By the time a classical server farm has crunched 100,000 simulations to give a precise risk number to the Chief Risk Officer, the market data is already stale. You are effectively driving a high-speed vehicle while exclusively looking in the rearview mirror.

This operational anxiety led me to research Quantum Amplitude Estimation (QAE). QAE offers a theoretical quadratic speedup over classical Monte Carlo methods. Instead of scaling at the inverse square root of  $N$ , it scales at the inverse of  $N$ .

While this sounds highly abstract, the operational impact on a trading desk is profound. It means a risk engine could potentially run a simulation with the precision of 1,000,000 scenarios in the exact same time it currently takes to run 1,000.

For the industry, this unlocks real-time, intra-day risk monitoring. Funds could move away from slow, overnight batch processing. Furthermore, complex exotic options, like the deeply structured products we utilized to hedge client exposure at StoneX, could be priced with the exact same speed as simple vanilla calls and puts.

This does not just make the risk management process faster, it makes the entire use of capital more efficient. If a firm can calculate its risk with a much higher degree of precision, it requires a much smaller capital buffer, or margin, to cover its uncertainty. Quantum computing, in this sense, directly unlocks trapped liquidity.

## 25.5 Beyond Correlation: Quantum Machine Learning for Hedging

Standard hedging policies rely heavily on linear correlations. We analyze historical data to determine that one asset is 80 percent correlated with another, and we build our hedge ratios accordingly. We utilized this approach extensively, and successfully, to reduce foreign exchange exposure by 75 percent.

However, the failure mode of this classical approach is well documented. In a crisis, historical correlations fail entirely. The linear regression models that worked beautifully in a low-volatility bull market provide absolutely zero protection in a sudden market crash.

Classical machine learning models, like the Random Forests, Gradient Boosting, or Scikit-learn models I have built and validated, struggle immensely when the number of variables becomes too high. If we attempt to model the simultaneous interaction of global interest rates, geopolitical sentiment, supply chain logistics, and currency flows, classical models almost always succumb to the curse of dimensionality, leading to severe overfitting.

Quantum Machine Learning introduces the revolutionary concept of the “Quantum Feature Space.” A quantum computer has the unique ability to map classical data into a high-dimensional Hilbert space that is exponentially larger than what any classical supercomputer can manage.

In this vastly expanded space, non-linear relationships, patterns that look like completely random noise to a classical computer, suddenly become linearly separable and highly predictable.

A quantum model could potentially identify “market regimes”, such as the subtle, underlying shift from low-volatility inflation to high-volatility stagflation, much earlier than any classical indicator by detecting microscopic shifts in the non-linear interactions of thousands of assets. Instead of relying on a static hedge ratio, Quantum Machine Learning could enable a dynamic, probabilistic hedge that adjusts in real-time to the changing state of the market, offering a level of protection that remains robust even during unprecedented Black Swan events.

## 25.6 The Prerequisite of Data Hygiene

It is imperative to ground this forward-looking discussion in the deeply unglamorous reality of data engineering. My recent experience functioning as the key business

analyst at Citiloc Systems in Vancouver serves as a critical, grounding case study for quantum readiness.

At Citiloc, I was tasked with analyzing historical payment patterns to devise a data-driven collections strategy. I utilized Python to build and validate a payment-delinquency model that achieved a 95 percent accuracy rate, resulting in a 92 percent reduction in aged accounts receivable. However, we achieved this high accuracy not because the mathematical model was incredibly complex, but because of the extreme rigor we applied to the data cleaning process. I spent weeks normalizing historical payment patterns, removing outliers, formatting timestamps, and structuring the raw inputs.

We are currently operating in the NISQ era, which stands for Noisy Intermediate-Scale Quantum. Quantum qubits are incredibly fragile entities. They interact with their surrounding environment and lose their quantum state, a process known as decoherence, within fractions of a microsecond. If we feed “noisy” financial data, data with missing values, extreme outliers, or inconsistent formats, into a “noisy” quantum computer, the result is nothing but pure, uninterpretable static. In the classical world, bad data is a nuisance that we can brute-force our way through. In the quantum world, bad data is a complete showstopper.

Therefore, my most urgent advice to the Boardroom is that the path to Quantum Readiness does not actually begin with quantum physics, it begins with classical Data Engineering. The automated ETL processes I developed to standardize multi-formatted client data at StoneX, and the data validation pipelines I built at Citiloc, are the exact prototypes required for Quantum State Preparation. Financial institutions must invest heavily in cleaning, structuring, and maintaining their data lakes right now. When fault-tolerant quantum computers finally become commercially viable, the firms with pristine, “Quantum-Ready Data” will be able to deploy algorithms immediately, while their competitors will be forced to spend years scrubbing their databases.

## 25.7 The Hybrid Operating Model and the Role of the Translator

The transition to quantum finance will not be a sudden, binary switch. We are not going to unplug our CPU-based servers on a Friday and plug in a Quantum Processing Unit on a Monday. The future of quantitative finance is undeniably Hybrid.

Just as modern computers use a standard CPU for general logic and a specialized GPU for graphics rendering, the financial infrastructure of the future will utilize a QPU as a highly specialized co-processor.

The classical CPU will continue to handle data ingestion, user interfaces, Power BI dashboards, and simple linear regressions, tasks where classical computing already excels. The QPU will be called upon entirely and exclusively for the specific, intractable bottlenecks, the combinatorial optimization of the PRO Systematic

portfolio, or the high- fidelity Monte Carlo simulation of a massive currency book.

This complex hybrid architecture requires a completely new breed of financial professional. It requires what I call a “Quantum Translator.”

My personal career path, moving from studying accounting in Colombia to executing futures trades in Miami, and finally building machine learning models in Canada, has been defined by translation. I have constantly had to translate between languages, Spanish, English, and Italian, and more importantly, I have had to translate complex stakeholder needs into rigid technical requirements.

The industry currently has brilliant physicists who do not understand market liquidity or margin calls, and brilliant traders who do not understand superposition or quantum entanglement. The future belongs to the practitioners who can bridge this massive gap. We need analysts who understand the deep business logic of a trade, the regulatory requirements of the NFA or FINRA, the mechanics of settlement, and who can also formulate that exact financial problem into a mathematical Hamiltonian that a quantum algorithm can process. This is the exact intersection where my personal interests and my professional expertise now converge.

## Conclusion: The Asymptotic Edge

The financial industry has pushed classical computing to its absolute, physical limit. We have squeezed every conceivable drop of alpha out of Scikit-learn, out of GARCH volatility models, and out of Mean-Variance Optimization. We have reached the final asymptote of what is possible with binary logic.

The ongoing success of programs like PRO Systematic CM proves beyond a doubt that we know exactly how to build robust, intelligent strategies. We possess the financial logic. We possess the historical data. What we fundamentally lack is the underlying physics to optimize these strategies perfectly in real-time.

My journey from the University of the Andes to the forefront of quantitative data analytics has been a continuous search for better, more accurate ways to map the chaos of the financial markets. Quantum Computing offers the ultimate ladder to climb over the classical wall. It promises a paradigm shift from heuristic, “good enough” approximations to probabilistic, mathematical certainties. It offers the unprecedented ability to see non-linear risks that are currently completely invisible, and to allocate capital with a level of efficiency that is currently mathematically impossible.

For the members of the Quantum Finance Boardroom, the ultimate message is clear. The operational risk is no longer in adopting this cutting-edge technology too early. The true, existential risk is in continuously optimizing for a classical world that is rapidly, and permanently, becoming obsolete.

# Chapter 26

## Quantum Neural Networks

### *Application for credit risk assessment and fraud detection*

Sebastian Zajac & Krzysztof Kuba

#### 26.1 Introduction

Quantum computers have introduced a new paradigm in computing, offering transformative potential across various domains such as physics, chemistry, logistics, and finance [1-3]. Within the financial sector, one of the most critical challenges remains the accurate assessment of credit risk [4].

Formally, credit risk is defined as the probability that a debtor will fail to fulfill their obligations to a financial institution. Precise estimation of this risk is fundamental to the stability of the banking sector and the effective allocation of capital within the economy. Since the emergence of the modern banking system—and particularly in the wake of numerous financial crises in the 20th and 21st centuries—credit risk modeling has become a cornerstone of economic, financial, and regulatory research. Furthermore, regulatory frameworks such as Basel II and III have established the importance of quantitative risk measures and standardized the approaches used by financial institutions worldwide [5, 6].

Modern methods of credit scoring are based on advanced statistical techniques and machine learning models, which allow for the processing of huge volumes of data coming from disparate sources: financial, behavioral, transactional, and macroeconomic. The primary objective of these models is not only to classify

---

This chapter is a featured contribution to the monograph *A Portrait of Quantum Technologies in Finance*, edited by Oswaldo Zapata, PhD, and published by The Quantum Finance Boardroom (2026). All rights reserved by the chapter's authors.

clients according to their credit risk profiles but also to gain deeper insights into the underlying factors influencing default. Traditionally, classical credit risk models have relied on logistic regression, decision trees, random forests, and deep neural networks. In this chapter, we focus our attention on the fundamental architecture of neural networks, which serves as a necessary precursor to understanding their quantum counterparts: quantum neural networks.

The efficacy of these models is heavily contingent upon the availability and quality of data, as well as the underlying computational capacity. In modern financial institutions, datasets frequently scale to hundreds of millions of records and tens of thousands of variables, necessitating advanced parallel computing and cloud infrastructure. Yet, as financial markets grow in complexity, so too does the demand for model sophistication, leading to a significant surge in computational costs. While past technological advancements have kept pace by scaling the power of classical systems, further progress is increasingly hindered by the fundamental physical limitations encountered at the nano-scale [7, 8].

In this context, quantum computers offer a transformative alternative by leveraging phenomena such as superposition and entanglement. These properties enable the pursuit of *Quantum Advantage*—a state where quantum algorithms demonstrate superior speed or efficiency over classical methods. Such improvements in efficiency may be realized through the reduction of model parameters, which in turn significantly lowers the overall computational burden.

This evolution of computational paradigms is not merely a technical shift, but a testament to the profound convergence of fundamental physics and practical data science. For the authors, it represents a deliberate return to the roots of physical intuition within the complex landscape of modern finance.

The first author's journey to Quantum Machine Learning began with a PhD in theoretical physics, specializing in neutrino physics and the abstract beauty of category theory. For many years, the mathematical elegance of morphisms and quantum structures seemed far removed from the daily grind of credit risk modeling and MLOps. However, experience in the banking sector and at the Warsaw School of Economics (SGH) revealed a hidden truth: the vast, noisy streams of financial data are not just tables of numbers, but complex systems that mirror the dynamics of physical fields. By treating data points as excitations of quantum fields—a core focus of his current research—he seeks to bridge the gap between high-level theory and business utility. This approach is born from the conviction that an analyst's greatest strength lies in returning to these physical roots, where the laws of nature provide the ultimate framework for understanding market uncertainty.

Complementing this vision, Krzysztof Kuba brings a perspective that bridges the gap between the microscopic and the macroeconomic. With a background starting in nanoengineering and evolving into quantitative finance, his path reflects the same transition from the physical world to the abstract world of markets. Having passed Level I of the CFA Program, further grounded his knowledge in finance. Therefore, he ensures that these quantum-inspired theories are not just elegant abstractions, but robust, scalable tools capable of meeting the demands of real-world financial engineering.

At finQbit, we do not view QML as a “black-box” replacement for classical AI. Instead, we see it as a beautiful synthesis: an opportunity to apply the most sophisticated tools of theoretical physics to the very human problems of risk and value. This chapter is an invitation to look at financial modeling not just as an exercise in statistics, but as a realization of quantum mechanics in action.

While quantum computers are still in the stage of intensive research, they are already finding applications in optimization and machine learning—a field known as *Quantum Machine Learning*. In this chapter, we will explore the fundamental concepts of QML, from data encoding to the measurement of quantum states. Following this, we will examine a practical model for credit card fraud detection, providing a foundation for broader applications of quantum machine learning in finance. However, to fully appreciate these quantum concepts, it is essential to first understand their classical counterparts.

## 26.2 Classical Machine Learning

Technological advancement and digitization have made data a common and valuable resource [3]. It is generated and processed in both structured and unstructured formats. The structuralization of data led to the development of various models, commonly referred to as *Machine Learning*. However, the processing of unstructured data—such as text, images, or videos—contributed to the development of *Deep Learning Models*. Both of these approaches are often collectively called *Artificial Intelligence*. They were created primarily to recognize patterns, but increasingly, they are used to model and generate entirely new data.

### 26.2.1 Back to the feature: the machine learning process

To appreciate the quantum leap in financial modeling, one must first master the classical framework that precedes it. Understanding the machine learning (ML) process is not merely a procedural requirement; it is about grasping how raw information is transformed into actionable intelligence. This process can be divided into five fundamental stages:

1. Data collection
2. Data cleaning and preparation
3. Model training
4. Model validation
5. Model deployment and monitoring

The journey begins with **data collection**. In the financial realm, data is the lifeblood of the model, sourced from internal bank records, market indices, macroeconomic indicators, and even the sentiment of social media. We distinguish between *structured data*—information that fits neatly into relational tables—and *unstructured data*, such as text or images, which requires more sophisticated handling to reveal its patterns.

Next is **data cleaning**, a stage governed by the fundamental principle: “garbage in, garbage out.” If a model is fed low-quality, noisy data, its outputs will inevitably be flawed. Real-world financial data is notoriously messy, often riddled with missing values or extreme outliers that can significantly skew the learning process.

Once the data is refined, we select and **train the model**. The choice of the model depends on what we are trying to predict. If our goal is to assign observations to discrete categories—such as identifying a transaction as either “fraudulent” or “legitimate”—we are dealing with a *classification* problem. If, however, we aim to predict a continuous value, such as a stock price or an option’s fair value, we utilize *regression*.

The distinction further depends on the availability of **labels**. In *supervised learning*, the model learns from historical examples where the outcome is known (e.g., predicting customer churn based on past behavior). In contrast, *unsupervised learning* seeks to uncover hidden structures in data without predefined labels, such as clustering clients with similar risk profiles. A third paradigm, *reinforcement learning*, involves an agent that learns to make decisions by receiving feedback (rewards or penalties) from its environment.

After training, the model must undergo rigorous **testing and validation**. This stage ensures the model generalizes well to unseen data and that the underlying business logic remains sound. We must ask: is the level of inaccuracy acceptable? In the high-stakes world of banking, a poorly validated model can lead to catastrophic financial losses or severe reputational damage.

Finally, the model is **deployed** into a production environment. However, deployment is not the end of the journey. Constant monitoring is required to detect “drift” and ensure the model remains aligned with shifting market dynamics. It is crucial to view this entire lifecycle as iterative; a failure in validation often necessitates a return to data cleaning or feature engineering.

## 26.2.2 Neural networks: optimization as the new paradigm

Neural networks (NN) are a class of machine learning models inspired by the structure and function of the human brain. They consist of interconnected layers of nodes, or “neurons,” that process and transmit information by assigning weights to input data and adjusting these weights through learning algorithms.

In the modern landscape of artificial intelligence, machine learning can be fundamentally viewed as the process of defining a specific **network architecture paired with an appropriate loss function**. The challenge is no longer about finding a closed-form analytical solution to a problem—which classical statistics often sought—but rather about navigating a high-dimensional landscape to find the optimal set of parameters. This involves critical decisions: choosing the right architecture, selecting an efficient optimization algorithm, and defining a cost function that accurately captures the nuances of the business problem.

Perhaps the most widely used architecture is the feed-forward neural network, as

shown in Figure 1. To understand how this works in practice, let us consider the credit card fraud detection problem. Here, the goal is to classify transactions as either legitimate (0) or fraudulent (1). The input layer receives variables such as transaction amount, time, and location. These signals are multiplied by weights, passed through non-linear *activation functions*, and propagated through hidden layers until a prediction is reached.

The difficulty of this task is twofold. First, financial data is often heavily imbalanced—fraudulent transactions are rare, making it hard for the model to learn their patterns without biased results. Second, we face the dilemma of misclassification: a *false positive* (blocking a valid payment) causes customer frustration, while a *false negative* (missing a theft) leads to financial loss.




Figure 26.1: A simple neural network design. Each circle acts as a “neuron,” and each neuron from a given layer is connected to every other neuron in the next layer. The bigger the weight of connection, the darker the shade.

The true power of a neural network lies in its ability to act as a universal approximator, fitting virtually any arbitrary function. However, this power comes at a high computational price. Training a model involves back-propagation [9], an algorithm that calculates the “gradients” of the loss function with respect to every weight in the network. As we increase the number of layers and parameters to capture more complex patterns, the search for the global minimum of the loss function becomes exponentially difficult.

This is where the weakness of classical computing emerges: the sheer volume of floating-point operations required for such high-dimensional optimization starts to hit the ceiling of current hardware capabilities. Classical systems struggle to find an efficient path through a massive, non-convex loss landscape. This bottleneck is exactly what motivates our transition into the quantum realm. Where a classical computer struggles to find an analytical or numerical path through a massive, non-convex loss landscape, quantum systems offer a new way to explore and represent these complex spaces.

## 26.3 Quantum Machine Learning

The current phase of quantum technology development is widely referred to as the Noisy Intermediate-Scale Quantum (NISQ) era [10]. The term “Intermediate-Scale” means that contemporary quantum processors have a limited number of qubits to operate, typically a few hundred. For comparison, most daily-use laptops have 16 GB of RAM, which is roughly 128 billion bits. This constraint means that quantum computers currently remain unsuitable for general-purpose computation. Moreover, the “Noisy” term means that the outputs of quantum processors are subject to error. This error can be partially mitigated by Quantum Error Correction (QEC) algorithms. However, these schemes are not yet fully implemented. A further limitation stems from the requirement that quantum gate operations be performed on timescales much shorter than the characteristic decoherence times of the qubits, thereby restricting the achievable circuit depth and hindering the implementation of large-scale, highly complex quantum algorithms.

Despite these limitations, NISQ devices are anticipated to enable the demonstration of Quantum Advantage for computational tasks of practical relevance. In contrast, experimental realizations of Quantum Supremacy to date have been restricted to highly specialized benchmark problems—most notably random circuit sampling—that are not directly applicable to real-world domains such as quantitative finance.

A central class of methods in this regime is based on Parameterized Quantum Circuits (PQCs), which we will explore in this section. In such approaches, quantum circuits containing tunable parameters are trained using classical optimization algorithms to determine the most optimal parameter values. This paradigm, which combines classical optimization routines with quantum information processing, is commonly referred to as hybrid quantum-classical learning.

### 26.3.1 Parameterized quantum circuits (PQC)

PQCs are the cornerstone of hybrid quantum-classical algorithms. Each gate in the circuit can depend on a set of continuous parameters, which are optimized using classical methods.

Parameterized quantum circuits are realized through a predetermined arrangement of quantum *gates*, augmented by tunable, parameter-dependent gates. These gates are operations performed on a qubit, enabling highly expressive, non-linear transformations of quantum states [11, 12]. Quantum Machine Learning (QML) methodologies are being developed on the basis of such circuits and encompass a broad class of so-called Variational Quantum Algorithms (VQAs). These models can be implemented using Python-based software development kits, such as IBM Qiskit, PennyLane, or Cirq, as well as the specialized `fnqbit` library, and can be executed on both simulators and contemporary quantum processing units (QPUs).

### 26.3.2 Data as field excitations: a new encoding paradigm

Traditional QML approaches often treat data encoding—the process of mapping classical data  $x$  into a quantum state  $|\psi(x)\rangle$ —as a mere technical necessity, often using simple angle or amplitude encoding. However, as established in our introduction, we propose a more fundamental perspective: treating classical financial data as excitations of quantum fields.

In this framework, the initial state preparation is not just a transformation of coordinates, but the creation of a field configuration. If we consider the financial variables (such as credit history or market volatility) as local field values, the process of encoding them into a quantum circuit becomes equivalent to simulating the interaction of these fields with the qubits. This perspective allows us to utilize the rich mathematical structure of Quantum Field Theory (QFT) to design more robust feature maps. By viewing the input data as a set of excitations, we can better capture the correlations and “entanglement” inherent in financial markets, which are often missed by classical statistical methods. This theoretical bridge ensures that the PQC does not just process numbers, but involves a physical representation of the underlying financial reality.

### 26.3.3 Quantum neural network model

The implementation of a Quantum Neural Network (QNN) follows a process structured into three main sequential stages:

1. Data preparation (Encoding)
2. Parametrized Quantum Circuit (The Ansatz)
3. Measurement

The **data preparation** step serves a purpose analogous to embeddings in classical neural networks and requires prior steps like data exploration, cleaning, and standardization. In the context of quantum computers, we must translate classical data into quantum states—a process formally known as *data encoding*. This is precisely where our conceptualization of data as *field excitations* comes into play. Rather than merely mapping a numerical value to an angle, we treat data encoding as the preparation of a physical field state, embedding classical financial variables into the exponentially large Hilbert space of the quantum system. The schematic representation of this step is shown in Figure 2. Moreover, if the dataset is highly dimensional, additional dimensionality reduction techniques may be required before encoding.

Next, we introduce the quantum equivalent of a “hidden layer”: the parametrized quantum circuit. In quantum jargon, this layer is often called an *ansatz* (a term derived from German, meaning an “educated guess” or “assumption”). Because we rarely know the perfect circuit structure in advance, we postulate a sequence of gates and iteratively optimize their parameters. Choosing this structure is equivalent to defining the architecture (layers and neurons) in classical deep learning.

To get a sense of how the ansatz operates, observe Figure 2. There are two main components: the single-qubit rotation gates (e.g., U3 gates, represented as boxes) and

the two-qubit entangling gates (such as CNOT gates, which connect the horizontal qubit lines). You can think of the U3 gates as the nodes (neurons) in a neural network. These gates contain tunable *parameters* that are updated during the learning process, allowing the model to capture the underlying structure of the data. Furthermore, the CNOT gates are responsible for generating quantum *entanglement*, functioning similarly to the weighted connections that link neurons across different layers.

Designing an effective ansatz involves balancing expressibility (modeling capacity) against trainability. While more parameters and complex entanglement allow the model to capture highly non-linear data patterns, overly deep circuits risk encountering a phenomenon known as *Barren Plateaus*. In this scenario, the gradients vanish exponentially, making the training process entirely ineffective.

Last but not least is the **measurement** stage, which is analogous to the classical output layer. Measuring a qubit causes its quantum state to collapse into a classical bit (either 0 or 1). Returning to the credit card fraud problem discussed earlier: similar to a classical NN, if the result of the measurement is 0, the transaction is classified as normal. Conversely, a measurement of 1 alerts us to a potential fraud.

Formally, there is one additional component required to close the learning loop: the *cost function* (or loss function). The purpose of this function is to evaluate how much our model’s prediction differs from the true label. The output of this cost function is then passed to a classical optimizer—often utilizing techniques like the *parameter-shift rule* (the quantum analogue of back-propagation)—which fine-tunes the U3 gate parameters to fit the data.




Figure 26.2: Representation of the Quantum Neural Network, from the Input Layer (Data Encoding), via the hidden layer (the ansatz), to the output layer (measurement of the qubit).

Having discussed the architecture of quantum neural networks, let us dive deeper into the intuition behind data encoding and the ansatz. Currently, our classical credit card data lives in a 2-dimensional table on a flat piece of paper. Imagine coloring the points representing normal transactions green, and the fraudulent ones red. Because this “paper” contains hundreds of thousands of densely packed dots, drawing a single straight line that perfectly separates the green dots from the red ones is impossible.

A classical linear classifier can try its best, but it will inevitably fail.

Now, instead of points confined to a 2-dimensional space, imagine these data points as marbles. Throw all these marbles onto a trampoline, bouncing them up into the air. Suddenly, the data is no longer flat! By popping the data into a new, higher dimension (the space above the trampoline), hidden patterns emerge. This additional dimension provides the spatial freedom to simply slide a flat piece of cardboard (a hyperplane) right between the red and green marbles.

To transfer this analogy to quantum computing: *data encoding* is the process of turning flat dots into marbles, the *ansatz* acts as the trampoline launching them into a high-dimensional quantum Hilbert space, and the *measurement* is the flat cardboard plane that cleanly distinguishes between normal and fraudulent behavior. This provides a vivid intuition of what happens when we map data into higher dimensions, illustrating the core power of quantum machine learning.

### 26.3.4 Is it all any useful?

So far, we have seen how quantum neural networks are different from classical ones. Now the question arises: can we do anything useful with it? The answer is yes, of course! There are several applications of quantum machine learning for classification problems, including, predicting customer behavior, grouping clients and fraud detection.

To begin with, the variety of financial products and their customization requires the ability to predict customer behavior. Banks desperately want to anticipate when a client will buy their first home, when they might abandon their account for a competitor, or which specific investment product they actually need right now. Unfortunately, human decisions are influenced by many different factors, and sometimes are even random and unpredictable. The decrease in monthly spending may happen just by accident, or it may signal upcoming financial problems of an individual. The aim of machine learning models is to distinguish between these two scenarios. Bombarding a customer with loan offers of financial products at the wrong time can discourage the client from the bank. On the other hand, missed opportunity can cause the customer to switch to a competitor.

With the wide verity of highly customizable products, banks may be interested in assigning each customer to a specific group, called *cluster*. By doing so, the bank can advertise product to group of clients with similar habits, needs and risk profiles. In this way, we can offer more personalized services rather than irrelevant, generic products that are almost instantly rejected by the customer. To visualize this problem, imagine trying to organize a massive library. A classical computer might simply group the books by broad categories like genre or author name. In finance, this refers to lumping customers together by basic demographics such as age or zip code, which completely misses the nuanced behavior of how people actually manage their money. As banks collect thousands of distinct data points on every individual, from daily coffee purchases to long-term investment strategies, the sheer volume of variables creates a multi-dimensional puzzle, just like with the green and red marbles. Classical algorithms struggle to calculate the mathematical “distance” or similarity

between customers when there are too many competing factors. In practice, this may result in inaccurate grouping, where a risk averse retiree might be mistakenly offered the same high-risk investment product as an aggressive day trader.

And last but not least, the credit card fraud problem. We have already introduced this problem in the context of neural networks. Let us dive deeper into this problem and investigate the results of a classical neural network model and a quantum version. To show that quantum models can possibly have better capabilities than the classical neural network, we trained both models on the credit card fraud dataset [13]. First, we trained a simple neural network with 30 neurons in the input layer, each neuron corresponding to a different variable. There were also 2 hidden layers, with 8 and 4 neurons each. Finally, the output layer consisted of only one neuron, with output between 0 and 1. This gives a total number of parameters equal to 276. The results of the validation step of the model are shown in Table 1.

The quantum model was essentially the same as the one presented in Figure 2, with 4 qubits, with 4 encoded variables. However, there were 4 hidden layers, instead of only 1 layer, in order to improve model capabilities to recognize patterns. With each  $U_3$  gate containing three parameters, and one gate applied to each of the four qubits across four layers, the architecture possesses a total of  $3 \times 4 \times 4 = 48$  trainable parameters. As for large language models, we are told that the more parameters, the better. We will see if this is actually the case for quantum machine learning.

One final remark on the training stage of both models - the trick we used to overcome the issue of imbalanced datasets was to train both models on a smaller dataset, which was balanced. By picking the same number of fraudulent and non-fraudulent transactions, the models had better ability to capture patterns and nuances. It may sound counter-intuitive, as we are told that the more data, the better. Well, not always, and it is the role of machine learning engineer to come up with different solutions to not so obvious problem and see what works best. However, this technique is not methodologically correct when we perform model validation. Each model should be evaluated on separate data that the model had not seen before, which ensures model ability to generalize. If the model performs well on training data, but fails on unseen examples, we might stumble on the problem of *overfitting*. Overfitting should be avoided because it means that the model learned data we provided for training in great detail, but the results are not applicable to unseen, real-world data. Imagine a elementary school student who learns how to add. One student memorized all answers from the exercises:  $2 + 3 = 5$ ,  $3 + 4 = 7$ , and so on. The other student actually learned how to add numbers. The question on exam is  $10 + 11$ . Who is going to pass and who is going to fail? The same logic can be applied to overfitting. Having described the architectures of the models, let us put some numbers into perspective.

Table 26.1: Confusion Matrix: Classical Neural Network (NN)

Prediction	Normal (0)	Fraud (1)
True: Normal (0)	54,820	2,044
True: Fraud (1)	9	89

The tables above show the number of transactions that the model correctly and

Table 26.2: Confusion Matrix: Quantum Machine Learning (QML)

Prediction	Normal (0)	Fraud (1)
True: Normal (0)	56,762	102
True: Fraud (1)	12	86

incorrectly classified as normal or fraudulent. The truly correct examples are shown in the rows, while predictions of the models are shown in columns. These tables also directly shows the imbalance of the data set. There were almost 55,000 transactions that the model classified as normal, which were in fact, correct. This is called true negative. Here, the quantum model correctly predicts more true negatives than the neural network, which is the first sign of improved performance. Moreover, the number of false negatives (transaction that were classified as normal, where in reality they were frauds) is similar - 12 in case of quantum model and 9 in case of neural network. However, the false positives is the place where the quantum model reveals its potential. Look at the numbers: 102 versus 2,044! These numbers correspond to the case, where the transaction would have been blocked, even though it was completely normal. In reality, these would be almost 2,000 clients complaining about their card being blocked, and potentially leaving their current bank. Remember that these results were achieved with fewer variables (4 versus 30) and fewer parameters (48 versus 276).

### 26.3.5 Reality check

The empirical success of the quantum model in our credit card fraud example clearly demonstrates its potential to outperform classical architectures. However, does this mean quantum models will *always* reign supreme? To critically assess this claim, we must weigh the inherent advantages and limitations of both paradigms.

One of the primary advantages of classical neural networks is their profound ability to learn and adapt. Their architecture allows them to generalize complex, non-linear dependencies across massive datasets. Coupled with the highly efficient back-propagation algorithm, they have become the workhorses of modern artificial intelligence, powering everything from computer vision to Large Language Models (LLMs).

Furthermore, classical networks are highly scalable. While a relatively simple task like recognizing handwritten digits might require only a few dozen neurons, LLMs seamlessly scale to billions of parameters and hundreds of layers. Yet, this scalability is a double-edged sword. As financial data grows in complexity, the computational power required to train these massive models approaches absolute physical and thermal processing limits, creating a severe bottleneck for future development.

Additionally, classical networks notoriously struggle with imbalanced datasets. They often fail to correctly identify the underrepresented class, leading to high rates of false positives and false negatives. In our baseline test, the classical neural network produced an unacceptable rate of false positives (2,044 normal transactions incorrectly blocked), risking severe customer dissatisfaction. While a naive solution

might be to simply add more neurons, this does not guarantee better performance and significantly increases the risk of *overfitting*—a scenario where the model memorizes the training data but fails to generalize to the real world.

A critical limitation shared by both classical and quantum paradigms is the “black-box” problem of interpretability. In credit scoring, for instance, financial regulations often require banks to explicitly explain why a loan was declined in order to prevent discrimination and ensure equal access to capital. Consequently, a model’s parameters must provide interpretable insights into which variables drove the decision. Because neither deep classical NNs nor QNNs natively offer this transparency, their deployment is often restricted to internal tasks (like fraud detection) rather than customer-facing credit decisions. To address this, researchers are actively developing frameworks broadly referred to as XAI, where the “X” stands for *Explainable* Artificial Intelligence.

Quantum neural networks share several of these broad advantages and limitations, but they introduce a fundamental shift in how information is processed. Their true power lies in their dimensionality. By mapping financial data into a high-dimensional quantum feature space—much like bouncing our scattered marbles off a flat surface into the air—QNNs can identify subtle, intricate relationships that classical models miss. Crucially, they achieve this with a fraction of the parameters. In our study, the quantum model achieved superior predictive results using a highly efficient architecture of just 4 qubits, 4 variables, and 48 parameters, slashing the number of false positives to a mere 102.

Despite these compelling theoretical and empirical advantages, practical QML is currently constrained by the Noisy Intermediate-Scale Quantum (NISQ) era. Modern quantum processors possess a limited number of qubits and suffer from high error rates due to environmental noise and decoherence. While researchers routinely use classical computers to flawlessly simulate quantum circuits up to around 30 qubits—which conveniently allows for the use of classical back-propagation to drastically speed up training—this *hybrid* simulation approach hits a hard physical memory limit beyond that scale.

When transitioning to actual quantum hardware (QPUs), we gain access to more qubits, but we lose the ability to use back-propagation. Because inspecting intermediate quantum states to calculate gradients would cause the wavefunction to collapse, hardware training must rely on gradient estimation techniques (like the parameter-shift rule) which require multiple circuit evaluations per parameter. Therefore, while QPUs offer an exponentially large state space, training time remains a significant bottleneck for large-scale quantum models today.

## 26.4 Summary

In this chapter, we have traversed the evolving landscape of machine learning, from foundational classical approaches to the frontier of quantum models. We began by identifying credit risk, credit scoring, and fraud detection as paramount challenges for modern banking. While classical models—such as feed-forward Neural Networks (NN)—have historically managed these complex datasets through rigorous machine

learning pipelines, they are increasingly constrained by the physical, thermal, and computational limits of classical hardware.

To overcome these limitations, we explored the mechanics of hybrid quantum-classical learning. Just as classical networks rely on input nodes, hidden layers with tunable weights, and non-linear activation functions, Quantum Neural Networks (QNNs) operate through analogous stages: Data Encoding (translating classical data into high-dimensional quantum states, akin to physical field excitations), The Ansatz (a parametrized quantum circuit utilizing tunable rotation gates and entanglement), and Measurement.

Our comparative analysis of a credit card fraud detection model provided a tangible demonstration of this potential. By mapping financial data into a quantum feature space, the QNN achieved superior precision, drastically reducing false positives (from 2,044 down to 102) while utilizing a mere fraction of the parameters required by its classical counterpart.

In conclusion, this chapter highlights the very real promise of *Quantum Advantage* in finance. It is not merely a theoretical concept, but a practical evolution: the ability to harness the fundamental laws of physics and high-dimensional geometry to solve complex, real-world classification problems more efficiently, accurately, and elegantly.

## References

- [1] R. Orús, S. Mugel, and E. Lizaso. “Quantum computing for finance: Overview and prospects.” *Reviews in Physics* 4 (2019): 100028.
- [2] D. Herman et al. “Quantum computing for finance.” *Nature Reviews Physics* (2023): 1–25.
- [3] Frontiers Editorial. “Quantum computing: foundations, algorithms, and emerging applications.” *Frontiers in Quantum Science and Technology* (2025). Available at: <https://www.frontiersin.org/journals/quantum-science-and-technology/articles/10.3389/frqst.2025.1723319/full>.
- [4] D. Kaszynski, B. Kaminski, and T. Szapiro. “Credit Scoring in Context of Interpretable Machine Learning.” Warsaw: SGH, 2020.
- [5] Basel Committee on Banking Supervision. *International Convergence of Capital Measurement and Capital Standards: A Revised Framework—Comprehensive Version*. Bank for International Settlements, June 2006. Available at: <https://www.bis.org/publ/bcbs128.htm>.
- [6] Basel Committee on Banking Supervision. *Basel III: A Global Regulatory Framework for More Resilient Banks and Banking Systems*. Revised version (originally published December 2010). Bank for International Settlements, June 2011. Available at: <https://www.bis.org/publ/bcbs189.htm>.
- [7] M. M. Waldrop. “The chips are down for Moore’s law.” *Nature* 530.7589 (2016): 144.

- [8] G. E. Moore. “Cramming more components onto integrated circuits.” *Electronics* 38.8 (1965): 114–117.
- [9] D. E. Rumelhart, G. E. Hinton, and R. J. Williams. “Learning representations by back-propagating errors.” *Nature* 323.6088 (1986): 533–536. DOI: <https://doi.org/10.1038/323533a0>.
- [10] J. Preskill. “Quantum Computing in the NISQ era and beyond.” *Quantum* 2 (2018): 79. DOI: <https://doi.org/10.22331/q-2018-08-06-79>.
- [11] A. P. Lund, M. J. Bremner, and T. C. Ralph. “Quantum sampling problems, BosonSampling and quantum supremacy.” *npj Quantum Information* 3.1 (2017): 15. DOI: <https://doi.org/10.1038/s41534-017-0018-2>.
- [12] A. Harrow and A. Montanaro. “Quantum Computational Supremacy.” *Nature* 549 (2017): 203–209. DOI: <https://doi.org/10.1038/nature23458>.
- [13] “Credit Card Fraud Detection Dataset.” Kaggle. Available at: <https://www.kaggle.com/datasets/mlg-ulb/creditcardfraud>.

## Chapter 27

# Musings Between Breakthroughs & Builds

### *POVs & FAQs for a future that can't build itself*

Sierra Clouse

At a recent MIT quantum hackathon, I was one of few participants whose everyday work & learning didn't include writing code.

Most participants were impressively well versed in pulse sequences, calibration limits, & error mitigation. I usually live in the world of “decks & docs”—commercialization models, integration strategies, & long-term roadmaps. I joined the hackathon as a deliberate, tactical move to stay close to the build & the builders. Moving from slide decks into Discord threads & GitHub repos is a learning exercise, to be sure. But it's how we ensure our venture & services work remains grounded in high-fidelity execution.

In a live repo, theory becomes tactile. There is no room for abstraction: it either runs, or it doesn't.

What I saw was a room & online spaces full of people trying to bridge a gap in expectations. Researchers want their work to outlive a whitepaper. Developers want tools that don't break in production. Far from hobbyists, many of them are on the frontline of a new industry trying to figure out how breakthroughs survive the “real world” contact: procurement cycles, regulatory scrutiny, & enterprise scale integration/interoperability.

**Commercialization is a contact sport.** You don't get the customer insights or the execution needed to monetize a breakthrough top down from a boardroom. You get them bottoms up & middle out with the people building & supporting the stack.

---

This chapter is a featured contribution to the monograph *A Portrait of Quantum Technologies in Finance*, edited by Oswaldo Zapata, PhD, and published by The Quantum Finance Boardroom (2026). All rights reserved by the chapter's author.

## 27.1 The Pattern of Progress

I didn't come to quantum technology as a physicist or through academia. I came as someone who has spent years helping emerging technologies move from possibility to implementation.

Whether it's AI, distributed systems, or the cloud, every innovation wave follows a similar S-curve. It starts with the steep climb of technical discovery, but it only hits the "inflection point" of mass adoption when the focus shifts from what it is to how it integrates.

The inflection point in this cycle is rarely just technical; it is organizational. Quantum tech isn't a single "eureka" moment. It is a layered ecosystem forming in parallel across academia, startups, enterprises, national labs, & defense agencies.

At our studio, we see commercialization as a specific path of conversion:

- **Research** provides the scientific foundation; there is always more to learn here.
- **Roadmap** is the engineering that turns that foundation into something specific & useful.
- **Revenue** is the market's confirmation that the engineering solved a problem (or unlocked an opportunity) worth paying for.

When these three move together, industries mature.

& you are likely reading this because you're interested in these implications for the Finance industry where this conversion is particularly high stakes & vital. We are dealing with the structural plumbing of our global economy, after all. It's an area where we have to balance two equally important goals: building a strong **Defense** to keep our systems safe, & developing a creative **Offense** to make those systems smarter & more resilient.

## 27.2 The Quantum Divide: Building for the Real World

The global finance industry, like a few others, is navigating a "Quantum Divide"—a gap between our long-term aspirations & the immediate realities of our systems. Bridging this gap means looking at the plumbing of the economy through a very practical lens.

### **Defense: looking out for the "breakers"**

Right now, there is a bit of a disconnect in how we think about security. Many of us are focused on the long game—waiting for perfectly stable hardware. But the

“breakers”—those looking to find gaps in our systems—don’t follow a roadmap. They aren’t waiting for a specific milestone before they start experimenting.

They use what is functional right now: quantum annealing, hybrid configurations, & cloud-access hacks.

- **The Middleware Gap:** Often, the vulnerability isn’t the complex math; it’s the simple implementation. Think of the *2016 Bangladesh Bank heist*. The “exploit” wasn’t a mathematical breakthrough; <sup>1</sup>it was a legacy system relying on hardcoded tokens or static seeds. A quantum annealer today can search for those patterns significantly faster than classical systems.
- **Adversarial AI:** Breakers are already using AI to map out vulnerabilities in seconds. By adding quantum’s ability to stress-test trust systems, the challenge becomes multi-dimensional (Schwartz et al.). Our best defense is to evolve alongside them, using AI-led monitoring to spot these high-speed patterns before they become a problem.

## Offense: finding the opportunity in resilience

In a high-stakes market, Offense is the ability to remain decisive & liquid when others are blinded by volatility. It is the pursuit of systemic resilience as a proactive capability. This creates a unique window of opportunity to capture value while the rest of the market is stuck in a reactive loop.

- **The Window to Act** Resilience converts into a score through timing & readiness. This depends on having the organizational protocols required to act within the small window of time between the initial detection of a pattern & the total market realization of a shock. The offensive advantage belongs to those who have the infrastructure to trust their signals & move—whether rebalancing a portfolio or securing liquidity—while competitors are still diagnosing the problem.
- **The Information Advantage** Classical tools excel at identifying historical patterns, but they can struggle with coordinated activity that spans multiple institutions. We see this challenge addressed in central bank experiments like Project Hertha (Bank for International Settlements) & Project Leap (Bank for International Settlements and Bank of England), collaborations between the Bank of England & the BIS Innovation Hub. These initiatives show that during periods of high-dimensional complexity, traditional systems can produce results that obscure emerging risks. Success in an offensive context is the technical ability to see through that complexity before it settles.
- **The Quantum-Native Edge** This is the toolset for gaining a predictive lead. Quantum-enhanced machine learning (QML) is uniquely capable of analyzing the multi-dimensional transaction networks of global finance (Patel #). By identifying subtle, non-linear correlations that indicate a coordinated market shift or a systemic shock, institutions can gain a functional lead time.

---

<sup>1</sup>Often a COBOL-based system.

Organizations including JPMorgan Chase (JPMorgan Chase) & Goldman Sachs (Kundu et al.) have already begun exploring these quantum-classical hybrids to optimize high-stakes areas like credit risk & complex portfolio rebalancing.

- **Trust & Clarity** In an applied environment, a playbook is only as effective as its governance. For an offensive strategy to be actionable, it must survive the scrutiny of risk committees & regulators. This requires a move toward explainable, quantum-native models. Winning in finance involves having the analytical depth to see the invisible paired with a transparent framework that gives human leaders the confidence to act when the window of opportunity opens.

## 27.3 The Question(s) in Front of Us

The transition from a weekend sprint to an industrial revolution takes a different kind of endurance. While the talent in the room is working through Total Readiness Levels (TRLs), the breakers are already operating & tinkering with threat vectors defined by current functionality. As we bridge that gap, we have to decide:

- Will we treat quantum as a speculative cycle, or as infrastructure to be built deliberately?
- Will we invest in literacy before urgency forces our hand?
- Will we remember what prior revolutions taught us about scale & stewardship?

## 27.4 Make These Your FAQs

Building & cultivating a perspective & capabilities with quantum tech doesn't have to feel like a chore; it's actually an engaging way to find your own path through the noise & string in the maze. Because the field is so “lumpy”—with access, awareness, & readiness varying wildly between different countries, companies, & even individual roles—staying curious is your best tool for staying oriented.

Think of these questions as a supportive starting point to spark better conversations, lean into a bit of “beginner’s mind,” & discover new ways to connect with the broader ecosystem of people & partners growing alongside you:

1. **If the topic is “Workforce Development”:** Ask “Which *three specific job titles* currently have the most critical proficiency gaps? (e.g., Quantum Algorithm Developers for non-convex optimization, Quantum-Classical Systems Architects, or PQC Migration Leads).” Use this question to gauge awareness of value chain needs & required proficiencies. This helps move the conversation from abstract talent needs to a methodical hiring & training roadmap for upskilling/cross skilling.
2. **If the topic is “Quantum Security”:** Ask “How are we looking at *functional vulnerabilities* today? If a breaker uses hybrid-classical-quantum tools available

now, where is the entry point?” (Hint: Look at *legacy banking mainframes* or *SWIFT-connected middleware* where the vulnerability is often a legacy reliance on static seeds or hardcoded tokens).

3. **If the topic is “Product Readiness” (aka “it’s years away”):** Ask “What is your product-specific critique? What is or isn’t ready, valuable or accessible with tools like NVIDIA’s Brev or NVQlink, IBM’s Qiskit, or AWS Braket, etc? What specifically about the current noise profile, circuit depth, or general functionality is the bottleneck for your roadmap or how you would like to use these tools?”

## 27.5 Bonus Round: Orient, Learn, and Act

There is a special kind of “click” that only happens when you stop reading about the tech & start interacting with it. Use these tasks to find your footing, test your assumptions, & get a better feel for the actual road ahead from the people who are paving it:

### Task 1: Map the Missing Link

- **The Prompt:** Who is the one person, team, or skill I need to engage with right now to bridge a specific technical, commercial, or knowledge/data gap?
- **The Action:** Don’t just name a role; find a living example of that expertise in the wild. Identify a specific project, an active repo, or a partner organization where that skill set is currently solving a problem. Reach out, follow their work, or study their documentation to see how they’ve bridged the gap you’re facing.

### Task 2: Trace Your Digital Footprint

- **The Prompt:** As a digital citizen, how “quantum-safe” is the journey of the data you interact with every day?
- **The Action:** Use your own daily digital activities as a learning lab for readiness. Pick one service you rely on—a SaaS provider, a bank, or a communication tool—& look for their stance on post-quantum cryptography (PQC) & data privacy. Are they moving toward sovereign-stack protections or dual-use compliance? By asking, “How are the systems I rely on protecting my data against future decryption, forgery, & fraud?” you can move from a passive user to an informed participant who understands the baseline of the current ecosystem & adjust accordingly.

### Task 3: Test Drive Available Tools

- **The Prompt:** Beyond the papers & pitch decks, what do the current products, tools, & available platforms actually deliver for my purposes?
- **The Action:** Spend an hour “test-driving” the claims of a specific tool or provider rather than just reading about them. Sign up for a cloud platform

like *IBM Quantum* or *AWS Braket* & run a sample circuit using a compiler like *Qrisp* or *Qiskit*. Make *Github* your curiosity accomplice. Pay attention to what is missing or frustrating—like queue times, noise levels, or gaps in documentation. This hands-on approach allows you to decide what you actually find useful & which products are ready for your roadmap, so you never have to just take a vendor’s word for it.

Every technological age carries a weight of responsibility alongside the thrill of discovery. Discipline & wonder both matter. Quantum technology expands the horizon of what is achievable. & how we choose to address the current challenges faced by builders, exploited by breakers, & borne by end users will define the era.

## Works Cited

Bank for International Settlements. “Project Hertha: Identifying Financial Crime Patterns in Real-Time Retail Payment Systems.” *BIS Innovation Hub Reports*, vol. 2025, no. 96, 2025, p. 1. *Bank for International Settlements*, <https://www.bis.org/publ/othp96.pdf>.

Bank for International Settlements and Bank of England. “Project Leap Phase 2: Quantum-Proofing Payment Systems.” *BIS Innovation Hub Reports*, vol. 2025, no. 107, 2025, p. 1. *Bank for International Settlements*, <https://www.bis.org/publ/othp107.pdf>.

JPMorgan Chase. “Global Technology Applied Research: Quantum Computing for Finance.” *J.P. Morgan Technology Insights*, vol. 2025, no. 1, 2025, p. 1. *JPMorgan Chase & Co.*, <https://www.jpmorganchase.com/about/technology/research/applied-research>.

Kundu, Niranjana, et al. “A Detailed, End-to-End Assessment of a Quantum Algorithm for Portfolio Optimization Released by Goldman Sachs and AWS.” *AWS Quantum Computing Blog*, 2023, <https://aws.amazon.com/blogs/quantum-computing/a-detailed-end-to-end-assessment-of-a-quantum-algorithm-for-portfolio-optimization-released-by-goldman-sachs-and-aws/>. Accessed 19 February 2025.

Patel, B. “Quantum Machine Learning for Financial Forecasting and Portfolio Optimization: Algorithms, Applications, and Future Prospects.” *International Journal of Science and Applied Technology*, vol. 10, no. 4, 2025, p. 1. *International Journal on Science and Technology (IJSAT)*, <https://www.ijstat.org/papers/2025/4/9033.pdf>.

S., Kim, et al. “Quantum annealing for combinatorial optimization: a benchmarking study.” *npj Quantum Information*, vol. 11, no. 77, 2025, p. 77, <https://doi.org/10.1038/s41534-025-01020-1>.

Schwartz, Reva, et al. “Adversarial Machine Learning: A Taxonomy and Terminology of Attacks and Mitigations.” *NIST AI 100-2e2025*, vol. 2025, no. 2025, 2025, p. 1. *NIST Information Technology Laboratory*, <https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-2e2025.pdf>.

White, Geoff. *The Lazarus Heist: From Hollywood to High Finance: Inside North Korea’s Global Cyber War*. Penguin Business, 2022.

# Chapter 28

## Quantum and Security

*Protecting vital functions when physics changes the cost of attack*

Tatiana Mitrova

### 28.1 Introduction: When “Quantum” Stopped Being a Future Topic for Me

For a long time I treated quantum the way many security experts treat deep scientific advances: important, but not urgent. My early mental model was linear — quantum computing matures, cryptography breaks, we swap algorithms, and the rest of security carries on. Two realizations changed that.

First, security rarely fails because of a single “breakthrough day.” It rather fails — and is rebuilt — through transitions: mixed systems, uneven adoption, incompatible assumptions, and the friction of real infrastructure. Even if a disruptive quantum computer is years away, incentives move now: data is collected now, architectures are chosen now, and standards and procurement decisions are also made now.

Second, I realized that “quantum” is not one capability. It is a family — computation, communication, sensing and timing — each landing on different security dependencies. Some are still emerging; others have shaped real operations quietly for decades.

This chapter is a basic, security-first explanation of (1) what security is at its most fundamental level and (2) how quantum technologies both challenge and protect it. The goal is to replace vague excitement or vague fear with a practical way of

---

This chapter is a featured contribution to the monograph *A Portrait of Quantum Technologies in Finance*, edited by Oswaldo Zapata, PhD, and published by The Quantum Finance Boardroom (2026). All rights reserved by the chapter’s author.

thinking: the pathway from stress, through vulnerability, to consequence, and the defensive levers that change outcomes.

So what is security, in the most general sense? Across military security, food security, energy security, health security, cyber security, and economic security, the same definition works:

*Security is the ability of a system to keep delivering vital functions under stress.*

This framing is deliberately “systems-first.” It avoids defining security as the absence of threats — an impossible standard — and instead defines it as performance under pressure.

Three questions sit inside the definition:

- What must continue? (the vital function or mission).
- Under what stresses? (adversarial and non-adversarial shocks).
- With what tolerated degradation? (acceptable loss, downtime, and recovery time).

A note that matters in practice: “security” often emphasizes intentional adversaries while “safety” emphasizes accidents. Real systems rarely allow that separation — accidents create openings; adversaries exploit chaos; crises accelerate unsafe decisions.

To keep analysis clear, treat risk as a pathway:

$$\text{Risk} = \text{Stress} \times \text{Vulnerability} \times \text{Consequence}$$

- Stress is the pressure applied (attack, coercion, accident, disaster, market shock).
- Vulnerability is how that pressure gets leverage (single points of failure, weak controls, dependency concentration, governance gaps).
- Consequence is what breaks when leverage succeeds (loss of life, loss of capability, cascading outages, loss of trust).

Security strategy then reduces risk through three universal levers:

1. Reduce stress likelihood (deterrence, prevention, disruption, de-escalation).
2. Reduce vulnerability (hardening, segmentation, diversification, agility).
3. Reduce consequence (redundancy, reserves, graceful degradation, rapid recovery).

Most security debates across all domains become clearer if you track five variables:

1. Secrecy horizon: how long information must remain confidential to matter.
2. Authenticity assurance: how reliably identity, authorization, and “what is real” can be verified.

3. Operational resilience: how well the system functions when dependencies fail or are attacked.
4. Detection balance: whether concealment or sensing has the advantage in context.
5. Governance leverage: who can exclude whom via standards, procurement, supply chains, and jurisdiction.

Quantum matters because it moves all these variables — sometimes in opposite directions at once.

I used to ask, “When will quantum arrive?” The better security question is, “Where are our assumptions brittle today — and what happens to them during a messy transition?” That shift — from timelines to dependencies — changed how I rank risks and what I treat as “strategic.”

A more personal note: in my day-to-day work, “future risk” becomes real when it hits ordinary constraints — budget cycles, vendor roadmaps, and systems that can’t be patched quickly. I’ve sat in meetings where “we use encryption” sounded like the end of the conversation — until someone asked the uncomfortable follow-ups: which encryption, where exactly, who owns the keys, and what breaks if we change it? Quantum didn’t create those questions. It just makes it harder to keep postponing them.

## 28.2 Where Quantum Enters the Security System

Quantum enters security through three capability families:

- Quantum computation: changes the cost of solving certain problems.
- Quantum communication: changes how keys and secure links can be engineered.
- Quantum sensing and timing: changes what can be measured, detected, navigated, and synchronized.

One point that comes through strongly in the cyber-defense community is that quantum is rarely “one big moment.” It is a slow pressure that forces upgrades in standards, tools, and habits — without breaking the systems that depend on them.

The rest of the chapter examines each domain through the same lens: stress changes, vulnerability changes, consequence changes, defensive leverage.

## 28.3 Quantum Computation: The Trust-Layer Shock

Quantum computation matters to security because it changes the cost of attack against foundational trust mechanisms. The right analysis is not “quantum breaks en-

Quantum capability	Primary security pressure	Primary defensive opportunity	Typical system risk in practice
Computation	Secrecy horizon and authenticity assurance	Post-quantum migration, crypto-agility	Hybrid seams, misconfigurations, shared-vendor fragility
Communication	Secure coordination on selected links	High-assurance key establishment (selective)	Infrastructure and endpoint dependencies; availability risks
Sensing/timing	Detection balance and operational resilience	Resilient timing/PNT and monitoring	New critical dependencies; calibration/trust-chain targets

ryption,” but rather: quantum reshapes secrecy horizons and authenticity assurance, and forces a transition that creates new seams.

The most security-relevant stresses are:

- Retrospective decryption incentives (“collect now, decrypt later”): if an actor expects future capability, collecting encrypted traffic and stored archives becomes strategically rational today. The attack is time-shifted — collection now, exploitation later.
- Signature and identity compromise (in a fault-tolerant future): if widely used signature schemes become forgeable, attackers can impersonate trusted entities, authorize actions, and distribute malicious updates that look legitimate.

This is why authenticity deserves equal billing with confidentiality. A world where you cannot reliably verify identity and authorization is not merely “less private”; it is less governable.

Even without a large quantum computer, quantum computation increases vulnerability through the transition it forces:

- Mixed architectures (legacy and post-quantum) produce interoperability edges and “downgrade” behaviors.
- Crypto inventory gaps surface late — especially where cryptography is embedded deep in devices, protocols, and vendor stacks.
- Dependency concentration grows: a small number of libraries, hardware modules, and vendors become shared failure domains during migration.
- Human factors worsen: rushed change increases misconfiguration, brittle roll-outs, and false confidence.

This part is very familiar in practice: the hard work is rarely choosing the “new” algorithm. It is finding where the “old” one is hiding — in vendor products, legacy appliances, and one-off integrations that quietly became permanent — and then getting clear ownership to change it without taking production down. In security

terms, this is a familiar pattern: the system becomes more complex faster than assurance can keep up.

If confidentiality fails, you lose secrets. If authenticity fails, you can lose control. Consequences can become systemic because identity and signatures sit upstream of many vital functions:

- secure updates and patching,
- device and user authentication,
- authorization of commands in critical systems,
- legal and economic non-repudiation,
- trust in institutional communication.

This is where “global security” becomes tangible: these trust infrastructures are cross-cutting and correlated, so failures do not stay local.

Quantum also strengthens security. The defensive story is not “wait for quantum.” It is:

- Crypto-agility: the ability to swap primitives without redesigning systems.
- Phased post-quantum migration: staged rollout with testing, fallback modes, and explicit ownership of residual risk.
- Shorten secrecy horizons where possible: minimize retention, reduce data replication, and treat long-lived archives as a special class of risk.
- Treat integrity as first-class: hardened code-signing pipelines, auditable update mechanisms, and strong roots of trust.

## 28.4 Quantum Communication: Strengthening Links Without Pretending Systems Are Perfect

Quantum communication is often introduced with a promise of “unbreakable encryption.” A security framing keeps it honest: QKD (quantum key distribution) and related techniques can strengthen key establishment for some links, but they also introduce infrastructure dependencies and do not eliminate classic failure modes.

Where quantum communication can help is narrow but meaningful:

- On certain high-value channels, QKD can make passive interception harder to perform without detection, changing the attacker’s cost-benefit calculation.
- In a world worried about future decryption, some actors may prefer link-level key establishment that does not depend on computational hardness.

But stress can also shift in the opposite direction:

- Denial remains powerful: quantum channels can be easier to disrupt than to intercept; jamming, degradation, and physical disruption remain practical pressure tools.
- Attackers move to endpoints: insiders, compromised devices, and supply chain manipulations stay dominant because they bypass “perfect channel” assumptions.

This matches what appears again and again in security reviews: when “the channel” gets strong, attackers do not give up — they go after the people and devices at the ends, or the supplier path in the middle.

QKD does not eliminate trust; it relocates it. Typical vulnerability changes include:

- Trusted node assumptions in long-distance terrestrial networks (physical and organizational security becomes central).
- Endpoint dependence: if the device is compromised, “perfect” key exchange does not save you.
- Operational complexity: calibration, maintenance, monitoring, and incident response become part of security posture.
- Jurisdiction and governance: who owns and controls nodes, fibers, satellites, and ground stations becomes a security variable.

Used well, quantum communication can reduce consequence for a specific function: secure coordination on a critical link. Used as a geopolitical instrument, it can also increase systemic consequence by encouraging “trusted corridors” that fragment interoperability, creating strategic dependencies on infrastructure providers, or shifting alliance architectures around who can securely communicate under stress.

A defensible posture looks like this:

- Treat QKD as a targeted control for high-assurance links, not a universal upgrade.
- Pair it with post-quantum cryptography for scalable protection across broad networks.
- Invest in endpoint security, physical security, and operations — because real-world systems fail there.
- Design for continuity: secure communication is not only secrecy; it is also availability and graceful degradation.

## 28.5 Quantum Sensing and Timing: Security Begins with Time and Visibility

Quantum sensing and timing often deliver security impact earlier than quantum computing because they are closer to deployable instrumentation. Their strategic

effect comes from two levers: improving resilience in denied environments and shifting the detection balance.

Improved sensing can create new stresses:

- Detection shifts: in some environments and for some targets, improved measurement can erode concealment and increase the cost of hiding.
- Tempo shifts: resilient navigation and timing can increase operational tempo and compress decision cycles — sometimes stabilizing operations, sometimes increasing crisis brittleness.

A recent media story about a secret tool nicknamed “Ghost Murmur,” reportedly used to help locate a missing airman in Iran, is a useful illustration — even if the details should be treated with caution. The basic idea is simple: use extremely sensitive sensing combined with heavy signal processing to pull a very faint signal out of noise. Whether or not that specific account is fully accurate, it points to the direction of travel: quantum sensing is about making the “barely detectable” detectable under the right conditions. That changes search-and-rescue, surveillance, and the logic of hiding.

At the same time, sensing can reduce stress from adversaries who rely on deception or navigation denial, because defenders gain alternative ways to orient and operate.

As sensing and timing become central to resilience, they become attractive targets. Vulnerability shifts include:

- dependency on timing distribution and calibration chains,
- hardware trust and tamper concerns in fielded sensors,
- integration risk: higher precision can create false confidence if systems assume the sensor is infallible.

Quantum timing and sensing affect consequences in ways that matter beyond any single sector:

- If resilient timing and navigation exist, the consequence of GPS denial is reduced — operations degrade instead of collapsing.
- If detection improves in a specific domain, survivability assumptions may change, potentially affecting deterrence and escalation dynamics.
- If crisis decision time compresses, mistakes can propagate faster and de-escalation windows can narrow.

This is where the perception of quantum shifts most. Thinking of quantum less as “exotic computation” and more as “who owns time and measurement” reveals its deeper security relevance. Once timing and sensing are understood as foundational to coordination under stress, the security implications of quantum stop being speculative.

## 28.6 Governance, Chokepoints, and the Transition Problem: Why “Global” Security Often Means “Who Controls What”

A security analysis that stops at technology misses where global security is often decided: governance and chokepoints. This is where standards, procurement, supply chains, and jurisdiction determine who can deploy capability, who can trust it, and who can be excluded.

Procurement and vendor concentration are not merely operational concerns — they are part of the threat model. If you cannot buy, patch, certify, or replace something during a shock, it does not matter how elegant the design was on paper.

In global security competition, stress is not only kinetic or cyber. It can be economic and institutional:

- denial of access to critical components,
- exclusion from trusted supplier regimes,
- restrictions on collaboration or deployment,
- influence over standards and certification.

Quantum technologies — because they touch trust, timing, and secure coordination — naturally attract these instruments.

Governance choices can reduce vulnerability through assurance and trusted supply, or create it through concentration. Common vulnerability patterns include:

- Vendor concentration: too few suppliers for critical cryptographic or quantum infrastructure components.
- Opaque update and maintenance channels: the ability to push updates is the ability to alter behavior at scale.
- Misaligned incentives: speed to deploy versus assurance to trust.

The Anthropic–US military debate is useful here not because it is quantum, but because it surfaces a recurring security pattern: when a capability becomes strategically sensitive, trust disputes become public, procurement becomes a tool, and arguments about control and failure modes move from theory to policy. Mapped to quantum, the structural questions are familiar:

- Who supplies the cryptographic and quantum infrastructure others must rely on?
- Who controls updates to security-critical modules and nodes?
- What hidden failure modes exist, accidental or intentional, and who can verify them?

- What happens to security when access is restricted during crisis?

This is governance as security engineering: it shapes stress, vulnerability, and consequence just as surely as technical design does.

If there is one mechanism that turns quantum into a global security issue, it is chokepoints — places where control or scarcity translates into leverage. Typical chokepoints include:

- specialized hardware and manufacturing capacity,
- satellite and ground-station access,
- critical fiber routes and secure facilities,
- certification regimes and standards that define “trusted,”
- the talent pipelines that keep systems operable.

The more a society depends on a small set of chokepoints for trust, time, and secure coordination, the more security becomes about protecting those chokepoints and managing dependence on them.

Here, defense is less about acquiring quantum capability and more about reducing systemic fragility:

- diversify critical suppliers and avoid single points of failure in trust infrastructure,
- make update mechanisms auditable and verifiable; reduce the risk of silent control,
- coordinate standards and migration paths to avoid interoperability collapse,
- treat the hybrid transition as a security operation: staged rollout, testing, fallback modes, explicit risk ownership.

## 28.7 Conclusion: A Disciplined Way to Talk About Quantum and Security (Without Hype)

If the chapter’s argument is compressed into one security-grade statement, it is this: quantum technologies change security by shifting secrecy horizons, authenticity assurance, operational resilience, detection balance, and governance leverage — thereby changing the stress–vulnerability–consequence pathway across domains.

The two common mistakes are symmetrical: people either treat quantum as a future “event” rather than a present transition dynamic, or they treat any single quantum technology as a complete solution rather than a subsystem with dependencies.

A pragmatic security view is calmer and sharper: quantum raises the standard for defense and raises the stakes because it touches foundational layers — trust, time,

and visibility. The work of security is to manage that shift so vital functions continue under stress, not only in controlled environments, but in the real world where systems are mixed, people are imperfect, and governance is part of the battlefield.

# Chapter 29

## The Quantum Edge in Finance

### *Inside finQbit's origins*

Tomasz Ćwik

#### 29.1 The World Before the Revolution: When Quantum Was “Science Fiction”

The history of finQbit begins at a time when quantum computing wasn't yet a component of banks' technology strategies, and often didn't even feature prominently in board discussions. The financial world was dominated by topics related to post-financial crisis regulations, capital optimization, and process digitization. Quantum computers existed primarily as an academic curiosity.

In parallel, technologies that today seem obvious, such as LLM models, were developing. When Microsoft experimented with GPT for natural language programming, many experts considered it a futuristic vision. Today, it's difficult to imagine an analyst's work without the support of such tools.

This experience teaches humility in the face of technological breakthroughs. A quote from Neven Hartmut of Google Quantum AI, “It looks like nothing is happening, nothing is happening, and suddenly ... we are in a different world”, perfectly captures the nature of a quantum breakthrough. At finQbit, we've made the assumption that we don't want to be surprised.

---

This chapter is a featured contribution to the monograph *A Portrait of Quantum Technologies in Finance*, edited by Oswaldo Zapata, PhD, and published by The Quantum Finance Boardroom (2026). All rights reserved by the chapter's author.

## 29.2 Roots in Banking: Domain Advantage from Day One

Most quantum companies begin their journey in solid-state physics laboratories, with research teams working on qubit stability, error correction, or superconducting processor architectures. Their first problem is decoherence, not CVA. Their natural habitat is scientific publications and research grants.

FinQbit was founded in a completely different environment than most quantum startups, its roots lie directly in banking. It wasn't a university spin-off or a project stemming from a physics lab. It was the initiative of individuals who had spent years working within large financial institutions, co-creating pricing models, developing computational architectures, implementing solutions compliant with international regulations, and participating in dialogue with regulators and management committees.

This experience of working “from the inside” of a bank, with its processes, business responsibilities and decision-making pressure, has shaped the DNA of finQbit and the way the company approaches building technology.

The founders gained experience at institutions such as BNP Paribas, Moody's Analytics, Credit Agricole, and Deloitte, consistently operating at the intersection of business and technology. They combined the worlds of financial modeling with systems architecture, working on interest rate models, derivatives pricing, contract exposure analysis, credit models, and technological transformation projects in the banking environment.

For years, Tomek served as R&D team leader at BNP Paribas Bank, responsible for the development and testing of groundbreaking technologies in financial applications. In this role, he combined strategic and operational perspectives, defining directions for technological development, initiating research projects, and building teams capable of experimenting at the intersection of innovation and banking practice. It was in this environment that both LLM and quantum computing technologies were explored simultaneously, even before they became widely discussed in the industry.

On the other hand Rafał Pracht is a quantum algorithm architect and financial mathematician, a PhD candidate in Quantum Finance at the Polish Academy of Sciences. He completed the MIT xPRO program in Quantum Computing and is an IBM Qiskit Advocate. He has over a decade of experience in stochastic modeling and derivatives pricing. His research, published in publications including *Wilmott Magazine*, provides the scientific foundation for the technology developed by finQbit, a technology designed to be not only innovative but, above all, mathematically rigorous and scalable.

As you can see, the founders didn't function solely as theoreticians or engineers, their role was to translate mathematics and technological concepts into operational solutions that must function in the real-world environment of a financial institution. They understood the process of approving a new model, the expectations of validation and decision-making teams, and the importance of combining mathematical elegance

with regulatory requirements and business practice. They fully understood that even the most sophisticated mathematical constructs are worthless unless they pass formal approval and integrate into the financial institution's processes.

This genesis led finQbit from the outset to view quantum computing not as an end in itself, but as a potential tool for solving specific banking problems. Instead of asking “how can qubits be used?”, the team asked “which element of a bank's current computing architecture is the most costly, time-consuming, or capital-intensive, and can it be optimized with a quantum approach?”

This approach builds a natural connection with clients. Conversations begin not with a technology presentation, but with a shared process analysis: how long does overnight Monte Carlo take? What are the costs of HPC infrastructure? What are the bottlenecks in derivatives portfolio valuation? How does the risk profile change with increasing market volatility?

Thanks to this, finQbit doesn't have to “explain banking” to its customers, it's part of it. This shared experience bridges the gap and builds trust, which is crucial in a regulated sector.

### 29.3 The Hackathon That Changed Everything: The Birth of the Team

But every startup begins with a meeting of the founders. And this is the story of the Hackathon at the bank. The idea of organizing a quantum hackathon in a bank was an experiment that was supposed to answer one question: are there people in financial structures capable of combining quantum thinking with the practice of risk modeling?

The hackathon wasn't a PR event. The tasks were designed in collaboration with the country's top experts to identify individuals who truly understand quantum computing and can approach the problem creatively.

The winner, Rafał, later founder and CTO of finQbit, presented solutions that not only met the requirements of the task but also proposed an alternative framework for approaching the problem. This wasn't a demonstration of academic knowledge. It was an attempt to translate theory into an implementable architecture.

Following the hackathon, intensive collaboration with the risk teams began. Various approaches were tested, potential accelerations were analyzed, and practical applications within the bank were explored. This was the stage where quantum computing ceased to be an abstract concept and became a topic of real-world business discussions.

The hackathon had another effect, it built the foundation of the organizational culture. It was already clear then that the future team would need to be interdisciplinary. They needed not only physicists and programmers, but also quants and individuals who understood the realities of banking. This combination of skills became finQbit's

DNA.

## 29.4 From Corporation to Startup: The Decision to Leap into the Unknown and Building a New Identity

Leaving the corporate world was one of the most groundbreaking moments in finQbit's history, not only professionally but also mentally. For years, we worked within the structures of large financial institutions: BNP Paribas, Moody's Analytics, Credit Agricole, and Deloitte. We were part of organizations with a global reach, with access to infrastructure, budgets, and teams of specialists.

At the same time, we witnessed the constraints of a regulated financial institution. Innovation at a bank must align with the annual budget, IT roadmap, risk policy, compliance requirements, and regulatory expectations. Each project passes through multi-level decision-making committees. Naturally, response times are longer.

Quantum computing, as an emerging technology, requires a completely different pace, experimentation, rapid prototyping, hypothesis testing, and sometimes failure. At some point, we realized that to fully leverage this technology's potential in finance, a more flexible structure than a bank was needed.

### 29.4.1 Change of mentality, from manager to entrepreneur

Transitioning from a leader and expert role in a bank to that of a startup founder required a shift in mindset. In a corporation, one is responsible for a specific area, technology strategy, team, project budget. In a startup, the founder is responsible for everything:

- technology,
- sale,
- fundraising,
- recruitment,
- marketing communications,
- contract negotiations,
- financial liquidity management.

We had no prior experience building a startup from scratch. We didn't know how to structure a funding round, negotiate a term sheet, build a cap table, or manage a runway with limited resources. We learned by doing, often under time pressure.

At the same time, our corporate experience proved to be a huge asset. We understood how banks think. We knew what documents would be required during the vendor onboarding process, what an IT security audit looks like, and how model validation proceeds.

Many startups emerging from labs lack this knowledge. We had it from day one.

### 29.4.2 Personal and professional risk

The decision to abandon a secure career path meant real risk, financial and reputational. Quantum finance wasn't yet a buzzword. For many, it was a technology "for our children or grandchildren."

We knew that there would be no stable revenues in the first few years, that we would have to convince investors to explore this niche area, and that the market might react skeptically.

It was a decision to step outside of her comfort zone. But it was also a decision to be consistent with her own vision, that quantum computing could be for finance what classical computers were decades ago: a catalyst for entirely new models and instruments.

### 29.4.3 Building culture from day one

From the very beginning, we wanted to build a company that wasn't a classic technology spin-off. FinQbit wasn't meant to be a research lab or a consulting firm.

Our ambition was to create an organization that:

- combines scientific rigor with business pragmatism,
- focuses solely on finances,
- builds solutions with real implementation in the bank in mind,
- treats the client as a strategic partner.

Our organizational culture was shaped by our banking experiences. Respect for processes, responsibility for quality, and awareness of the consequences of a risk model error, these are values we brought from the corporate world to the startup.

At the same time, we strived to maintain the flexibility and speed of operation that large organizations lacked. Short decision-making cycles, direct communication, and a lack of unnecessary bureaucracy, these are the elements that allowed us to experiment and develop technology at a pace impossible within a banking structure.

### 29.4.4 Identity: from a bank, not a lab

One of the most important aspects of this stage was defining our identity. FinQbit is not a company that "looks for applications for technology" We're a team that comes from a banking background and understands its challenges from the inside. We see quantum computing as a tool, a potentially groundbreaking one, but still a tool for solving specific challenges:

- shortening the valuation time,

- reducing computing infrastructure costs,
- improving the accuracy of risk estimation,
- construction of new generations of models.

This fundamentally differentiates us from many quantum startups that start with technology and then seek a market. We started with the market, and only then chose technology as the answer.

### **29.4.5 The first months, intensity and humility**

The first few months of finQbit’s operation were a mixture of euphoria and humility. On the one hand, the excitement of building something of our own, and on the other, the awareness of limited resources and the enormous scale of the challenge.

We had to simultaneously:

- develop a technology platform,
- build relationships with customers,
- apply for acceleration programs,
- seek financing,
- recruit the first team members.

Every decision mattered. Every conversation with a potential investor or client was a lesson.

It was this period that shaped finQbit into a company resilient to uncertainty. Deep tech isn’t a sprint. It’s a marathon with many unknowns. Leaving the corporation was the first step in this marathon, a step that defined the company’s entire subsequent history.

## **29.5 Accelerators, Structure, Pace, and Global Perspective of Deep-Tech Construction**

For finQbit, entering acceleration programs meant more than just earning a prestigious “logo” on the investor slide. It marked the transition from a phase of enthusiastic technological vision to the professional, international development of a deep-tech company.

### **29.5.1 EIT Digital, putting the foundations in order**

Applying to the EIT Digital program was a bold decision, as finQbit was still a very young project at the time. We had banking experience, domain knowledge, and initial technological concepts, but lacked startup experience.

EIT Digital provided us with something incredibly valuable: an outside, critical perspective. The program’s experts weren’t thrilled with the idea of quantum finance itself, they asked tough questions:

- Who will pay and when?
- What does the implementation path look like at the bank?
- What are the specific KPIs for the pilot?
- How do you minimize technological risk on the client’s side?

This was the moment when we had to translate our enthusiasm into a structured action plan. We defined customer segmentation, refined our value proposition, and began building a product roadmap based on realistic milestones.

EIT Digital was a kind of “strategic bootcamp” for us, it helped us understand that in deep tech, it’s not enough to be technologically right. You have to be business-wise.

### 29.5.2 Techstars, MBA at an accelerated pace

Barely after finishing my adventure with EIT, a Techstars representative contacted us. Initially, we took it as a curiosity. Techstars is the second-largest startup accelerator in the world. It accepts fewer than a percent of applications. In the startup world, it’s known as the “Mercedes” of acceleration programs.

Applying was risky, devoting several months to the program meant freezing some operational activities and moving to Berlin for a dedicated deep-tech program. However, we recognized that if we wanted to build a global company, we needed to think beyond the local context.

The deep-tech program in Berlin was extremely intense. Mentor Madness, a series of dozens of meetings in a short period of time, forced continuous refinement of the narrative. Each mentor approached the program from a different perspective: technological, investment, sales, and regulatory.

We had to learn to answer questions in a way that:

- understandable to an investor who is not a physicist,
- credible for a banker who doesn’t trust hype,
- precise for a technology expert.

This was the moment when our story about quantum finance stopped being a story about the future and became a story of evolution, step by step, from an idea to a fully scaled business.

### 29.5.3 Network of relationships and first investments

But Techstars isn’t just about mentoring. It’s above all a global network of relationships. It was there that we met our first investors, who not only understood deep

tech but were also willing to trust a team from Central and Eastern Europe building something as niche as quantum finance.

Many of our mentors had experience building and selling technology companies to banks and financial institutions. Thanks to them, we were able to avoid common mistakes made by young companies, such as underestimating the sales cycle length or regulatory constraints.

For us, Techstars was like an accelerated MBA for founders, with the difference that each day ended with specific decisions that impacted the company's survival.

## 29.6 Deep-Tech Financing, Selectivity, Valuation, and Strategic Courage

Funding a deep-tech startup in quantum finance is a multi-layered challenge. The technology is emerging, the financial sector is conservative, and the payback horizon is long-term.

We understood early on that our path would require a combination of public and private capital.

### 29.6.1 Public funds, fuel for building the first platform

Participating in the PARP competition was a natural step. We weren't afraid of the word "innovation", because what we were building was innovative on a global scale. Taking the podium and securing funding allowed us to finance the initial development of the platform and conduct our first research and development projects.

The grant, however, was just the beginning. In deep tech, public funding can help initially, but it won't replace the private capital needed to scale. That's why we also had to engage in talks with private investors.

### 29.6.2 Two groups of investors, two worlds of misunderstanding

During our conversations with investors, we quickly saw a clear division:

1. **Technology investors (deep-tech)**, understood the potential of quantum computing, but often considered the banking sector too difficult, regulated and long-term.
2. **Fintech investors**, they understood banking, but were afraid of the technological risks and the long horizon of quantum development.

Our task was to find those few who understood both worlds, quantum and finance. This required hundreds of meetings, attending conferences, pitching in various

countries, and constantly fine-tuning communication.

### 29.6.3 Chemistry and Trust, An Underrated Foundation

Many founders talk about valuation. Less often discussed is the chemistry between founder and investor. For us, this was fundamental.

A startup is a bumpy road. Projects don't always go according to plan, experiments fail, and customers delay decisions. In these moments, mutual trust and a shared understanding of the time horizon are crucial.

We rejected an investment offer from a VC fund, even though it represented a real risk to the company's survival. The reason? A lack of long-term strategic understanding and the pressure for a quick return, inappropriate for the nature of deep tech.

### 29.6.4 Valuation, the early stage trap

The second critical element was the shareholder structure. In deep tech, development doesn't end with a single round. There will be subsequent rounds of financing, and too much dilution at the outset can stifle future growth.

We rejected another offer, this time because the investor's stake was excessive. It was a difficult decision, but strategically necessary.

We finally found investors who:

- understood quantum finance as a long-term game,
- they offered fair conditions,
- they had been through the entrepreneurial journey themselves and knew the pressure of building a company from scratch.

This combination of capital, trust and a shared vision has allowed finQbit to build a stable financial foundation without losing control of its strategic direction.

### 29.6.5 Deep-tech is a marathon, not a sprint

Quantum finance requires patience, from both founders and investors. Technology takes time, but it must also deliver value here and now.

Our strategy, combining grants, selective private capital and focusing on real projects with clients, allowed us to avoid two extremes:

- excessive dependence on public financing,
- aggressive dilution in the name of rapid growth.

Thanks to this, we can build finQbit as a long-term company, with the ambition to be a leader in the field of quantum finance, and not just a short-term technological experiment.

## 29.7 Advisory Board and Team Building, Intellectual Capital as a Real Competitive Advantage

While building a company that aspires to be a leader at the intersection of risk modeling and quantum technologies, we quickly realized that technology alone, even the most innovative, is not enough. Quantum finance is an area where substantive credibility is as important as code quality. In a regulated world, reputation and expert authority are the most valuable currency.

Therefore, in parallel with securing financing, we focused on building a strong advisory board. This wasn't about names for investor presentations. It was about real substantive support and the opportunity to consult on technological development directions with people who co-created the foundations of modern financial engineering.

The advisory team included, among others, Professor Dariusz Gatarek, co-creator of the BGM (Brace-Gatarek-Musiela) model, also known as the LIBOR Market Model, which remains the foundation for pricing interest rate instruments in many banks worldwide. Consulting on the architecture of future models with someone who co-created one of the most important mathematical constructs in finance provides a unique perspective. Who better understands the limitations of a model than its co-creator?

Józef Wancer, a legend in the Polish banking sector and former CEO of BNP Paribas and Raiffeisen Bank, among others, has also joined the advisory board. His strategic experience and understanding of the dynamics of managing large financial institutions have brought a managerial perspective to finQbit, crucial when engaging with C-level banking professionals.

Philippe de Brouwer, responsible for risk at HSBC, brought a global perspective on risk management in one of the world's largest banking groups. Zofia Dzik, a supervisory board member and active participant in organizational transformations, brought a strategic dimension and a comprehensive perspective from beyond the sector.

This combination of academic, managerial, and operational experience created a foundation upon which we could build credibility with clients and investors. For the bank, knowing that solutions are consulted with individuals who understand both model theory and the realities of financial institutions significantly reduces risk perception.

At the same time, we were expanding our operations team. From the beginning, finQbit's philosophy was simple: surround yourself with people smarter than yourself. In deep tech, mediocrity doesn't guarantee an advantage. We need people who can combine scientific rigor with engineering pragmatics.

The finQbit team includes individuals with PhDs in physics, academic lecturers, winners of global quantum hackathons, and experts with years of experience in risk modeling on Wall Street. This unique combination of deep theoretical background and market practice allows us to focus exclusively on the intersection of quantum and finance, without distracting ourselves from other industries.

It is this concentration, both in terms of competences and sectors, that constitutes one of the pillars of our competitive advantage.

## 29.8 Product Strategy, Hybrid as a Responsible Path to the Future

One of the biggest mistakes you can make in quantum computing is assuming the revolution will be instantaneous. The history of technology shows that breakthroughs often come in leaps and bounds, but preparing for them is a long-term process.

FinQbit has adopted an evolutionary strategy from the outset. Instead of waiting for full “quantum advantage” in banks’ production environments, we are building a quantum-ready platform that delivers value today by optimizing classical computations and preparing the architecture for future QPU integration.

In the field of Monte Carlo, it is known that quantum algorithms, such as Quantum Amplitude Estimation, theoretically offer quadratic speedup. However, translating this theory into practice requires a deep understanding of both quantum algorithms and the structure of financial problems.

It’s not enough to implement an algorithm. You need to know how to embed it into the risk model architecture, how to reduce the cost of QPU calls, how to optimize data representation, and how to integrate the results into the bank’s existing systems.

That’s why our platform operates in a hybrid mode. We leverage HPC and GPUs to accelerate classical computations today, while simultaneously building an abstraction layer that transparently allows us to offload more and more computation to quantum hardware as it matures.

Importantly, we remain hardware agnostic. We don’t know which architecture, superconducting, ionic, photonic, or otherwise, will emerge as dominant and cost-effective. Therefore, we enable benchmarking across vendors and minimize the risk of vendor lock-in for the customer.

Our goal isn’t to “push” quants into the world of quantum computing. Our goal is for quantum computing to serve quants, in their language, with their definitions, and in the context of their processes.

## 29.9 Quantum Machine Learning, Between Hardware Limitations and Scientific Ambition

The second pillar of finQbit’s business is Quantum Machine Learning (QML). This area is particularly attractive to the financial sector, where classic ML models are already widely used, from credit scoring to volatility prediction.

However, QML faces a significant challenge: the volume of data. Business problems

in banking operate on large data sets. Current quantum devices have a limited number of qubits and a high level of noise.

Therefore, designing efficient architectures and optimizing software becomes crucial. Our CTO's decision to build the platform in Julia, contrary to the prevailing Python trends, was a strategic one. Julia offers performance similar to C while maintaining a high level of abstraction.

The results were measurable. In one project for a top European financial institution, we achieved ten times greater performance than other platforms. This allowed us to process larger volumes of data and apply QML to a real-world business problem.

Moreover, the quantum neural network model outperformed the classic XGBoost model in a specific prediction task. This isn't proof of QML's overall superiority, but rather a sign that the technology can offer real value under certain conditions.

However, we didn't avoid costly lessons learned. A single night of model training on real hardware cost approximately \$30,000, which caught the attention of AWS headquarters in the US. This experience demonstrated the importance of a hybrid approach and cost control.

## 29.10 Creative Destruction Lab, Global Validation in the Most Demanding Environment

You can have an ambitious vision, advanced technology, and deep domain knowledge. However, without external validation by the most demanding experts, it's difficult to realistically consider conquering global markets, especially in such a competitive and capital-intensive field as quantum computing.

Therefore, participating in the Creative Destruction Lab (CDL) program in North America was a natural step for finQbit toward international validation of its technology and strategy. The CDL Quantum Stream, run in partnership with the University of Toronto (Rotman School of Management) and HEC Montréal, is one of the most prestigious and selective acceleration programs for quantum startups in the world.

The selection process is rigorous and multi-stage. Hundreds of companies from around the world apply for the program, both spin-offs from leading universities and teams backed by renowned deep-tech funds. The evaluation encompasses not only the quality of the technology but also its scalability, commercial potential, competitive advantage, and the team's ability to implement an ambitious roadmap. Acceptance into the program itself constitutes a form of international validation.

The CDL framework is based on clearly defined objectives set for each round. Companies are systematically evaluated by mentors and investors, and progression to the next stage is not guaranteed, it requires proof of real technological and business progress. This is an environment where declarations are irrelevant, what matters is evidence, numbers, and concrete milestones.

The program's mentors are world-class entrepreneurs, investors, and scientists, in-

cluding Nobel Prize winners in physics and leaders of global technology companies. The discussions are conducted at a level that requires not only excellent substantive preparation but also the ability to defend the adopted strategic assumptions. Each element is subject to intensive analysis.

Completing the CDL Quantum Stream program among six global companies was a breakthrough for finQbit. It meant that its strategy focused on the financial sector, rather than building a generic “quantum stack”, and its hybrid approach (HPC + quantum) were recognized as coherent, realistic, and scalable in an international context.

CDL isn’t a marketing or networking program in the traditional sense. It’s an environment where startups are confronted with brutally honest feedback. If something isn’t working, it’s communicated directly. If the business model isn’t compelling, it’s deconstructed. If the technology doesn’t have a clear path to commercialization, the questions are relentless. For finQbit, participating in CDL was not only a validation but also a catalyst for organizational maturity. The program forced them to precisely define their market position, refine customer segments, and clearly map their competitive advantage over traditional solutions and other quantum startups.

CDL Quantum Stream alumni status has become one of the key milestones in the company’s history, a symbol of the transition from an ambitious European startup to a group of globally recognized entities in the quantum for finance space.

It is also a clear signal to partners and financial institutions: finQbit’s strategy has been tested in one of the most demanding innovation ecosystems in the world, and it has passed the test.

## 29.11 Customer Relationships, Rebuilding Trust in a Post-Hype World

One of the biggest challenges we encountered was banks’ hesitation towards quantum technology. Importantly, this wasn’t solely due to hardware limitations or the immaturity of current quantum processors. In many cases, it stemmed from previous experiences with vendors who promised a revolution but delivered only demonstrations. The financial sector, especially in the risk sector, has a long memory for unfulfilled promises.

Banks operate in a highly regulated environment, under pressure from capital and reputation. Every new technology must undergo a complex evaluation process, from business teams and IT to validation and compliance units. In such an environment, slogans like “quantum advantage” without hard data evoke skepticism rather than enthusiasm.

We realized very quickly that our biggest competitor was not another quantum startup, but a lack of trust in the entire technology category.

That’s why we’ve embraced the principle of radical transparency from the outset.

We don't sell empty promises. We don't promise "here and now" computational superiority if we can't demonstrate it in a controlled environment. We define realistic pilot goals, clearly communicate the limitations of current NISQ devices, and design experiments so that their results, positive or negative, are measurable and useful for business.

This methodology changes the dynamics of the relationship. Instead of presenting a "quantum solution" we conduct a research experiment with the client in a controlled business environment. As a result, the bank feels less like a marketing recipient and more like a co-creator of the verification process.

Our approach resulted in collaboration with the global risk center of one of Europe's largest financial groups, a unit responsible for models used internationally. This environment is characterized by mathematical rigor, methodological documentation, and formal validation of every model element.

Simultaneously, our research was published in *Wilmott Magazine*, one of the most recognizable journals in the quantitative finance community, which sent an important signal that our work meets the substantive standards expected by the quant community. Our appearances on the CQF Institute stage also allowed us to confront our findings with market practitioners, individuals responsible for pricing and risk models in global institutions on a daily basis.

This demonstrates that a strategy that focuses on the real problem, not the technology, builds lasting relationships and long-term credibility. Banks aren't looking for a "quantum revolution" They're looking for stable, measurable improvements in areas like Monte Carlo, XVA, and derivatives pricing amidst growing market volatility.

By positioning itself as a partner in the experimentation and validation process rather than a provider of a trendy solution, finQbit is gradually changing the perception of quantum computing in the financial sector, from a hype category to an engineering tool that can have concrete applications in well-defined use cases.

In the regulated sector, trust is a currency more valuable than capital. Building it has proven to be one of the most crucial elements of our journey.

## **29.12 Business Model and Long-Term Vision: Quantum Finance as a New Layer of Financial Infrastructure**

Importantly, in the path of creating global innovations, one cannot rest on one's laurels at any stage. A deep-tech startup can't rely solely on grants, accolades, and media attention. The ultimate test of any technology is its ability to generate revenue and solve real customer problems. In the field of emerging technologies, such as quantum computing, this challenge is twofold: it must simultaneously build a market and deliver value here and now.

The finQbit strategy is based on this dual perspective.

On the one hand, we deliver measurable value today, through the optimization of classical computation and hybrid models, and the parallel development of quantum competencies on the client side. On the other hand, we are consistently developing a long-term advantage in the area of fully quantum financial models, which will be able to take over an increasing share of critical computations as hardware matures.

In other words, we are building an evolutionary path from classical financial engineering to hybrid architecture, and ultimately to quantum infrastructure.

We focus on two fundamental areas that constitute the core of modern financial engineering:

- Monte Carlo calculations,
- development of Quantum Machine Learning for risk modeling applications.

These are not random choices. Monte Carlo is at the heart of derivatives pricing, XVA, scenario simulations, and capital management. Machine learning is increasingly influencing credit models, anomaly detection, volatility forecasting, and high-frequency data analysis. If quantum computing is to become a viable part of the financial infrastructure, it must enter precisely these areas, where the scale of computation is largest and the cost of error is highest. In the following chapters, we will explain the principles of Quantum Monte Carlo and Quantum Machine Learning.

Our solutions are designed to evolve with hardware development. Today, they operate in a hybrid model. Tomorrow, as device quality improves, this proportion may change. The architecture we build is future-proof, not dependent on a single hardware generation.

Most importantly, however, finQbit isn't building a generic platform "for all industries". We're not trying to be another universal AI tool or quantum stack provider. We're 100% focused on finance.

This specialization implies a deep understanding of interest rate models, XVAs, regulatory requirements, validation processes, and the architecture of banking IT risk systems. It also implies a clear identity, we are a company that builds quantum finance, not "general purpose quantum applications".

In this sense, finQbit is pioneering the unknown. Just as artificial intelligence pioneers built language models before the market fully understood their potential, we are building a computational layer that could become the next stage in the evolution of financial infrastructure.

For us, quantum finance isn't a theoretical experiment or a marketing slogan. It's a natural continuation of the history of financial engineering, just as classical computers enabled the development of complex derivatives and global capital markets decades ago. Today, we stand on the threshold of another transformation.



# Chapter 30

## The Coverage Banker

*A story of unlearning to relearn*

Tse Loong Chin

### 30.1 The Humble Ledger

My story doesn't begin on a trading floor or in a boardroom. It begins in the quiet, methodical world of trade finance operations. In that role, I was a guardian of details, a processor of paper. Every day, I navigated the intricate dance of letters of credit, bills of lading, and documentary collections. It was a world governed by precise rules, where a single misaligned date could halt a shipment worth millions. It was a humble beginning, but it taught me a foundational lesson that has never left me: discipline. The discipline to check, to verify, and to ensure that the gears of global commerce turned smoothly, one transaction at a time.

But even in that world of paper, I felt the first stirrings of curiosity. What if we could make this faster? What if we could make it less prone to error? That curiosity was a quiet whisper, one I decided to follow.

### 30.2 The Architect of Automation

That whisper led me from operations to the dynamic world of project management. I joined a team dedicated to automating the very mechanisms I had once manually processed. This was my first real experience in “unlearning.” I had to unlearn the comfort of routine and embrace the uncertainty of building something new—the

---

This chapter is a featured contribution to the monograph *A Portrait of Quantum Technologies in Finance*, edited by Oswaldo Zapata, PhD, and published by The Quantum Finance Boardroom (2026). All rights reserved by the chapter's author.

automation of trade finance processes via FITAS, a legacy trade finance system from the 1990s, managed through the project management waterfall methodology.

This technical and structural understanding of banking machinery then naturally evolved into a more strategic role: trade finance advisory. Here, I stepped out of the engine room and onto the bridge. I wasn't just building the ship; I was helping to chart its course for clients. I advised on structuring deals, mitigating risks, and finding opportunities within the complex web of international regulations. This was a deeper form of unlearning—moving from process to purpose.

### 30.3 The Banker and the Break

The next logical step was to take my operational and advisory knowledge directly to the client. As a Credit & Marketing Officer, a corporate and commercial coverage banker, I became the trusted face of the bank. I learned the art of the relationship, the nuance of negotiation, and the critical responsibility of credit assessment. I saw the business plans behind the balance sheets, the ambitions behind the applications. This phase was about preservation—preserving the bank's capital while nurturing a client's growth. It was a high-wire act of trust and prudence.

Then came 2007, a year that would mark a 180-degree shift in my career trajectory. I made the leap from corporate and commercial banking into the world of institutional coverage. Suddenly, my clients were no longer businesses in the traditional sense; they were banks, brokerages, insurance companies, asset managers, and even central banks. It was a different language, a different pace, and a different set of complexities. This was my first real taste of the FIG (Financial Institutions Group) world, long before I would return to it after my break. The experience planted a seed, one that would lie dormant for years but never truly fade.

Then, in 2020, came the career break. A full year away from the banking treadmill. To some, a career break might be a pause, a rest. For me, it was a launchpad. I called it my “unlearning and relearning journey.” I looked at my 20+ years in finance and realised that while my experience was deep, the world around it was shifting seismically. I made a conscious decision: I would not just update my skills; I would overhaul my mindset.

I didn't scroll through social media; I devoured knowledge. I looked at the projects I had led and formalised that expertise, earning my CSM and CSPO certifications, solidifying my project management identity. The key takeaway was a new perspective: I wasn't just processing trade; I was deconstructing it, understanding its logic down to the smallest user story, and then rebuilding it with code. I learned to facilitate, to prioritise, to bridge the gap between the bankers who knew the “what” and the developers who knew the “how.” This phase taught me diligence. It wasn't just about doing the job right; it was about designing the job to be right, thinking ten steps ahead for the user.

But I wanted to go deeper. I looked at the bank through the eyes of a regulator and completed ICA UK's Advanced Certificate in Regulatory Compliance. This gave me

a formal risk lens, not just from the first line of defence, but from the perspective of the rule-makers themselves. This was about dedication—a dedication to seeing my industry from every possible angle.

I saw the world waking up to a new existential threat and a new opportunity. I studied Sustainability, Financing, and Climate Risk, earning certifications from the Chartered Banker Institute and the GARP Sustainability and Climate Risk (SCR) program. I was learning to see risk not just in a credit file, but in the physical world and the transition to a greener economy.

And then, there was the rabbit hole of AI and Blockchain. This wasn't just about fintech hype; it was about the fundamental rewriting of trust and efficiency in finance. I devoured courses, white papers, and thought leadership. I saw that my world of trade finance, with its complex multi-party trust issues, could be fundamentally reshaped by these technologies. This wasn't a break from banking; it was an investment in its future and my own.

## 30.4 The Institutional and the Spark of Quantum

Returning from my break, I was a different banker. I stepped into the world of covering banks: financial institutions are the ultimate relationship game. It requires a deep understanding of their business models, their regulatory pressures, and their strategic goals. I moved between coverage, product sales, and back to coverage, each rotation strengthening my ability to not just offer a product, but to provide a solution that fit the intricate puzzle of their institution.

My journey of learning never stopped. It became part of my identity: the discipline to maintain my curiosity. And that curiosity eventually led me to the edge of the map: quantum computing.

It started with a question: What happens to our impregnable encryption, the very foundation of digital trust in banking, when a sufficiently powerful quantum computer exists? What happens to the blockchain solutions I had been studying? What happens to our models for complex risk calculations?

My immediate reaction wasn't panic; it was the familiar thrill of a new puzzle. I knew I didn't need to become a quantum physicist. I didn't need to write quantum code. What I needed was to understand the application and its impact. I needed to understand the “so what?” for finance.

I began to study the basics: superposition, entanglement, qubits. I learned the difference between quantum annealing and gate-based models. More importantly, I started following the work of organisations like the GARP Risk Institute on quantum risk, reading papers on post-quantum cryptography, and understanding the timeline and the threats.

## 30.5 Drawing from Mentors' Paths

This journey of continuous reinvention is not one I walked in isolation. Along the way, I have been profoundly influenced by individuals whose resilience, perseverance, and vision have illuminated my path.

I found great inspiration in the story of Eric Sim, a fellow banker who transitioned from a successful career in institutional banking to becoming a renowned author, a LinkedIn influencer with millions of followers, and a self-entrepreneur. Eric Sim's story resonated with me because he didn't just change jobs; he transformed his identity. He took the deep skills he cultivated in banking relationship building, communication, complex problem-solving and applied them to a new canvas. He proved that a banker's expertise is not confined to a bank. It is a portable toolkit of insights about human nature, markets, and strategy.

I have also drawn strength from observing Royce Wee, law trained and technologist, whose thoughtful commentary on public policy and leadership consistently demonstrates the power of reflective practice. The works of Dr. James Ong, Andeed Ma and Siok Siok Tan authors of *AI for Humanity*, have expanded my understanding of how artificial intelligence can serve a greater good. While their entrepreneurial journeys have shown me the true meaning of resilience and perseverance, seeing them navigate challenges with unwavering determination has been a lesson in itself.

I must also express my sincere gratitude to my Head of ICG Singapore, Guan Sim. She has been incredibly supportive of my journey from pursuing AI and becoming an AI Ambassador for HSBC, to venturing into digital assets and cryptocurrencies ("DAC"). Her trust and empowerment have given me the space to explore, to question, and to grow. It is not every day that one gets a leader who actively encourages you to chase the frontiers of your curiosity. For that, I am truly thankful.

Most significantly, I must acknowledge the profound influence of Philip Intallura, Global Head of Quantum Technologies at HSBC. Philip generously shared his expertise in ways that ignited my understanding of why quantum matters for finance. Through the virtual sharing session, he showed me that one doesn't need to be a quantum physicist to grasp the strategic implications. His insights transformed my casual curiosity into a structured pursuit of knowledge, and he is the reason I can now call myself a quantum-ready banker.

And then there is Yun Er, my closest friend, the one who has stood unwaveringly behind me through all of this. Just a quiet, steady presence that has given me strength on days when the path felt uncertain. Yun Er has been my sounding board, my anchor, and my biggest believer. I am deeply grateful for Yun Er's support.

There are, of course, many others who have helped me along the way, too many to name, but never too many to forget. To everyone who has offered a word of encouragement, shared a piece of advice, or simply walked beside me during uncertain times, I carry your kindness with me. This journey is as much yours as it is mine.

Each of these individuals has contributed a thread to the tapestry of my thinking. They have shown me that the path is never walked alone, and that the greatest

insights often come from the generosity of those who share their light.

## 30.6 The Preservation of Curiosity

My story is not a straight line from A to B. It is a constellation of experiences connected by a single, consistent thread: the preservation, diligence, dedication, and discipline to maintain my curiosity.

- **Preservation:** Preserving the foundational lessons from my trade finance ops days, remembering that even the most complex structures are built on simple, accurate transactions.
- **Diligence:** The diligence I learned as a project manager, ensuring that every new idea is examined from all angles and built on a solid foundation of requirements.
- **Dedication:** My dedication during that one-year break, investing in myself to understand the tectonic shifts in my industry, from climate risk to AI.
- **Discipline:** The daily discipline to keep learning, to ask “what’s next?”, and to have the courage to step into the unknown, even into the strange, probabilistic world of quantum mechanics.

