



universität
wien

MASTERARBEIT | MASTER'S THESIS

Titel | Title

How Quantum Computing Escapes the Maze
Microsoft's Case of Narrative Construction and the Limits of Knowledge in
Emerging Technologies

verfasst von | submitted by

Mag.rer.soc.oec. Stephan Scipal

angestrebter akademischer Grad | in partial fulfilment of the requirements for the degree of
Master of Arts (MA)

Wien | Vienna, 2026

Studienkennzahl lt. Studienblatt |
Degree programme code as it appears on the
student record sheet:

UA 066 906

Studienrichtung lt. Studienblatt | Degree pro-
gramme as it appears on the student record
sheet:

Masterstudium Science-Technology-Society

Betreut von | Supervisor:

Univ.-Prof. Dr. Ulrike Felt

KURZFASSUNG

Diese Diplomarbeit untersucht, wie Microsoft durch Narrative Ansprüche auf technologische Reife im Bereich Quantencomputing konstruiert und legitimiert. Labore für Quantencomputer sind physisch unzugänglich und epistemisch undurchdringlich, da das dort produzierte Wissen fast ausschließlich durch insidervermittelte Berichte zirkuliert. Externe Beobachter*innen können die technischen Behauptungen, die Unternehmen über Fortschritte, Zeitpläne oder Durchbrüche aufstellen, nicht unabhängig überprüfen. Diese Strukturbedingung, die in dieser Arbeit als narrative Abhängigkeit bezeichnet wird, bildet den zentralen Gegenstand der theoretischen und empirischen Untersuchung. Narrative Abhängigkeit entsteht, wenn vier Bedingungen zusammentreffen: physische Unzugänglichkeit; technische Spezialisierung, die Verständnisbarrieren erzeugt; Unternehmenskontrolle, die Transparenz einschränkt; sowie öffentliche Bedeutsamkeit, die erheblichen Druck auf eine akkurate Wissensproduktion ausübt. Unter diesen Bedingungen fungieren insiderproduzierte Narrative als die primäre verfügbare Ressource, durch die externe Beobachter*innen überhaupt ein Verständnis der technologischen Entwicklung gewinnen können. Die Arbeit entwickelt narrative Abhängigkeit als ein originäres epistemologisches und politisches Konzept innerhalb der STS. Sie baut auf Czarniawskas (2004) Einsicht auf, dass Organisationen narrative Wirklichkeiten konstituieren, sowie auf Knorr Cetinas (1999) Analyse epistemischer Kulturen — und erweitert beide, um diese spezifische Situation zu adressieren. Darüber hinaus wird narrative Abhängigkeit in Beziehung gesetzt zu Latour und Woolgars (1986/1979) literarischer Inskription, Hilgartners (2015) Vanguard-Framework, der Soziologie der Erwartungen (Borup et al., 2006; van Lente & Rip, 1998), soziotechnischen Imaginären (Jasanoff & Kim, 2009), Boundary Work (Gieryn, 1983), Collins' (1992/1985) Experimenters' Regress sowie der Performativitätstheorie (Austin, 1962; Callon, 1998; MacKenzie, 2006). Die Arbeit analysiert drei Dokumente von Microsoft auf Grundlage der Methode der praxisorientierten Dokumentenanalyse von Asdal und Reinertsen (2022). Als diachronische Sequenz konstituieren diese Dokumente eine Vanguard-Vision, die Hilgartners synchronen Rahmen temporal erweitert. Die vier Beiträge der Arbeit — narrative Abhängigkeit als transferierbares Konzept, diachronische Vanguard-Analyse, integrierte diskursive Legitimationsstrategien und strategische Imaginärdynamiken — bieten analytische Ressourcen für das Studium aufkommender Technologien unter Zugangsbeschränkungen. Die Ergebnisse zeigen: technologische Reife ist keine objektive Eigenschaft, sondern eine narrative Errungenschaft, die materielle Folgen in Form von Investitionen, politischer Aufmerksamkeit und öffentlicher Legitimität erzeugt.

Thank you ...

... to everyone who accompanied me on this journey. Nobody writes a thesis alone.

First and foremost, I wish to thank my supervisor, Prof. Ulrike Felt, for her invaluable feedback and scholarly guidance throughout this project. I am also grateful to Prof. Sarah Davies, whose seminar first sparked the curiosity that eventually grew into this thesis — without her encouragement to develop that initial seminar paper further, this work would not exist.

My deepest personal gratitude goes to my partner, Alexandra Hanzl, who supported me mentally through the most difficult moments, motivated me to continue my studies when my resolve wavered, and contributed practically to the research by assisting with the transcription of Troyer's colloquium. Her presence throughout this process has been indispensable.

I am grateful to Anna Sophie Windisch for her assistance in methodical data collection from Microsoft's YouTube archive.

For their careful and generous feedback on the final version of the thesis, I thank my dear brother Klaus Scipal and my dear friends Robert Schanzer, Alexander Schmidt, and Ian Smith.

The research proposal that preceded this work was shaped by the thoughtful responses of post-doc researchers Fredy Mora Gámez and Andrea Schikowitz, and fellow students Paul Wunderlich and Michaela Zuckerhut, to whom I am sincerely grateful.

I also wish to thank the fellow students who accompanied me through the seminar *Coping with Trouble* and offered feedback at various stages of the thesis: Rachel Bell, Roxana Demeter, Sofie Klevstrand, Kerstin Mitterhuber, and Pablo Valenzuela-Astudillo. Your engagement with this work made it considerably stronger.

Finally, I extend my gratitude to the teachers, fellow students, and staff of the Department of Science and Technology Studies at the University of Vienna for creating a safe space in which to engage with STS scholarship and to think openly about all that it entails.

TABLE OF CONTENTS

1 Introduction: Quantum Computing’s Narrative Dependency Problem1

1.1 The Epistemic and Political Challenge of Laboratory Confinement1

1.2 Research Questions and Approach3

1.3 Thesis Roadmap5

2 Theoretical and Conceptual Framework: Situating Narrative Dependency in STS
Scholarship.....7

2.1 Quantum Computing as an Analytical Object7

2.2 The Laboratory Confinement Problem and Narrative Dependency8

2.3 The State of STS Research on Quantum Computing 10

2.3.1 Sociotechnical Imaginaries and the Construction of Quantum Futures 10

2.3.2 Science and Technology Policy and the Sociology of Expectations 11

2.3.3 Responsible Research and Innovation 12

2.3.4 Sociotechnical Vanguard and Vision-Building..... 13

2.3.5 The Sociology of Scientific Knowledge and Credibility..... 14

2.3.6 What Existing STS Research Leaves Unaddressed 15

2.4 An Analytical Framework for Narrative-Dependent Technologies 16

2.4.1 The Constitution of Facts: Narrative as Practice 16

2.4.2 The Dynamics of Vision-Building: Vanguard and Expectation Management..... 17

2.4.3 The Mechanisms of Legitimation: Boundary Work and Performativity 18

2.4.4 The Construction of Futures: Imaginaries and Temporal Framing 18

2.4.5 Framework Integration: A Shared Constructivist Commitment..... 19

2.5 This Thesis’s Contribution 20

3 Methodology: Analysing Narratives Under Conditions of Limited Access22

3.1 Document Analysis as Analytical Consequence of Narrative Dependency.....22

3.2 Case Selection: Three Texts, Eleven Years.....23

3.2.1 Microsoft’s Quantum Communication Landscape23

3.2.2 Selection Logic and Search Process.....24

3.2.3	Why Three Documents Are Sufficient	25
3.2.4	The Three Documents	27
3.3	Analytical Tools and Techniques	28
3.3.1	Primary Analytical Framework	28
3.3.2	Use of Large Language Models as Interpretive Aids	29
3.3.3	Transcription.....	30
3.4	Reflexive Positioning and Limitations	30
3.4.1	My Positioning Revisited	30
3.4.2	Methodological Limitations	30
4	The Frontier Imaginary: Station Q (2014).....	32
4.1	Context and Material Form	32
4.2	Critical Assessment of Issues	34
4.3	Narrative Analysis: Five Stories of Quantum Pioneering	36
4.3.1	Story 1: "The Quest for a Quantum Future"	36
4.3.2	Story 2: "Particles Gone Wild"	37
4.3.3	Story 3: "Now Arriving at Station Q"	39
4.3.4	Story 4: "Hunting Quasiparticles"	40
4.3.5	Story 5: "From Here to Infinity"	42
4.4	STS Interpretation: Constructing the Vanguard	43
4.4.1	Constructing Authority at the Frontier Stage	44
4.4.2	Constructing the Expansive Imaginary	45
4.4.3	Performing the Vanguard into Existence	46
4.5	Synthesis: The Frontier Imaginary	47
5	Managing Expectations: Navigating the Quantum Maze (2021).....	49
5.1	Context and Material Form	49
5.2	Critical Assessment of Issues.....	51
5.3	Narrative Analysis: At the Entrance to the Maze	53
5.3.1	Opening: Establishing Motivation and Identity	53

How Quantum Computing Escapes the Maze

5.3.2	Defining the Quest: "Practical Quantum Advantage"	54
5.3.3	Introducing the Maze: Quantum Applications Landscape	55
5.3.4	Establishing the Methodology: The Race Framework	56
5.3.5	Debunking Big Data: The I/O Problem.....	57
5.3.6	Addressing QRAM: Dismissing Theoretical Solutions	58
5.3.7	The Constant Slowdown Problem	58
5.3.8	Eliminating Quadratic Speedups: Grover's Algorithm Dismissed	59
5.3.9	Clearing the Maze: Dropping Out Applications	59
5.3.10	The Leaf in Spring: Quantum Simulation as Viable Path	60
5.3.11	From Analogue to Digital: The Scalability Argument	61
5.3.12	Chemical Reactions: Barrier and Pathway Metaphors.....	61
5.3.13	Quantum Tunnelling: Optimisation Metaphor.....	62
5.3.14	Temporal Framing: Multiple Future Horizons	63
5.3.15	Addressing NISQ: Dismissing Competitors' Approaches.....	63
5.3.16	Benefits Along the Journey: Classical Computing Improvements	64
5.4	STS Interpretation: Escaping the Quantum Wishing Well	65
5.4.1	Reconstructing Authority Through Expert Assessment	65
5.4.2	Managing Expectations and Constraining the Imaginary	66
5.4.3	Performativity, Narrative Dependency, and the Absent Laboratory	67
5.5	Synthesis: The Pragmatic Turn	69
6	Claiming Breakthrough: Majorana 1 (2025)	71
6.1	Context and Material Form.....	71
6.2	Critical Assessment of Issues	73
6.3	Narrative Analysis: From Maze to Milestone	76
6.3.1	Opening: The Revolutionary Invention	76
6.3.2	Part Two: Scientific Vindication.....	77
6.3.3	Part Three: Transformative Applications.....	79
6.4	STS Interpretation: Materialising the Future	80

6.4.1	Constituting Breakthrough: Literary Technologies and Authority Accumulation ..	81
6.4.2	Strategic Imaginary Re-Expansion.....	82
6.4.3	Performativity, Scientific Controversy, and the Limits of Narrative Authority	83
6.5	Synthesis: The Breakthrough Imaginary.....	85
7	Discussion: The Evolution of Quantum Narratives and the Limits of Knowledge	87
7.1	Synthesising Narrative Evolution.....	87
7.2	Addressing the Primary Research Question.....	90
7.3	Theoretical Contributions.....	91
7.3.1	Contrib. 1: Narrative Dependency as Epistemological and Political Concept.....	91
7.3.2	Contribution 2: Diachronic Vanguard Analysis.....	92
7.3.3	Contribution 3: Integrated Discursive Legitimation Analysis.....	93
7.3.4	Contribution 4: Strategic Imaginary Dynamics.....	94
7.4	The Breakthrough Controversy Revisited	95
7.5	Implications for STS Methodology.....	96
7.5.1	Document Analysis Under Access Constraints.....	96
7.5.2	Productive LLM Use.....	97
7.5.3	Insider-Outsider Positioning	97
7.6	Closing the Discussion.....	97
8	Conclusion: Towards Better Understanding of Emerging Technologies	99
8.1	Summary of Findings.....	99
8.2	Limitations.....	100
8.3	Future Research Directions	101
8.4	Broader Implications.....	103
8.5	Final Reflection	103
	Appendix I – Transcript of Troyer’s Colloquium	105
	Appendix II – Questions to Large Language Models on Quantum Computing.....	125
	References	128

1 INTRODUCTION: QUANTUM COMPUTING'S NARRATIVE DEPENDENCY PROBLEM

Quantum computing laboratories are, for many observers, a maze. The technology is physically inaccessible — housed behind cryogenic barriers a hundredth of a degree above absolute zero — and epistemically opaque, producing knowledge that circulates only through insider-mediated accounts. William Henry Matthews (2014/1922), in his classic study of maze structures, distinguishes between two fundamentally different relationships to a maze: the aerial view, from which the solution is obvious, and ground-level navigation, where confusion reigns and false paths proliferate. Microsoft's quantum researchers occupy the aerial view; external observers — including STS scholars, policymakers, and the public — are navigating from the ground. Corporate narratives function, in this context, as Ariadne's thread: the insider-produced signposts that external observers must follow, because no independent path through the maze exists.

But this thesis's central argument goes further than pointing to an epistemological inconvenience. It argues that Microsoft's quantum computing narratives actively *construct* claims of technological readiness under conditions where those claims cannot be independently verified — and that this construction has political as well as epistemic consequences. Between 2014 and 2025, Microsoft produced a remarkably smooth account of its own progress: from frontier exploration, through pragmatic expectation management, to the announcement of breakthrough. This thesis examines how that smoothness is made, what it suppresses, and what it reveals about the conditions under which emerging technologies acquire public legitimacy. The analysis does not ask whether the narrative leads to the exit; it asks who built the maze and how the walls were arranged.

1.1 THE EPISTEMIC AND POLITICAL CHALLENGE OF LABORATORY CONFINEMENT

Quantum computing has been presented, by corporations and governments alike, as a transformative technology: one capable of breaking encryption systems, accelerating drug discovery, modelling climate interventions, and reshaping global security. These are claims with substantial public consequences. Billions in private investment and public funding flow in response to them (Soller et al., 2025). Policy frameworks addressing quantum security, quantum workforce development, and quantum supply chains are already being designed in anticipation of capabilities that have not yet been demonstrated at scale. Yet the knowledge on which such frameworks rest is produced almost entirely within corporate laboratories that external observers cannot enter, scrutinise, or independently evaluate. A technology that

promises to shape societies cannot be meaningfully accessed by those societies — a condition that is both epistemologically and politically consequential.

I term this condition *narrative dependency*. It arises when four circumstances coincide: the technology under study remains physically inaccessible to external observers; the expertise required to evaluate its progress is asymmetrically concentrated among insiders; the primary producer of public knowledge about the technology is simultaneously a commercial actor with interests in particular assessments of its progress; and the technology's claimed societal significance creates substantial public stakes in accurate understanding. Under these conditions, external observers — scholars, policymakers, and citizens alike — are compelled to rely predominantly on insider-produced narratives, not as one option among others but as the primary available epistemic resource.

The concept draws on Czarniawska's (2004) insight that narratives in organisational contexts do not merely describe realities but actively constitute them, and on Latour and Woolgar's (1986/1979) analysis of how literary inscription transforms laboratory observations into publicly circulating facts. Neither framework, however, directly addresses the situation in which ethnographic access to the laboratory is structurally foreclosed rather than merely difficult. Narrative dependency names that specific condition and identifies it as the central methodological and analytical challenge facing STS scholarship on quantum computing. As an epistemological concept, it characterises what external observers can and cannot know; as a political concept, it characterises who controls the terms in which the technology is publicly understood; and as a methodological concept, it identifies the analytical consequence — systematic document analysis — that follows when other forms of access are unavailable.

The technical requirements of quantum computing — including ultra-low temperature conditions “a hundredth of a degree above absolute zero” (Rietsche et al., 2022) and highly specialised equipment — constitute one barrier to outside access. Corporate proprietary interests constitute another. Microsoft's strategic partnership with the United States military in quantum computing research (Boyle, 2023) illustrates how geopolitical significance compounds these barriers further: access restrictions reflect not merely organisational preference but the classified status of research at the nexus of national security and scientific advance.

The political stakes of this condition deserve explicit acknowledgement. When technologies of claimed civilisational importance are funded by public resources, shaped by public policy,

How Quantum Computing Escapes the Maze

and justified through appeals to public benefit, yet subject to no independent epistemic scrutiny, a structural problem of democratic governance arises. The struggle over interpretive authority — over what counts as a breakthrough, what counts as hype, whose assessment is credible — is not merely an academic dispute about classification. It shapes the conditions under which quantum computing acquires, or fails to acquire, social legitimacy; and it does so in ways that favour those with privileged access to the laboratory over those who must govern the technology from outside it.

My dual positioning — as an STS scholar and a Microsoft employee with no direct involvement in the company's quantum computing programme — shapes how I read the documents. It creates both interpretive affordances and analytical constraints that require careful reflexive engagement. The relevant methodological context for this positioning, including its implications for document selection, potential bias, and the maintenance of analytical distance, is addressed fully in chapter 3.4. Readers are invited to consult that section if they wish to assess the positioning's implications in advance.

1.2 RESEARCH QUESTIONS AND APPROACH

Under conditions of narrative dependency, systematic narrative analysis becomes the primary available means for understanding how technological readiness is constructed, how authority over that construction is established, and how public resources are mobilised in response. This thesis takes three documents from Microsoft's quantum computing programme as its empirical objects, asking what they reveal about the construction of technological credibility under conditions of limited access.

Czarniawska (2004) argues that narratives in organisational contexts are not secondary accounts of processes that occur elsewhere but primary sites where those processes are socially constructed and legitimised. This perspective transforms the epistemological constraint identified above from a methodological problem into an analytical opportunity: if corporate narratives constitute technological realities rather than merely reflecting them, then analysing those narratives offers a productive route to understanding quantum computing's sociotechnical development.

The central question this thesis addresses is:

Primary Research Question: How do corporate narratives construct and legitimise claims of technological readiness for an emerging technology whose development remains confined to inaccessible laboratories?

Whilst the primary research question is framed broadly, quantum computing is not treated here as an interchangeable instance of emerging technology in general. Several features make it analytically distinctive. The extreme technical specialisation required to evaluate quantum claims produces an epistemic distance between insiders and outsiders that exceeds most comparable fields: limited access compounds into limited comprehension. Research capacity is concentrated in a handful of corporate and state actors worldwide, meaning insider-produced narratives face little independent correction. Claimed applications — particularly the threat to existing encryption infrastructure — carry present societal stakes even before the technology matures. And entanglement with national security interests adds a layer of institutionalised secrecy that compounds corporate proprietary closure. These features do not merely instantiate narrative dependency; they intensify it in ways that make quantum computing a particularly revealing case for examining the epistemological and political conditions under which technological credibility is constructed.

The primary question is addressed through four subsidiary questions that structure the empirical and theoretical inquiry:

SQ1: How have Microsoft’s quantum computing narratives changed between 2014 and 2025, and what do these changes reveal about the dynamics of corporate vision-building in emerging technology contexts?

SQ2: What discursive strategies do these narratives employ to construct authority and demarcate technological approaches, and how do these strategies relate to one another?

SQ3: What sociotechnical imaginaries do these narratives construct, and how do those imaginaries change over time in relation to Microsoft’s strategic positioning?

SQ4: What does the reliance on corporate narratives — in the absence of direct laboratory observation — reveal about the epistemological conditions for studying technologies under access constraints, and what are the methodological implications for STS?

To address these questions, I analyse three documents produced by Microsoft across eleven years. Jennifer Warnick’s “Station Q” (2014) is an ethnographic portrait of Microsoft’s quantum computing laboratory, constructing what I term the “frontier imaginary” of quantum computing as scientific exploration. Matthias Troyer’s “Towards Practical Quantum Advantage” colloquium (2021) is a technical presentation navigating between quantum hype and practical possibility, employing the metaphor of a maze to distinguish productive research

How Quantum Computing Escapes the Maze

paths from dead ends. Microsoft's "Majorana 1 chip" announcement (2025) is a corporate communication claiming breakthrough achievement and near-term commercial deployment, repositioning quantum computing from speculative research to imminent technology.

These three documents mark three positions inside the maze: the initial construction of the laboratory as a frontier worth entering (2014), the systematic elimination of false paths (2021), and the announcement that the exit has been found (2025). Tracing that sequence diachronically — across eleven years and three distinct communicative genres — is what allows the thesis's central argument about narrative dependency, authority, and the limits of independent knowledge to become visible.

1.3 THESIS ROADMAP

The thesis proceeds through eight chapters following this introduction.

Chapter 2 establishes the theoretical and conceptual framework. It opens with the laboratory confinement problem and narrative dependency before surveying five bodies of STS scholarship — sociotechnical imaginaries, the sociology of expectations, responsible research and innovation, Hilgartner's vanguard framework, and the sociology of scientific knowledge — identifying their contributions and the gap this thesis addresses. It then assembles an integrated analytical framework organised around four dimensions: the constitution of facts through narrative practice; the dynamics of vanguard vision-building over time; the mechanisms of legitimation through boundary work and performativity; and the construction of technological futures. A closing section positions the thesis's contributions within this landscape.

Chapter 3 elaborates the methodological approach, explaining why document analysis follows analytically from conditions of narrative dependency, how the three primary texts were selected, how large language models were used as interpretive aids for technically complex material, and what the limitations of this approach are.

Chapters 4, 5, and 6 form the empirical core, each analysing one of the three primary texts. These chapters follow a consistent structure: examining the document's context and material form, critically assessing its issues following Asdal and Reinertsen (2022), conducting detailed narrative analysis, and providing STS theoretical interpretation.

Chapter 7 synthesises findings across the three texts, develops the thesis's theoretical contributions — particularly the concept of narrative dependency and the diachronic evolution of vanguard visions — revisits the controversy surrounding Microsoft's breakthrough

claims, and discusses implications for STS methodology when studying emerging technologies under conditions of limited access.

Chapter 8 concludes by summarising key findings, acknowledging limitations, and suggesting future research directions — particularly the need for ethnographic access to quantum computing laboratories as the technology transitions towards deployment.

Throughout, this thesis demonstrates that quantum computing's development cannot be understood through technical progress alone but requires attention to how narratives construct possibilities, manage expectations, establish authority, and legitimise technological pathways. In doing so, it contributes to STS scholarship both substantively — advancing understanding of quantum computing's sociotechnical development — and methodologically — demonstrating how to study emerging technologies when direct observation remains impossible.

Readers should attend, across all three empirical chapters, to how authority is constructed, how temporal framings are deployed, and how the boundary between legitimate science and hype is drawn — for these are the recurring mechanisms through which readiness is narratively constituted. As a navigational device, I appropriate Troyer's own maze metaphor — which I analyse in Chapter 5 — to structure the reader's journey. Troyer uses the maze to describe the landscape of quantum computing applications; I use it reflexively to describe the landscape of corporate narratives through which this thesis moves. The metaphor is not decorative: it is theoretically motivated. Just as external observers of quantum computing must navigate by insider-produced signposts because laboratory access is impossible, so the reader of this thesis must navigate by the signposts I provide — a structural enactment of the narrative dependency argument itself.

2 THEORETICAL AND CONCEPTUAL FRAMEWORK: SITUATING NARRATIVE DEPENDENCY IN STS SCHOLARSHIP

This chapter serves two related but distinct purposes. The first is to situate the concept of narrative dependency — introduced in Chapter 1 as this thesis’s central analytical contribution — within the landscape of existing STS scholarship, showing what prior research illuminates, where it reaches its limits, and what analytical work remains to be done. The second is to assemble the theoretical framework that the empirical chapters deploy: not a collection of theories applied in sequence, but a set of complementary resources that share a common constructivist commitment and address distinct dimensions of a single underlying process. The chapter proceeds in four parts. Chapter 2.1 establishes quantum computing as an analytically distinctive object for STS inquiry. Chapter 2.2 develops the laboratory confinement problem and introduces narrative dependency as the concept that names the epistemological condition this creates. Chapter 2.3 surveys STS scholarship across five intersecting fields, identifying what each contributes and where, collectively, a gap remains. Chapter 2.4 assembles the analytical framework and argues for its coherence. Section 2.5 positions this thesis’s contributions within this landscape.

2.1 QUANTUM COMPUTING AS AN ANALYTICAL OBJECT

Quantum computing represents an emerging technology that researchers suggest could provide computational advantages for specific problem classes. Preskill (2018) describes quantum computers with 50–100 qubits as potentially able to “perform tasks which surpass the capabilities of today’s classical digital computers” (p. 79), whilst acknowledging that noise limitations constrain near-term applications.

Proposed applications include cryptography and drug discovery through molecular simulation (Blunt et al., 2022; Cao et al., 2019). The cryptographic concern is particularly consequential: Shor’s (1994) algorithm demonstrates that a sufficiently powerful quantum computer could factorise large integers exponentially faster than any known classical method, which would render RSA encryption — the asymmetric cryptography underpinning most current internet security — vulnerable to decryption. This threat is not merely theoretical; it motivates present-day policy responses, including the post-quantum cryptography standards provided by the National Institute of Standards and Technology in 2024 (NIST, 2024), precisely because data encrypted today could be used as soon as decryption with quantum computers is available.

For STS scholarship, however, quantum computing's analytical interest lies less in what it promises than in the conditions under which those promises are produced, evaluated, and contested. As established in Chapter 1, quantum computing combines extreme technical specialisation, a highly oligopolistic research landscape, claimed civilisational stakes, and institutionalised secrecy in ways that make the relationship between insider knowledge and public representation unusually consequential. These are not incidental features but structurally determining conditions for how quantum computing is socially known. The technology's claimed significance is calibrated to its inaccessibility: the larger the promise, the more completely its evaluation depends on trusting the promisors. It is from this constitutive tension that the thesis's analytical questions arise.

2.2 THE LABORATORY CONFINEMENT PROBLEM AND NARRATIVE DEPENDENCY

The appropriate starting point for assembling a theoretical framework is not the frameworks themselves but the problem they must address. Narrative dependency — the condition under which external observers are compelled to rely on insider-produced accounts because no independent route to the object of study exists — does not map neatly onto a single existing theoretical framework. It names, rather, a condition that several traditions illuminate partially and from different angles, and that drawing on their combination makes analytically tractable. But narrative dependency is more than a methodological predicament for the analyst: it is a political condition for the technology's publics. When access to a laboratory is physically and epistemically foreclosed, insider narratives do not merely supplement independent understanding — they substitute for it. The power to narrate becomes, under these conditions, the power to constitute: to determine what counts as progress, what counts as a breakthrough, and what counts as hype. This concentration of interpretive authority in the hands of a technology's producers is not incidental to quantum computing's development but structurally inherent to it and analysing how that authority is constructed and exercised is one of the central tasks this thesis undertakes.

The foundational tradition for understanding this condition is laboratory studies. In *Laboratory Life*, Latour and Woolgar (1986/1979) established a social constructivist account of scientific knowledge production, demonstrating that facts are not discovered but manufactured: “the set of statements considered too costly to modify constitute what is referred to as reality” (p. 243). Central to their methodology was the capacity for direct ethnographic observation of laboratory practice — watching how scientists transform material phenomena into inscriptions, and inscriptions into facts. Their innovation lay precisely in looking inside

How Quantum Computing Escapes the Maze

knowledge construction rather than accepting science's own account of itself. The laboratory could be entered; the ethnographer could be there.

Quantum computing forecloses this founding move. The physical requirements of quantum computing — cryogenic isolation, specialised fabrication environments, and measurement apparatus comprehensible only to those trained in condensed matter physics — present barriers to ethnographic access qualitatively different from those encountered in most prior laboratory studies. Corporate proprietary interests and national security entanglements add further restrictions on what can be observed, by whom, and under what conditions. The published STS ethnographic literature on quantum computing laboratories remains, consequently, extremely sparse. What remains analytically accessible is the literary output of the laboratory: the inscriptions, the narratives, the communications addressed to audiences who are also outside. Following Asdal and Reinertsen (2022), these texts are not secondary evidence about processes that occurred elsewhere; they are the primary available site for understanding how quantum computing is being socially constituted.

Knorr Cetina's (1999) concept of epistemic cultures deepens this analysis. Her comparative ethnography of high-energy physics and molecular biology demonstrated that different scientific fields operate according to fundamentally distinct knowledge-making practices, differing not only in methodology and instrumentation but in their characteristic approaches to evidence, validation, and the relationship between theory and observation. Quantum computing, combining theoretical physics, experimental condensed matter physics, computer science, and engineering, may constitute a distinctive epistemic culture — one whose knowledge-making practices are opaque not only to physical non-entry but to conceptual comprehension from outside. The asymmetry between insiders and outsiders is therefore not simply a matter of proximity but of irreducibly specialised expertise that cannot be fully bridged by reading narratives about it.

Together, these considerations define what this thesis terms narrative dependency: the epistemological and political condition that arises when physical inaccessibility, technical specialisation, corporate control, and public stakes combine to make insider-produced narratives the primary — and in practice often the only — available resource through which external observers can form any understanding of a technology's development. This is the condition the thesis analyses, and it is the condition that makes the theoretical framework assembled in Section 2.4 analytically motivated rather than merely applied.

2.3 THE STATE OF STS RESEARCH ON QUANTUM COMPUTING

STS scholarship on quantum computing has grown substantially over the past decade, drawing on several distinct but intersecting theoretical traditions. This section surveys five bodies of work, identifying what each contributes and noting where, collectively, they leave a gap that this thesis addresses.

2.3.1 Sociotechnical Imaginaries and the Construction of Quantum Futures

A significant strand of STS quantum computing research draws on Jasanoff and Kim's (2009) concept of sociotechnical imaginaries, defined as "collectively imagined forms of social life and social order reflected in the design and fulfilment of nation-specific scientific and/or technological projects" (p. 120). This framework has proven generative for understanding how quantum futures are constructed and contested across academic, corporate, and policy contexts precisely because it directs attention not to technical claims about what quantum computing will do but to the social visions those claims carry with them.

Roberson (2021) applies imaginaries analysis to Australian Research Council quantum computing grants from 2002 to 2020, demonstrating how expectations function as "powerful cultural resources" that shape research priorities and funding decisions. Nguyen (2025) extends this to policy, arguing that quantum technology export controls are shaped by the sociotechnical imaginary of quantum rather than its material reality: "it is not the material reality of quantum innovation — which is still mired in major engineering challenges — that informs export control efforts surrounding [quantum technology], but rather the 'sociotechnical imaginary' of quantum that serves as the 'muse' for law- and policy-makers" (p. 68). These analyses demonstrate that quantum imaginaries are productive forces shaping consequential decisions, not post-hoc reflections of technical realities.

Selin's (2008) sociology of the future extends imaginaries theory by analysing how futures are made present through narrative, how temporal constructions shape what appears inevitable versus contingent, and how different actors compete to establish their visions as authoritative. This temporal dimension is essential for understanding how corporate communications manage the relationship between current capabilities and promised futures: the assertion that quantum advantage is "years, not decades" away performs imminence whilst remaining deliberately imprecise enough to resist falsification.

2.3.2 Science and Technology Policy and the Sociology of Expectations

A second body of scholarship approaches quantum computing through frameworks developed in science and technology policy (STP) and innovation studies. Where imaginaries theory operates at the level of collective vision, this tradition addresses the mechanisms through which promises harden into requirements and shape technology trajectories at the systemic level.

Van Lente and Rip (1998a, 1998b) demonstrate that in emerging technology fields, promises are not merely expressions of optimism but performative statements that create obligations: once made, they must be answered, whether through fulfilment or through the production of new promises that account for the original's non-delivery. Borup et al.'s (2006) sociology of expectations shows that expectations are productive forces: they attract resources, coordinate activity, and constitute the social space within which innovation occurs. Brown and Michael (2003) add temporal complexity, showing how failed timelines are managed retrospectively — through what they term “the retrospective prospectation” — without abandoning the underlying narrative of progress. These frameworks provide direct resources for understanding the pattern observable in Microsoft's quantum narrative across eleven years. They illuminate the systemic context within which Microsoft's narrative work takes place: the company is not merely communicating about a technology but actively attempting to manage the conditions under which quantum computing transitions from research to commercial deployment.

Empirically, studies document the expectation dynamics specific to quantum computing. Deshpande (2022) highlights concerns about “potentially unrealistic expectations and timelines” (p.57) and the difficulty of distinguishing “the hype from market realities” (p.57). O'Shea (2022) notes the potential for speculative bubbles to arise from misinformation and unrealistic expectations, which may adversely affect long-term investment strategies.

Holter et al. (2024) find that whilst half of surveyed technologists, researchers, and policymakers acknowledged the value of generating societal enthusiasm, 84% believed that claims about quantum computing are frequently exaggerated in public discourse. These findings confirm that expectation management is not merely an academic category but a contested practical problem within the quantum computing community itself.

2.3.3 Responsible Research and Innovation

Responsible Research and Innovation (RRI) has emerged as a significant governance framework for emerging technologies, developed through the work of R. Owen et al. (2012) and institutionalised in European research policy through von Schomberg's (2013) formulation:

Responsible Research and Innovation is a transparent, interactive process by which societal actors and innovators become mutually responsive to one another with a view to the (ethical) acceptability, sustainability, and societal desirability of the innovation process and its marketable products (in order to allow a proper embedding of scientific and technological advances in our society). (René von Schomberg, 2013)

Applied to quantum computing by Inglesant et al. (2021), RRI analysis has begun to examine the visions and imaginaries that inform the technology's development and their implications for governance.

A more quantum-specific field — Quantum-ELSPI, addressing ethical implications, legal considerations, social impacts, and policy implications — has emerged as a distinct area of inquiry (Kop, 2023; Kop et al., 2024). They have developed governance principles organised around safeguarding, engaging, and advancing; anticipatory governance approaches have been proposed (Jong, 2022); quantum governance stack models articulated (Perrier, 2022); and constitutionalist perspectives developed (Wimmer & Moraes, 2022). This body of work reflects a growing recognition that quantum computing raises governance questions that cannot be addressed through technical standards alone.

Corporate actors have selectively adopted the language of responsible innovation. IBM's claim that "the era of quantum utility must also be the era of responsible quantum computing" (Wolf-Bauwens & Mandelbaum, 2024) illustrates how RRI frameworks migrate from academic governance discourse into corporate positioning. Microsoft's Quantum Safe programme and workforce development initiatives (Microsoft, 2025) operate within a similar register. For the purposes of this thesis, the RRI literature is analytically relevant less as a normative framework to be applied than as evidence of the contested field within which corporate quantum narratives operate: the very existence of RRI discourse around quantum computing is itself a sociotechnical achievement, produced through the work of scholars, institutions, and policymakers who are themselves responding to corporate narratives about the technology's imminence.

2.3.4 Sociotechnical Vanguard and Vision-Building

Hilgartner's (2015) analysis of sociotechnical vanguards provides a theoretical resource that bridges the imaginaries and expectations traditions whilst adding a dimension neither fully addresses: the organisational and rhetorical work through which specific actors position themselves as authoritative guides to technological futures that have not yet materialised. A vanguard, in Hilgartner's account, is an organisation that actively promotes a vision of a technological future, constructing the social conditions under which that future appears both desirable and inevitable. Vanguard work is not merely promotional in the colloquial sense; it is a form of world-building that establishes the cognitive and institutional scaffolding within which a technology is understood, funded, and governed.

Hilgartner identifies five characteristic operations through which vanguards conduct this work. The first is the construction of cultural identity: vanguards cultivate a distinctive image of their researchers, their institutional culture, and their relationship to prior technological revolutions, positioning themselves as the rightful inheritors of a scientific lineage and the legitimate custodians of a technological future. The second is the establishment of historical legitimacy: by narrating their work as the continuation of an established scientific tradition — connecting, for instance, quantum computing to the foundational achievements of quantum mechanics — vanguards claim a pedigree that pre-empts questions about credibility. The third is the creation of imaginary landscapes: vanguards map the terrain of a technology's possible applications, identifying where progress is possible and where it is not, which paths are worth pursuing and which are dead ends. This mapping is not neutral cartography but strategic positioning, defining the problem space in ways that favour the vanguard's own approach. The fourth is the temporal management of expectations: vanguards deploy promises, timelines, and milestones in carefully calibrated ways, maintaining sufficient credibility to attract investment and talent without committing to falsifiable predictions whose failure would undermine the vision. The fifth is the demonstration of capability: vanguards periodically produce material evidence — prototypes, experimental results, technical demonstrations — that serves to instantiate the imaginary, making the abstract vision tangible enough to sustain continued commitment from audiences who cannot directly observe laboratory progress.

What makes Hilgartner's framework particularly valuable for this thesis is its attention to the relationship between the vanguard's communicative work and the social conditions under which that work is received. Vanguard vision-building is not conducted in a vacuum but in a competitive landscape where multiple actors advance competing visions, where audiences bring their own interests and scepticisms, and where the credibility of any single vision

depends partly on the discrediting of alternatives. Microsoft operates in precisely such a landscape: its topological approach to quantum computing has been developed in explicit contrast to the superconducting architectures pursued by IBM and Google, and its communications consistently position this distinction as the basis for a decisive long-term advantage. Understanding these communications as vanguard work — as strategic performances addressed to investors, policymakers, potential recruits, and the scientific community simultaneously — is what allows the analysis to move beyond asking whether Microsoft's claims are true towards asking how they are made credible, to whom, through what operations, and at what cost to alternative framings.

Hilgartner's framework was developed primarily through synchronic analysis of vision-building at particular moments. The present thesis extends it diachronically, tracing how Microsoft's vanguard strategies shift across three distinct communicative stages spanning eleven years. This extension is analytically motivated by the sociology of expectations (Borup et al., 2006; van Lente & Rip, 1998a, 1998b): as promises age and encounter material constraints, the vanguard's communicative repertoire must adapt, and tracing that adaptation reveals dynamics that no single snapshot analysis could capture.

2.3.5 The Sociology of Scientific Knowledge and Credibility

The sociology of scientific knowledge (SSK) developed from the 1970s onwards as a systematic programme for treating scientific knowledge as any other form of belief. Its central claim — that the content of scientific knowledge, not merely its institutional organisation, is shaped by social factors — is analytically productive for understanding how credibility is constructed and contested in settings far beyond the laboratory. For this thesis, SSK provides the theoretical grounding for one of its most basic analytical premises: that Microsoft's claims about quantum computing are not self-validating, that their credibility must be produced through specific social operations, and that those operations are tractable to analysis.

Gieryn's (1983) concept of boundary work is the most directly applicable resource. Scientists, Gieryn argues, do not simply conduct research; they actively demarcate their activity from what lies beyond it — from non-science, pseudo-science, and competing forms of expertise — in ways that serve institutional interests and protect epistemic authority. This demarcation is strategic and context-dependent rather than the application of a fixed criterion. In quantum computing, boundary work operates simultaneously at several levels: between rigorous scientific assessment and hype, between viable and non-viable technological approaches, and between Microsoft's topological architecture and the competing superconducting

How Quantum Computing Escapes the Maze

approaches pursued by IBM and Google. What counts as serious science, credible progress, or premature promise is not given by the technical facts alone but is actively constructed through communicative work of precisely the kind this thesis analyses.

Collins's (1992/1985) work on expertise and the experimenter's regress adds a further dimension. In domains where experimental competence is concentrated in a small specialist community, outsiders cannot independently resolve disputes between competing expert claims: they lack the tacit knowledge required to evaluate experimental technique, and the very criteria by which a competent experiment would be recognised are themselves part of what is contested. This structural condition — which Collins terms the experimenter's regress — means that external audiences must navigate between appeals to authority whose credentials they cannot directly assess. It maps directly onto the situation of audiences for Microsoft's quantum communications: investors, policymakers, and the general public cannot evaluate the underlying experimental results and must instead rely on markers of credibility — institutional affiliations, publication records, peer recognition, rhetorical authority — that are themselves products of the social organisation of the field. Understanding how Microsoft's documents construct and deploy such markers is one of the analytical tasks the empirical chapters undertake.

2.3.6 What Existing STS Research Leaves Unaddressed

Taken together, these five bodies of scholarship illuminate important dimensions of quantum computing's social development. Imaginaries research reveals how quantum futures are constructed and contested. Innovation studies and the sociology of expectations explain how promises function as productive forces and how expectation management shapes technology trajectories. RRI scholarship identifies the governance challenges quantum computing raises. Science communication research illuminates the politics of how publics are addressed. SSK demonstrates that credibility is socially produced and analytically tractable.

What this scholarship does not fully address is the epistemological situation this thesis terms narrative dependency: the condition under which the usual resources for independent evaluation are structurally unavailable rather than merely practically difficult. Laboratory studies provide the most directly relevant tradition — developed precisely to study what happens inside laboratories rather than accepting science's own account of itself — but laboratory studies presuppose access. When that access is foreclosed by physical requirements, corporate control, and security classification, the tradition's founding methodological premise cannot be satisfied. The existing STS literature on quantum

computing has not yet systematically addressed this condition, nor proposed analytical resources for treating it as something other than a methodological limitation to be managed. Narrative dependency proposes that it should be treated, instead, as the central analytical object.

2.4 AN ANALYTICAL FRAMEWORK FOR NARRATIVE-DEPENDENT TECHNOLOGIES

The theoretical resources assembled above are diverse in origin but share a common analytical commitment: they treat scientific facts, technological futures, credibility, authority, and public understanding not as natural givens but as achievements — produced through social practice, contingent on institutional arrangements, and subject to contestation. This shared constructivist ontology is what makes their combination analytically coherent rather than merely eclectic. The frameworks are compatible not because they address the same object but because they address different dimensions of a single process: the social construction of technological credibility under conditions of limited access.

That process can be decomposed into four analytically distinct but practically intertwined dimensions. The framework is organised around these dimensions rather than around the research questions, because the same theoretical resource often bears on more than one question and because organising by question would obscure the relational character of the framework — the way in which each dimension presupposes and shapes the others.

2.4.1 The Constitution of Facts: Narrative as Practice

The primary research question asks how corporate narratives construct and legitimise claims of technological readiness. Answering it requires a theory of what narratives do: not what they say but how they operate as social practices that bring realities into being rather than merely reflecting them.

Czarniawska (2004) provides the foundational resource, arguing that narratives in organisational contexts are not secondary accounts of processes that occur elsewhere but primary sites where those processes are socially constructed and legitimised. This transforms documents from transparent windows onto reality into actors that perform work in the world — work that includes enrolling audiences, stabilising claims, and foreclosing alternatives. Asdal and Reinertsen (2022) develop a practice-oriented document analysis methodology on similar premises: documents are things that do things, and analysing them requires attending

How Quantum Computing Escapes the Maze

to how they position their producers, address their audiences, and make certain realities more or less available for contestation.

Latour and Woolgar's (1986/1979) concepts of literary inscription and modalization extend this to scientific knowledge production specifically. Modalization — the process through which claims move along a continuum from speculation to established fact, or back again — is directly observable in Microsoft's quantum communications: early presentations hedge claims with conditionals that later announcements retrospectively treat as having already been resolved. How this modalization is accomplished through specific rhetorical operations is one of the central questions the empirical analysis addresses.

2.4.2 The Dynamics of Vision-Building: Vanguard and Expectation Management

SQ1 asks how Microsoft's quantum narrative has changed across eleven years and what this reveals about corporate vision-building. Answering it requires resources capable of explaining diachronic change in how technological futures are constructed and managed.

Hilgartner's (2015) vanguard framework provides the primary analytical resource. Vanguard constructs cultural identity, establish historical legitimacy, create imaginary landscapes, manage expectations temporally, and demonstrate capability. This framework was developed to analyse vision-building at a given moment; the present thesis applies it diachronically, tracing how vanguard strategies shift across three distinct communicative stages as the promises of earlier moments encounter material and discursive constraints.

The sociology of expectations provides the theoretical account of why such shifts occur. Van Lente and Rip's (1998a, 1998b) promise-requirement cycle shows how initial promises create obligations that must be answered; Borup et al.'s (2006) analysis of how expectations attract resources and coordinate activity explains why vanguards cannot simply abandon commitments when they become inconvenient; Brown and Michael's (2003) retrospective prospecting illuminates how failed timelines are reframed as necessary steps on a still-credible path. Together, these frameworks account for the pattern observable in Microsoft's quantum narrative: the frontier confidence of 2014, the expectation-contracting pragmatism of 2021, and the re-expansive breakthrough announcement of 2025.

2.4.3 The Mechanisms of Legitimation: Boundary Work and Performativity

SQ2 asks what discursive strategies these narratives deploy to construct authority and demarcate technological approaches. Two complementary theoretical resources address the mechanisms through which legitimation is accomplished.

Gieryn's (1983) boundary work and Collins's (1992/1985) analysis of expertise address how distinctions between legitimate and illegitimate knowledge are drawn and maintained. In Microsoft's quantum communications, boundary work is visible at multiple levels simultaneously: Troyer's 2021 colloquium systematically demarcates viable from non-viable quantum applications, positions Microsoft's topological approach against competing architectures, and draws a distinction between rigorous scientific assessment and "hype." This work is not neutral description but strategic positioning that simultaneously advances Microsoft's approach and delegitimises alternatives.

Austin's (1962) speech act theory, extended to science and markets by Callon (1998) and MacKenzie (2006), provides the complementary resource. Performative utterances do not describe pre-existing realities but bring new realities into being through their declaration. When Microsoft announces a breakthrough, the performative effects — on investor expectations, on competitor positioning, on policy frameworks, on the social status of the underlying research — occur regardless of whether independent validation follows. MacKenzie's demonstration that financial models function as "engines, not cameras" is directly transferable to the quantum computing case: Microsoft's narratives do not reflect the state of quantum computing but actively constitute it, shaping what the technology is understood to be and what it is expected to become.

These mechanisms are not independent. Boundary work establishes the authority from which performative declarations derive their credibility; performative declarations, in turn, reshape the boundaries of what counts as legitimate scientific assessment. Analysing them together, as integrated legitimation strategies, illuminates how corporate authority is constructed and maintained across the three documents in ways that examining either mechanism in isolation would not.

2.4.4 The Construction of Futures: Imaginaries and Temporal Framing

SQ3 asks what futures Microsoft's narratives construct for quantum computing and how these change in relation to the company's strategic positioning. The relevant resources are those

How Quantum Computing Escapes the Maze

concerned with how futures are collectively imagined, institutionally stabilised, and contested.

Jasanoff and Kim's (2009, 2015) sociotechnical imaginaries framework is the primary resource here. Imaginaries organise expectations, mobilise resources, justify investment, establish legitimacy for technological pathways, and construct relationships between technology and society. The diachronic analysis this thesis undertakes reveals how Microsoft's quantum imaginary expands, contracts, and shifts across the three documents — changes that correspond to strategic imperatives and competitive pressures as well as, and sometimes rather than, changes in technical capability.

Selin's (2008) sociology of the future contributes the temporal dimension: how different futures are made present through narrative, how some trajectories are naturalised as inevitable whilst others are positioned as contingent, and how competing actors vie to establish their temporal framings as authoritative. This is directly relevant to understanding how Microsoft's communications manage the relationship between current limitations and projected futures — a relationship that is never simply reported but always strategically constructed. Knorr Cetina's (1995) attention to metaphors as epistemic resources — not merely rhetorical devices but active tools through which scientists navigate between disparate domains — adds a further dimension, particularly relevant to the analysis of Troyer's maze metaphor and its reappropriation by this thesis.

2.4.5 Framework Integration: A Shared Constructivist Commitment

The four dimensions outlined above address a single underlying process — the construction of technological credibility under conditions of narrative dependency — from different analytical angles. The constitution of facts (chapter 2.4.1) describes the micro-level operations through which claims are stabilised in text; the dynamics of vision-building (chapter 2.4.2) explains why those operations take the strategic forms they do over time; the mechanisms of legitimation (chapter 2.4.3) identify the discursive resources through which authority is established and contested; and the construction of futures (chapter 2.4.4) situates those mechanisms within the larger social project of imagining and institutionally anchoring desirable technological trajectories.

What holds these dimensions together is a shared constructivist commitment: the premise that facts, futures, authority, and credibility are achievements accomplished through practice rather than properties of the phenomena themselves. Czarniawska and Latour and Woolgar establish that narratives constitute rather than describe; Hilgartner and the sociology of

expectations establish that technological visions are strategically produced and managed over time; Gieryn, Collins, Austin, Callon, and MacKenzie establish that credibility and legitimacy are socially accomplished through specific discursive operations; and Jasanoff and Kim with Selin establish that the futures projected by those operations are politically consequential constructions, not neutral extrapolations from technical facts.

This shared commitment is not imposed on the analysis from outside but is the theoretical consequence of taking narrative dependency seriously as an analytical condition. If the objects of study — Microsoft's claims about quantum computing's readiness — cannot be evaluated against independent evidence by those outside the laboratory, then what remains tractable is precisely how those claims are constructed, by whom, through what mechanisms, and with what effects. A framework oriented towards construction rather than verification is not a methodological compromise; it is the analytically appropriate response to the epistemological situation the thesis confronts.

2.5 THIS THESIS'S CONTRIBUTION

By analysing three documents from Microsoft's quantum computing programme across eleven years, this thesis seeks to make four interconnected contributions to STS scholarship. These are developed and substantiated in Chapter 7; they are signalled here so that readers can attend to their emergence across the empirical chapters.

First, and most centrally, the thesis develops narrative dependency as an epistemological concept. It proposes that this concept names a distinctive condition — not reducible to general claims about the social construction of scientific knowledge — that arises at the intersection of physical inaccessibility, technical specialisation, corporate control, and public stakes. The concept draws on Czarniawska (2004) and Knorr Cetina (1999) but extends their frameworks to address the specific situation in which ethnographic access is structurally foreclosed rather than merely practically difficult. Whether this conceptual proposal proves useful beyond the quantum computing case is a question that must await future research; within this thesis, it functions as the organising analytical premise whose usefulness is demonstrated through application.

Second, the thesis applies Hilgartner's (2015) vanguard framework to a diachronic analysis spanning three distinct communicative moments. Most applications of the framework are synchronic; this analysis suggests that the strategies operative at one developmental moment — frontier confidence, expectation management, breakthrough announcement — differ

How Quantum Computing Escapes the Maze

substantially from those appropriate at another, and that understanding this variation requires attending to the interaction between strategic positioning and the promise-requirement dynamics identified by van Lente and Rip (1998a, 1998b) and Borup et al. (2006). The analysis is constrained to three documents from one company; the generalisability of this finding across other vanguard organisations and technology domains remains to be tested.

Third, the thesis treats boundary work, expectation management, and performativity not as isolated analytical mechanisms but as integrated legitimation strategies that operate together within corporate science communication. The empirical analysis of these three documents suggests that these mechanisms reinforce one another in ways that examining any single mechanism in isolation cannot reveal. This is a reading offered on the basis of a focused case study, not a general claim about corporate science communication.

Fourth, the thesis traces how quantum computing's projected futures change across the three documents in relation to Microsoft's strategic positioning. The analysis contributes to imaginaries research by attending to the strategic dimension of imaginary construction — the way in which imaginary expansion and contraction respond to competitive, financial, and reputational pressures alongside, and sometimes instead of, technical developments.

Methodologically, the thesis demonstrates one approach to conducting systematic document analysis under conditions of narrative dependency, including the productive use of large language models as interpretive aids for technically complex materials and the analytical possibilities opened by an insider-outsider positioning that is reflexively acknowledged rather than either concealed or treated as disqualifying.

Readers approaching the empirical chapters should carry from this framework a single overarching question: through what operations — narrative, rhetorical, performative, imaginary — does Microsoft establish the credibility of claims that cannot, by virtue of the conditions described in chapter 2.2 about narrative dependency, be independently verified? The chapters that follow show those operations at work across three documents, three moments, and eleven years of quantum computing's public life.

3 METHODOLOGY: ANALYSING NARRATIVES UNDER CONDITIONS OF LIMITED ACCESS

Before entering the empirical analysis, one further question must be addressed: not what this thesis is looking for, but how it can look at all. This chapter describes the analytical method — how one reads the walls of a maze one cannot physically enter. It explains why document analysis follows analytically from conditions of narrative dependency; how the corpus of Microsoft’s quantum communications was surveyed and three documents selected from it; why those three documents are sufficient for the claims the thesis advances; how large language models were employed as interpretive aids for technically complex material; and what the limitations of this approach are. Readers already familiar with document analysis methodology may focus particularly on chapters 3.2 and 3.3, which address the specific selection logic and its justification.

3.1 DOCUMENT ANALYSIS AS ANALYTICAL CONSEQUENCE OF NARRATIVE DEPENDENCY

The choice of document analysis as the primary method for this thesis is not a preference but an analytical consequence of the epistemological situation established in chapter 2. Under conditions of narrative dependency — where physical inaccessibility, technical specialisation, corporate control, and public stakes combine to foreclose independent observation — documents are not one source of evidence among others. They are the primary available site at which quantum computing’s development is publicly constituted.

Czarniawska (2004) argues that narratives in organisational contexts are not secondary accounts of processes that occur elsewhere but primary sites where those processes are socially constructed and legitimised. If corporate narratives constitute technological realities rather than merely reflecting them, then systematic analysis of those narratives is not a methodological compromise in the absence of laboratory access but a legitimate and analytically appropriate method for understanding how technological readiness is constructed. The epistemological constraint becomes, in this framing, an analytical opportunity: the documents are not a pale substitute for what could be observed inside the laboratory but the principal medium through which quantum computing’s development is made publicly available at all.

Asdal and Reinertsen (2022) develop a practice-oriented document analysis methodology that aligns well with these premises. They argue that documents should be treated not as transparent windows onto reality but as actors that do things in the world — that position their

How Quantum Computing Escapes the Maze

producers, address their audiences, make certain realities available and others invisible, and mobilise resources and commitments in response. Their methodology directs the analyst to attend to what documents make *present* and what they make *absent*: what claims are advanced, what evidence is offered, what qualifications are suppressed, what controversies are acknowledged or foreclosed. Applied to corporate quantum communications, this methodology transforms the narrative dependency condition from a limitation to be managed into the central analytical object.

This is not to suggest that document analysis is without limitation. It cannot access laboratory practices, internal decision-making processes, or the tacit knowledge through which scientists evaluate one another's claims. These limitations are genuine and are addressed in chapter 3.4. What document analysis can do — and what this thesis does — is reconstruct the public performative work through which quantum computing's credibility is constructed, managed, and occasionally contested, in the texts that circulate between the laboratory and its audiences.

3.2 CASE SELECTION: THREE TEXTS, ELEVEN YEARS

3.2.1 Microsoft's Quantum Communication Landscape

Microsoft's quantum computing programme operates across a wide and heterogeneous range of communication channels, and establishing what was available before explaining what was selected is essential for understanding the selection logic. The scale of Microsoft's public quantum communication is substantial. Between 2006 and 2025, the company contributed to more than 1,200 peer-reviewed scientific articles published in formal journals including *Nature*, *Nature Physics*, *IEEE Computer*, *Quantum*, and the *Proceedings of the National Academy of Sciences*. In the same period, 82 doctoral theses were produced within Microsoft's quantum computing laboratories, constituting a significant body of formal scientific knowledge production closely associated with the company's research programme. Alongside these formal scientific channels, Microsoft maintains an extensive informal public communication infrastructure. Several YouTube channels associated with Microsoft's quantum research contain hundreds of videos, ranging from introductory explainers to full-length academic presentations and researcher interviews. The Microsoft Source platform — the company's dedicated storytelling channel for in-depth features — publishes long-form journalism, ethnographic portraits of research laboratories, and narrative accounts of scientific developments addressed to educated general audiences. Research blogs and

official website materials provide further communication directed at investors, policymakers, potential recruits, and the technology community more broadly.

These informal channels offer a distinct and analytically valuable perspective that formal scientific publications do not. Peer-reviewed articles are subject to disciplinary conventions that constrain the expression of strategic vision, suppress rhetorical elaboration, and require claims to be presented in formats subject to expert evaluation. Informal channels — particularly the long-form features, and public presentations selected for this thesis — allow researchers to express ideas, visions, and assessments of the field with a flexibility and directness that the formal scientific literature forecloses. It is in these channels that the construction of sociotechnical imaginaries, the management of public expectations, and the performative work of vanguard vision-building are most directly visible. This is not a methodological concession but a theoretical justification: the documents best suited to analysing how technological readiness is narratively constructed are those in which narrative construction is least constrained.

3.2.2 Selection Logic and Search Process

From this broad landscape, the selection of three documents for sustained analysis followed a deliberate search logic organised around analytically critical moments in Microsoft's quantum computing programme rather than a sampling procedure aimed at representativeness. The goal was not to produce a comprehensive account of Microsoft's communications but to identify documents that would enable examination of how narrative strategies shift across distinct developmental stages — and that would therefore address the thesis's research questions with sufficient analytical purchase.

The first search focus was the *beginning* of Microsoft's public quantum computing programme: the moment at which the company first established a public presence for its quantum research and began the work of constructing an audience for a technology that was, at that stage, highly speculative. This search led to Jennifer Warnick's 2014 ethnographic portrait of Station Q, published on the Microsoft Source platform — the most extensive publicly available account of Microsoft's quantum laboratory at a formative moment when the company's topological approach remained deeply uncertain.

The second focus was a *moment of crisis*: the period following Microsoft's retraction, in 2021 (Zhang et al., 2021), of a 2018 *Nature* paper (Zhang et al., 2018) that had claimed to present experimental evidence for Majorana fermions, the quasiparticle on which Microsoft's entire topological quantum computing strategy depends. This retraction was a significant and public

How Quantum Computing Escapes the Maze

credibility event, attracting critical commentary within the quantum physics community and raising substantive questions about the robustness of Microsoft's approach. The crisis created an analytically productive moment: it made the work of expectation management visible in ways that a stable period would not. This search led to Matthias Troyer's 2021 Berkeley colloquium, delivered in the same year as the retraction and constituting, on analysis, a carefully calibrated communicative response to the credibility pressure the retraction had created. The colloquium's academic setting, technical sophistication, and explicit engagement with the "hype problem" in quantum computing made it an analytically rich document for examining how authority is reconstructed and expectations managed at a moment of institutional stress.

The third document was not sought but encountered. The Majorana 1 chip announcement, published by Microsoft in February 2025, appeared during the writing of this thesis and was recognised immediately as analytically indispensable. It represented a third and qualitatively distinct communicative configuration: the declaration of breakthrough, repositioning quantum computing from speculative research to claimed near-term deployment. Its appearance during the writing period is an instance of methodological responsiveness rather than opportunism: precisely because quantum computing's public life unfolds through documents and announcements, a document that arrives during the analysis is not an interruption but an extension of the empirical field. The methodological literature on qualitative research supports this kind of theoretical sampling, in which the analyst adjusts the corpus in response to emergent analytical significance (Flick, 2014). Including the Majorana 1 announcement allowed the thesis to trace the full arc of Microsoft's quantum narrative across three analytically distinct stages — frontier establishment, crisis management, and breakthrough declaration — in a way that the original two-document design could not have achieved.

3.2.3 Why Three Documents Are Sufficient

The claim that three documents over eleven years are sufficient to support this thesis's analytical arguments requires explicit justification. Two objections are foreseeable. The first is that the corpus is too small to support generalisable claims about Microsoft's quantum narrative. The second is that it is too narrow to support claims about quantum computing more broadly. Both objections rest on a misunderstanding of what the thesis claims.

The thesis does not claim to comprehensively describe Microsoft's quantum communications, nor to generalise about quantum computing's social development across

all actors and contexts. Its claims are of a different kind: it seeks to *demonstrate how* narrative dependency operates as an epistemological and political condition; to *show how* vanguard vision-building strategies shift across distinct developmental stages; and to *analyse how* legitimisation mechanisms — boundary work, expectation management, performativity — operate together in corporate science communication. These are claims about processes and mechanisms, demonstrated through a focused case study rather than established through representative sampling. For claims of this kind, analytical depth and contrast are more productive than breadth, and three documents selected for maximum analytical contrast are more valuable than a larger corpus of more similar texts.

The three documents provide precisely this contrast. They represent three analytically distinct communicative configurations: the frontier imaginary of 2014, in which speculative possibility is constructed as scientific vocation; the expectation-contracting pragmatism of 2021, in which credibility is rebuilt through systematic demarcation of the viable from the non-viable; and the breakthrough declaration of 2025, in which the imaginary is re-expanded and readiness is performatively asserted. Each configuration activates different combinations of the thesis's theoretical resources and makes visible different aspects of the vanguard's communicative repertoire. Together, they enable a diachronic analysis that no single document or set of similar documents could support.

This approach follows established precedent in STS scholarship. Jasanoff's (2005) comparative work on biotechnology governance across the United States, the United Kingdom, and Germany draws on a carefully selected set of key documents and institutional moments rather than a comprehensive survey; Hilgartner's (2015) vanguard analysis focuses on a small number of organisations and communications precisely because analytical depth is preferable to empirical coverage when the goal is to understand how processes work. The principle of theoretical saturation — well established in qualitative research (Flick, 2014) — holds that additional data should be collected until new categories cease to emerge, not until all available data has been processed. The three documents selected for this thesis collectively activate the full range of the analytical framework developed in chapter 2; additional documents would add width without generating new analytical categories.

The limitations that do follow from this selection are genuine and are addressed in chapter 3.4. The analysis cannot speak to what Microsoft communicates through other channels, what other quantum computing actors do with their narratives, or how the relationship between narrative and material development evolves beyond the 2025 endpoint. These are genuine

constraints; they are the appropriate price of the analytical depth the selected documents make possible.

3.2.4 The Three Documents

The following table summarises the three primary texts and their key characteristics. Fuller analysis of each document’s context, material form, and institutional positioning opens the corresponding empirical chapter.

Table 1: three primary texts and their key characteristics.

	Warnick (2014)¹	Troyer (2021)²	Bolgar (2025)³
Document type	Long-form ethnographic journalism, Microsoft Source platform	Academic colloquium, UC Berkeley Simons Institute (YouTube)	Corporate news announcement, Microsoft Source platform
Length / scope	c. 6,000 words; five narrative sections; researcher photographs	65-minute presentation; 36 slides; technical content; Q&A session	c. 3,500 words; three sections; chip photographs; researcher portraits
Authors / speakers	Jennifer Warnick (Microsoft writer); features Freedman, Das Sarma, and Station Q researchers	Matthias Troyer (ETH Zürich / Microsoft Distinguished Scientist)	Catherine Bolgar (Microsoft writer); features Nayak and Svore
Primary audience	Educated general public; potential recruits; investors	Academic researchers; quantum computing specialists; graduate students	Technology journalists; investors; policymakers; general public
Developmental stage	Frontier exploration: constructing quantum research as scientific vocation	Expectation management: navigating the quantum maze after the Nature retraction	Breakthrough declaration: asserting scalable quantum computing within reach
Availability	Publicly available, Microsoft Source (news.microsoft.com)	Publicly available, YouTube (Simons Institute channel)	Publicly available, Microsoft Source (news.microsoft.com)

¹ Warnick, J. (2014). *Station Q*. Microsoft Corporation. <https://news.microsoft.com/stories/stationq/>

² Troyer, M. (2021, March 16). *Towards Practical Quantum Advantage | Quantum Colloquium*. Simons Institute, UC Berkeley. <https://www.youtube.com/watch?v=WY3htdKUGsA>

³ Bolgar, C. (2025). *Microsoft’s Majorana 1 chip carves new path for quantum computing - Source*. Microsoft Corporation. <https://news.microsoft.com/source/features/innovation/microsofts-majorana-1-chip-carves-new-path-for-quantum-computing/>

3.3 ANALYTICAL TOOLS AND TECHNIQUES

3.3.1 Primary Analytical Framework

Drawing on Asdal and Reinertsen's (2022) document analysis methodology, I develop a five-step analytical approach adapted to the specific demands of this research. The five steps are applied consistently across all three empirical chapters, providing structural comparability whilst allowing each document's distinctive features to shape the analysis.

Step 1:
Context and
Material
Form

Each analysis opens by examining the document's publication context, institutional positioning, target audience, material presentation — layout, visual elements, typography, genre conventions — and circulation pathways. This establishes how documents position themselves before their content is addressed and attends to the material dimensions of narrative construction that purely textual analysis would miss.

Step 2:
Critical
Assessment
of Issues

Following Asdal and Reinertsen's (2022) emphasis on what documents make present and absent, I systematically examine what claims are advanced and what evidence supports them; what information, caveats, or qualifications are omitted; what controversies or alternative perspectives are acknowledged or foreclosed; what temporal framings are employed; and how authority and expertise are constructed and attributed.

Step 3:
Narrative
Analysis

Following Czarniawska (2004), I examine the story each document tells, attending to narrative structure — protagonists, antagonists, challenges, resolutions — metaphorical frameworks, temporal construction, character construction, and plot development. Particular attention is paid to how the relationship between current limitations and projected futures is narratively managed

Step 4:
Theoretical
Inter-
pretation

I apply the integrated analytical framework developed in chapter 2, examining how each document constitutes facts through narrative practice; how it conducts vanguard vision-building; what legitimation mechanisms it deploys and how they operate together; and what sociotechnical imaginaries it constructs and how these relate to the company's strategic positioning.

Step 5: Brief Synthesis

Each empirical chapter concludes with a synthesis that connects the specific document analysis to broader patterns in Microsoft's quantum narrative and to the thesis's central argument about narrative dependency, diachronic vanguard evolution, and the construction of technological credibility under conditions of limited access.

3.3.2 Use of Large Language Models as Interpretive Aids

A distinctive feature of this thesis's methodology is the productive use of large language models (LLMs) as analytical tools. This requires careful justification and explication. Troyer's 2021 colloquium presented significant comprehension challenges to me. As Peter Lee, Microsoft's research director, characterises such presentations, they are "notoriously complex that after a few minutes into one even he starts to feel a bit thick and small" (Warnick, 2014). As an STS scholar without extensive background in quantum physics, chemistry, or advanced computer science, understanding Troyer's technical content required assistance. Rather than limiting STS analysis only to topics where researchers possess deep disciplinary expertise, productive LLM use may enable STS scholars to engage with broader range of scientific developments whilst maintaining analytical rigour.

I employed two large language models—Microsoft 365 Copilot and Anthropic's Claude 3.5 Sonnet—to facilitate comprehension of technical concepts. I instructed these LLMs with explicit guidelines. You find a transcript of my conversations with these LLMs in the Appendix.

I did not simply accept LLM outputs uncritically. Rather, I used them as preliminary explications that I verified against scientific literature where possible, cross-checked between different LLM responses, integrated with my own developing understanding and treated as interpretations rather than definitive explanations. The LLMs functioned as technical dictionaries or expert consultants, helping me understand what Troyer was saying so that I could then analyse how and why he was saying it using STS frameworks.

LLMs can produce inaccurate or oversimplified explanations, particularly for cutting-edge scientific concepts. I mitigated this risk through using multiple LLMs to cross-check explanations, requesting and checking references, focusing on fundamental concepts rather than disputed claims and consulting published scientific literature where terminology or concepts remained unclear.

3.3.3 Transcription

Troyer's colloquium required transcription from video format. I used OpenAI's Whisper for initial transcription, then employed Microsoft Copilot to assist manual translation of the spoken language into more readable written language whilst preserving meaning. The full transcript appears in the Appendix. I have endeavoured to maintain fidelity to Troyer's intended meaning whilst making the text readable. Readers interested in precise wording can consult the original video.

3.4 REFLEXIVE POSITIONING AND LIMITATIONS

3.4.1 My Positioning Revisited

As elaborated in chapter 1.2, my positioning as Microsoft employee and STS scholar creates both opportunities and constraints. My employment did not provide special access to documents—all materials analysed are publicly available. However, my familiarity with Microsoft's communication culture may have enabled more nuanced interpretation of rhetorical strategies.

I have attempted to maintain critical distance by grounding analysis in established STS frameworks, systematically identifying "issues" (following Asdal & Reinertsen, 2022), and treating Microsoft as one actor amongst others in quantum computing development. I analyse only publicly available materials, making no claims about internal processes.

Despite these efforts, my employment may create unconscious biases—potentially more sympathetic readings of Microsoft's challenges or blind spots regarding competitors' framings. I have attempted to mitigate these risks through explicit acknowledgement and focusing on *how* narratives work rather than *whether* claims are true.

3.4.2 Methodological Limitations

Five principal limitations follow from the methodological choices made in this thesis and should be borne in mind when assessing its conclusions.

First, the analysis covers only three documents from one company. It cannot support claims about quantum computing's narrative development across the field as a whole, nor about Microsoft's communications through channels not represented in the selected texts. The conclusions advanced concern how a specific process — the construction of technological credibility under conditions of narrative dependency — operates in these three cases. Extending these findings to other actors and contexts is a task for future research.

How Quantum Computing Escapes the Maze

Second, this thesis cannot observe laboratory practices, internal decision-making, or the day-to-day scientific work through which quantum computing's material development unfolds. Document analysis follows from narrative dependency rather than despite it, but the relationship between the public narratives analysed here and the laboratory realities they purport to describe remains, by the thesis's own premises, inaccessible to independent verification.

Third, despite using LLMs as interpretive aids, my technical comprehension of quantum physics remains limited relative to specialist expertise. The STS analysis focuses on how narratives work rather than whether technical claims are correct, so this limitation does not undermine the analytical conclusions; but technical errors in my comprehension of specific quantum physics concepts, whilst mitigated through the verification procedures described in Section 3.3.2, cannot be entirely excluded.

Fourth, the Majorana 1 announcement was analysed whilst its significance was still actively being contested within the quantum physics community. The scientific controversy is addressed in Chapters 6 and 7, but the limited time elapsed since the announcement means that assessments of its significance may look different from a longer temporal vantage point.

A fifth transparency note concerns the use of large language models beyond their role as interpretive aids for technical content. LLMs were also employed to assist in identifying relevant literature and to improve the readability and stylistic consistency of the thesis text. Literature suggestions were verified against academic databases before inclusion, and no reference was accepted on the basis of an LLM output alone. Stylistic revisions were assessed and approved by the author, with final responsibility for all formulations remaining mine. These uses are disclosed here in the interest of full methodological transparency.

4 THE FRONTIER IMAGINARY: STATION Q (2014)

We enter the maze at its origin point. This chapter analyses the first of three Microsoft documents: Jennifer Warnick’s 2014 ethnographic portrait of Station Q, Microsoft’s primary quantum computing laboratory. Applying the five-step analytical framework described in Chapter 3, I examine how this document constructs what I term the “frontier imaginary” — a vision of scientific exploration at knowledge’s edges, conducted by exceptional individuals pursuing fundamental research. The chapter demonstrates how narrative readiness is constituted through charismatic authority, expansive sociotechnical imaginaries, and boundary work at the earliest stage of Microsoft’s public quantum narration. It primarily addresses SQ1 (how Microsoft’s narrative changes across stages), SQ2 (discursive strategies for constructing authority and demarcating technological approaches), SQ3 (the sociotechnical futures this document constructs), and SQ4 (epistemological and political conditions of narrative dependency at the frontier stage).

4.1 CONTEXT AND MATERIAL FORM

Despite access constraints, Microsoft has publicly disseminated documentation through their institutional communication channel that provides limited visibility into Station Q, their primary quantum computing laboratory during Microsoft’s foundation of their quantum computing program.

Step 1:
Context and
Material
Form

The establishment of Station Q as a collaborative research entity between Microsoft and the University of California at Santa Barbara (UCSB) represents a significant boundary organisation that bridges corporate, academic, and implicitly, military interests (UCSB, 2025).

Jennifer Warnick’s (2014) ethnographic account of Station Q, whilst crafted as corporate communication, nonetheless offers valuable historical insight into the sociotechnical imaginaries that shaped quantum computing research at Microsoft during its formative institutional period. The report is published on Microsoft’s “Source” platform under the “Story Labs” header, indicating it belongs to a series about stories from Microsoft’s laboratories.

Material Presentation: The webpage opens with a large photograph of a green chalkboard covered in white and yellow chalk equations and diagrams, immediately evoking the atmosphere of an old-fashioned university classroom. The tactile memory of chalk dust creates a nostalgic link to traditional academic environments—a sharp contrast to the cutting-edge technology discussed in the text. Above this image, Station Q’s logo prominently displays a large red “Q” with shadow effects, establishing the laboratory’s distinct visual identity.

How Quantum Computing Escapes the Maze




Figure 1: chalkboard at Station Q

The report consists of five interconnected stories, each offering unique perspectives on various aspects of laboratory life. The stories employ different visual separators—red graphical elements resembling quantum pathways—and feature multiple photographs of researchers and laboratory spaces. The photographs function as what Latour (1986) would term "inscriptions"—material representations that stabilise and circulate knowledge claims.

Publication Context: Jennifer Warnick was a Microsoft employee when she created this report, following Microsoft's publishing guidelines. As corporate communication, the document inevitably presents Microsoft favourably, emphasising successes whilst potentially minimising challenges. The report's primary goal is highlighting Microsoft's advances and positioning the company as a leader in quantum computing innovation.

Target Audience: The report addresses multiple audiences simultaneously:

- General educated public interested in cutting-edge science
- Potential scientific recruits who might join Microsoft's quantum programme
- Scientific community establishing credibility for Microsoft's approach
- Investors seeking confidence in Microsoft's quantum strategy
- Technology journalists seeking stories about innovation

This multiplicity of audiences shapes the report's rhetorical strategies, requiring simultaneous accessibility for general readers and credibility for specialists.

4.2 CRITICAL ASSESSMENT OF ISSUES

Issue 1: Simplified Technical Explanations

Quantum computing is extraordinarily complex, and simplifying technical concepts for general audiences creates risks of inaccuracy or oversimplification. Terms like "qubits" and "topological qubits" are explained accessibly but potentially sacrifice scientific precision. The "corn maze" metaphor (discussed below) illuminates quantum computing's advantage but obscures technical details about how quantum superposition actually functions.

Step 2:
Critical
Assessment
of Issues

[+] Strategic Presence: The document makes quantum computing comprehensible to non-specialists, essential for mobilising public support and investment.

[-] Strategic Absence: The document omits technical details that would reveal the immense challenges, uncertainties, and competing interpretations within quantum physics itself.

Issue 2: Unverified Claims About Majorana Particles

The report treats the detection of Majorana zero modes as accomplished fact. However, subsequent scientific controversy has emerged regarding the validity of these claims (Castelvecchi, 2021, 2025). In 2014, the detection was presented as "compelling evidence" that "made the cover of Science," yet ongoing debates within the scientific community suggest the interpretation of experimental results remains contested.

[+] Strategic Presence: The document establishes Microsoft's scientific credibility through peer-reviewed publication in prestigious journals.

[-] Strategic Absence: The document does not acknowledge alternative interpretations of experimental data or acknowledge that "compelling evidence" differs from "definitive proof."

Issue 3: Omission of Failed Experiments

The narrative presents a success story—researchers pursue a bold strategy, detect exotic particles, and advance quantum computing. However, Warnick briefly mentions that Kouwenhoven "had a few failures, made a few adjustments" before achieving results. The failures receive minimal attention compared to successes.

[+] Strategic Presence: Brief acknowledgement of setbacks enhances credibility by suggesting honesty.

[-] Strategic Absence: The document does not explore what failed, why it failed, or how long failures delayed progress. The narrative arc moves quickly from challenge to triumph.

Issue 4: Lack of Independent Verification

All quoted sources are Microsoft employees or collaborators. No independent scientists, external validators, or critical voices appear. This creates what Collins (1992/1985) would recognise as a "core-set" problem—the boundary between insider expertise and independent assessment becomes blurred when only corporate voices define progress.

[+] Strategic Presence: Microsoft researchers speak with authority about their work.

[-] Strategic Absence: Independent scientific assessment remains absent, potentially creating a closed circle of mutual validation.

Issue 5: Temporal Vagueness

The report describes quantum computing's transformative potential using conditional and future tenses: "could tackle problems," "could have wildest imagination-type applications," "could allow humans to understand." These formulations acknowledge uncertainty whilst maintaining enthusiasm.

[+] Strategic Presence: The document avoids over-promising by using conditional language.

[-] Strategic Absence: No specific timelines appear. When will quantum computers actually solve these problems? The temporal horizon remains deliberately vague, protecting against falsification.

Issue 6: Downplayed Competition

The report mentions other quantum computing approaches briefly, noting that "physicists don't think we're crazy people anymore" (suggesting previous scepticism) and that the topological approach "has become mainstream." However, competing corporate efforts (IBM, Google, IonQ) receive minimal attention.

[+] Strategic Presence: Brief acknowledgement that other approaches exist enhances credibility.

[-] Strategic Absence: The document does not engage seriously with alternative approaches' advantages or Microsoft's potential disadvantages. The narrative constructs a world where Microsoft's approach is obviously superior once you understand it properly.

4.3 NARRATIVE ANALYSIS: FIVE STORIES OF QUANTUM PIONEERING

Warnick's report employs sophisticated narrative techniques to construct quantum computing's meaning and Microsoft's role in its development. Each of the five stories performs distinct rhetorical work whilst contributing to an overarching narrative arc.

Step 3:
Narrative
Analysis

4.3.1 Story 1: "The Quest for a Quantum Future"

Narrative Function: This introductory story establishes setting, characters, and stakes.

Setting: Station Q is located on the California coast, blending "world-renowned experts" with "Hawaiian shirts and shorts," "faded wetsuits and battered, loaner surf boards." This geographical and cultural positioning is significant—California represents technological innovation, counterculture, and relaxed creativity. The juxtaposition of serious physics with surfing culture constructs Station Q as unconventional, creative, and distinctively Californian.

Character Introduction: Michael Freedman, Station Q's director, is introduced in almost mythical terms: "stately, fit, and well-tanned... looks a bit like heroic police chief Martin Brody." This comparison to the protagonist of Spielberg's "Jaws" (1975) is revealing. Brody confronts a terrifying threat (the shark) through a combination of determination, intelligence, and courage. The metaphor positions Freedman as confronting the "terrifying" challenge of quantum computing with similar heroic qualities.




Figure 2: Station Q Director Michael Freedman at the chalkboard (photo by Brian Smale for Microsoft)

Warnick continues the "Jaws" reference with the famous quote: "We're going to need a bigger boat." In the film, this line acknowledges that the characters' resources are inadequate for the

How Quantum Computing Escapes the Maze

challenge they face. Applied to quantum computing, it suggests that classical computers are inadequate—we need the "bigger boat" of quantum computing to solve certain problems.

Stakes: The document employs technological utopianism, describing quantum computing as potentially having "wildest imagination-type applications in fields such as machine learning and medicine, chemistry and cryptography, materials science and engineering." Freedman's vision suggests quantum computers "could allow humans to understand and control the very building blocks of the universe"—a remarkably ambitious claim positioning quantum computing as fundamental to human knowledge, not merely practical problem-solving.

Narrative Tone: The story is told with wonder and excitement. Researchers use words like "strange and unusual," "mind-bending," "exotic," "magical," and "beyond science fiction" to describe their work. This language constructs quantum phenomena as mysterious and almost supernatural, positioning quantum researchers as exploring unknown territories.

4.3.2 Story 2: "Particles Gone Wild"

Narrative Function: This story provides general education about quantum computing for non-specialist readers.

The Maze Metaphor: Warnick introduces a central metaphor that will reappear throughout Microsoft's quantum narrative:

The difference between quantum and classical computing is all in the approach. Classical computers attack problems like you would navigate a corn maze, those farm-size labyrinths popular in rural areas at harvest time. It proceeds down each long, stalk-lined corridor and at each fork, it picks one direction. If it reaches a dead end, it turns around, finds its way back, and tries another route until eventually it solves the maze (unless, of course, the maze is so massive that examining every route takes the lifetime of the universe). (Warnick, 2014, story 2, para. 14)

This metaphor performs crucial work. It constructs classical computing as sequential (one path at a time), methodical but potentially inefficient and potentially facing problems so large they become practically unsolvable.

The quantum alternative is vividly described:

Imagine that same corn maze, but instead of looking for the way out on foot, one turn at a time, you unleash a pack of high-octane, well-trained Tribbles. Those fuzzy, fictional Star Trek creatures would move out in every direction at once and, thanks to their tendency to

multiply at frighteningly exponential rates, explore every possible route simultaneously to quickly find the most efficient solution. (Warnick, 2014, story 2, para. 14)

This "Tribbles" metaphor is playful (referencing popular culture) whilst capturing quantum computing's essential feature—exploring multiple possibilities simultaneously through superposition. The "exponential multiplication" references quantum computing's exponential scaling advantage for certain problems.

Expert Voices: The story introduces Krysta Svore, discussing quantum computing simulation on classical hardware. This demonstrates Microsoft's multi-pronged approach—using classical computers to understand what quantum computers might achieve.

Historical Parallel: Peter Lee draws a comparison between quantum computing's potential impact and the transistor's historical significance. This parallel performs important work—it positions quantum computing within a narrative of revolutionary technological transformations that fundamentally reshape society. Just as the transistor enabled the digital age, quantum computing might enable a new age of computation.

Fundamental Transformation: Freedman concludes this story with a remarkable claim:

What is so amazing about quantum mechanics is it's the fundamental microscopic language of the universe. It's the way the universe talks to itself, and we don't think that way. We're more like classical computers... With this project to develop quantum computing, what we're really doing is making the transition as a species from our devices thinking in this very kind of clumsy, classical model, to our devices thinking in the fundamental language of the universe. So, we're going to be leaving behind this more primitive method of processing information, and we'll move into the quantum realm. .
(Warnick, 2014, story 2, para. 14)

This passage anthropomorphises nature ("the way the universe talks to itself") whilst positioning human development teleologically—we are evolving from "primitive" classical methods to quantum thinking. This framing suggests quantum computing is not just a new technology but a fundamental transformation in how humans relate to physical reality.

4.3.3 Story 3: "Now Arriving at Station Q"

Narrative Function: This story provides Station Q's founding narrative, establishing origins and mission.

The Founding Dialogue: The story centres on a conversation between Michael Freedman (mathematician) and Craig Mundie (corporate strategist):

Freedman: "Well Craig, I'm just a mathematician." Mundie: "Not anymore. Now you're a mathematician and a program manager." (Warnick, 2014, story 3, para. 9, 10)

This exchange performs boundary work. Freedman identifies with pure mathematics ("just a mathematician"), suggesting discomfort with management roles. Mundie's response redefines Freedman's identity—he must now bridge mathematics and organisational management. This dialogue captures a recurring theme in science studies: the tension between pure research and institutional/commercial objectives.

Naming as Identity: Freedman explains Station Q's name:

It's not a leisurely place, not a place for sipping tea. I thought it would go over much better as a research station. An outpost. Somewhere we'd go to just get the job done. (Warnick, 2014, story 3, para. 12)

The "outpost" metaphor positions Station Q as frontier exploration—similar to Antarctic research stations or space missions. Researchers are pioneers at the edge of knowledge, working in challenging conditions to advance human understanding. The rejection of "leisurely" work and "sipping tea" constructs Station Q as serious, focused, and mission-driven.




Figure 3: A text separator from Warnick's report looking like a maze or a knot

Validation Through Difficulty: Peter Lee notes that Station Q's conferences are "so notoriously complex that after a few minutes into one even he starts to feel a bit thick... and small." This acknowledgement of difficulty performs important work—it establishes the research's

sophistication whilst constructing a hierarchy of expertise. Even Microsoft's research director feels intellectually overwhelmed, reinforcing how advanced this work is.

Lee continues: "Up until recently, maybe two years ago, all of this was so speculative. Even in scientific circles this was considered out in the fringes... Now I think you could argue that the topological approach has become mainstream. Physicists don't think we're crazy people anymore."

This passage constructs a narrative of vindication. Station Q pursued a "fringe" approach whilst others were sceptical. But through perseverance and results, the approach gained mainstream acceptance. This narrative positions Microsoft as visionary—willing to pursue long-term, risky research that others dismissed.

Character Construction: The introduction of Sankar Das Sarma demonstrates how the report constructs scientific authority:

Buttoned-up with perfect posture, Das Sarma is an initially serious-seeming guy who it turns out is to his field what Neil DeGrasse Tyson is to the cosmos—the man can explain even the most complicated ideas about quantum physics like he's recapping last night's "Game of Thrones" episode. (Warnick, 2014, story 3, para. 18)

Comparing Das Sarma to Neil deGrasse Tyson (famous science communicator) positions him as both expert and accessible. The "Game of Thrones" reference suggests he can make complex physics entertaining and comprehensible—essential for public engagement.

4.3.4 Story 4: "Hunting Quasiparticles"

Narrative Function: This story narrates the central scientific achievement—detecting Majorana particles.

Scientific Narrative as Detective Story: The story employs a discovery narrative:

1. Theoretical Prediction: Alexei Kitaev proposes (in 2000) that Majorana particles could be used in quantum information processing. (Warnick, 2014, story 4, para. 1)
2. Theoretical Development: Researchers including Sankar Das Sarma, Michael Freedman, and Chetan Nayak develop the theoretical framework.
3. Experimental Challenge: How to actually observe these elusive particles?
4. Recruitment: Charlie Marcus joins Microsoft, bringing experimental expertise. He recruits Leo Kouwenhoven for his technical skills with nanowires.

How Quantum Computing Escapes the Maze

5. **Competitive Collaboration:** Marcus and Kouwenhoven are described as "friend-slash-rival" who "share recipes, but each of us still wants to be the best cook." This competitive friendship is framed as generative—"It's fun to compete. We're trying to advance the field." (Warnick, 2014, story 4, para. 7)
6. **Experimental Breakthrough:** After "a few failures" and "adjustments," Kouwenhoven's team detects compelling evidence of Majorana particles. The discovery "made the cover of Science."

This narrative structure follows classic discovery stories in science—theoretical prediction, experimental challenge, breakthrough, validation through publication. The story constructs scientific progress as collaborative yet competitive, requiring both theoretical insight and experimental skill.

The Donut Metaphor: Freedman explains topological protection using a donut metaphor:

If you have a donut with a hole in the middle, it's still a donut with a hole in the middle even if you take a large bite out of it. (Warnick, 2014, story 4, para. 14)

This metaphor translates abstract mathematical topology into everyday experience. A donut's fundamental property (the hole) persists despite local perturbations (taking a bite). Similarly, topological qubits' quantum information persists despite local environmental disturbances. The metaphor performs pedagogical work whilst also suggesting robustness and stability.

Corporate Positioning: Peter Lee situates this research within Microsoft's corporate strategy:

Our primary motivating factor is science... but we think that science could be developing the foundations for a new kind of technology—sort of a post-silicon age. (Warnick, 2014, story 4, para. 15)

This framing manages competing imperatives. Microsoft pursues "science" (suggesting purity of motive) but acknowledges commercial ambitions ("new kind of technology"). The "post-silicon age" positions quantum computing as succeeding the semiconductor era—as revolutionary as the transition from vacuum tubes to transistors.

Competitive Framing: Peter Lee describes the Majorana detection as "a landmark finding that kicked off a topological quantum computing craze in the physics community." This framing suggests Microsoft catalysed broader scientific excitement, positioning the company as leader rather than follower.

4.3.5 Story 5: "From Here to Infinity"

Narrative Function: This concluding story addresses broader implications and strategic positioning.

The Tortoise and the Hare: Craig Mundie employs Aesop's fable to contrast quantum computing approaches:

Some groups have built machines with quantum properties that might hit a wall in terms of their ability to scale up. Station Q's topological approach may be theoretically harder to get off the ground initially, but if and when they get it working, it will be very scalable. (Warnick, 2014, story 5, para. 2)

This metaphor positions competing approaches (NISQ devices) as fast initially but ultimately limited (the hare), whilst Microsoft's approach is slower but ultimately superior (the tortoise). The fable's moral—slow and steady wins the race—justifies Microsoft's longer timeline whilst suggesting competitors' early demonstrations are misleading.

Institutional Identity: Station Q is described as "a treasure trove of research" taking "a more soup-to-nuts approach to quantum problems than any other organization." This constructs comprehensiveness as advantage—Microsoft addresses the full spectrum of quantum challenges rather than narrow problems.

Dual Orientation: Mundie explicitly acknowledges commercial motivation:

Our pursuit is not strictly academic in nature. We dream a dream this will one day inform our product strategy and be of utilitarian and economic importance. (Warnick, 2014, story 5, para. 4)

This statement is significant. Whilst earlier passages emphasised pure science, Mundie here acknowledges commercial ambitions. The phrase "dream a dream" maintains visionary language whilst acknowledging profit motive.

Economic Imaginary: The report articulates quantum computing's economic significance:

The economic implications could be staggering... a second-coming of the computing age... would utterly dwarf the power of today's computers (Warnick, 2014, story 5, para. 5)

This constructs quantum computing as potentially generating immense economic value, justifying massive investment.

How Quantum Computing Escapes the Maze

Fundamental Uncertainty: Despite enthusiasm, the report acknowledges uncertainty. Quantum research is compared to "a bit of a boomerang" that researchers throw "with as much force as they can, but there's no guarantee of when or how it will come back (or whether it will at all)."

This acknowledgement of uncertainty is unusual in corporate communication. It suggests honesty and humility, potentially enhancing credibility. The researchers defend basic research where "in many cases we have absolutely no idea" of applications. Marcus argues "The truth is always relevant eventually" and "Dormancies frequently follow important discoveries, which will be brought back up at their right time."

This defence of basic research without immediate applications positions Microsoft as patient, willing to pursue long-term science rather than short-term engineering.

Unknown Applications: The report concludes by emphasising that "nobody knows" what quantum computing applications might emerge. Das Sarma states:

I'm not being humble. We have no idea what a quantum computer could do. Factor numbers? Make better drugs and materials? These are all possible applications. There's a good case to be made that quantum computing could change the world. (Warnick, 2014, story 5, para. 10)

This combination of uncertainty ("no idea") with confidence ("could change the world") is striking. It acknowledges ignorance whilst maintaining transformative potential.

Philosophical Conclusion: The report concludes with Marcus's reflection:

We live with the belief that what we see is all there is. The world is a much richer place than your first observation tells you. (Warnick, 2014, story 5, para. 17)

This frames quantum research as revealing hidden realities, positioning quantum physics as expanding human understanding of what exists. The final phrase—"unravel things which are true but have been veiled"—constructs quantum research as unveiling fundamental truths, connecting it to broader human intellectual achievement.

4.4 STS INTERPRETATION: CONSTRUCTING THE VANGUARD

Warnick's 2014 report functions as what Czarniawska (2004) terms a constitutive narrative: it does not merely describe Station Q but actively constructs Microsoft's identity as a quantum computing vanguard. Three analytical dimensions illuminate how this construction works. The first concerns authority — how the report establishes legitimate expertise in the absence of

Step 4:
Theoretical
Inter-
pretation

demonstrated results. The second concerns imaginaries — how the report projects quantum computing’s future scope and significance. The third concerns performativity — how the report enacts the very realities it purports to document. These dimensions are analytically distinct but operate together as an integrated legitimation strategy, responding to the research questions in different registers.

4.4.1 Constructing Authority at the Frontier Stage

At the frontier exploration stage, where demonstrated results cannot anchor authority claims, the report relies on two closely related strategies: charismatic figuration and boundary work. Freedman’s portrayal — Fields Medal at 36, Berkeley at 16, rock climber, swimmer, and hiker — constructs what Hilgartner (2015) terms vanguard leadership through extraordinary individual qualities rather than institutional position alone. This heroic figuration is reinforced by the report’s careful staging of Station Q’s cultural identity: Hawaiian shirts, chalkboard equations, California coastline. The setting signals that scientific excellence need not follow conventional institutional forms, positioning Station Q in liminal space between academia and corporation, between established physics and emergent technology. As Gieryn (1983) argues, boundary work of this kind constructs legitimacy by demarcating an in-group of genuine experts from surrounding institutions.

The Freedman-Mundie dialogue performs this boundary work with clarity. Freedman’s self-identification as “just a mathematician” signals attachment to disciplinary purity and resistance to managerial framing, whilst Mundie’s response — repositioning Freedman’s identity as “an architect of computation” — navigates the tension between pure research and commercial ambition that runs throughout the document. The report carefully manages this tension: early passages emphasise scientific curiosity, whilst later passages acknowledge commercial aspiration through the distancing phrase “dream a dream.” Drawing on Latour and Woolgar’s (1986/1979) framework, these inscription practices transform laboratory work into circulating claims. The photographs of researchers at chalkboards function as immutable mobiles, stabilising and exporting Station Q’s epistemic culture to audiences who cannot observe it directly — which is precisely the narrative dependency condition this thesis analyses.

Authority is further constructed through expertise hierarchies that distribute credibility strategically. Lee’s admission that conference presentations make him “feel a bit thick... and small” establishes Station Q researchers as possessing knowledge exceeding even senior Microsoft leadership. Das Sarma’s characterisation as able to explain quantum physics “like

How Quantum Computing Escapes the Maze

he’s recapping last night’s Game of Thrones episode” simultaneously signals extraordinary expertise and the capacity to translate it outward — the expert as boundary-spanner. Peter Lee’s modal qualification that Majorana detection represents “not definitive proof, but very strong evidence” illustrates how claim status is carefully managed: maintaining credibility through appropriate hedging whilst positioning the work as substantively significant. Taken together, these strategies construct charismatic authority that can sustain vanguard positioning when institutional validation and demonstrated results are not yet available.

4.4.2 Constructing the Expansive Imaginary

The sociotechnical imaginary that the report constructs is characteristic of the frontier exploration stage: maximally expansive, temporally vague, and institutionally anchored. Following Jasanoff and Kim’s (2009, 2015) framework, imaginaries are collectively held visions that organise expectations and mobilise resources. The report articulates these visions through shared language — quantum computing as “mind-bending,” “exotic,” “magical,” and “world-changing” — across multiple researchers, constructing collective endorsement. This shared vocabulary performs institutional stabilisation: Microsoft’s long-term investment in laboratory facilities and researcher recruitment materialises the imaginary in concrete organisational form, lending it durability beyond individual declarations.

The imaginary’s application scope is deliberately broad. “Wildest imagination-type applications in fields such as machine learning and medicine, chemistry and cryptography, materials science and engineering” enrolls diverse stakeholder communities whilst the absence of demonstrated constraints permits indefinite expansion. This breadth is strategically functional: without specific claims to falsify, the imaginary can sustain enthusiasm across a wide range of potential investors, partners, and recruits. The corn maze metaphor — classical computing “down each long, stalk-lined corridor” versus quantum computing as “a pack of high-octane, well-trained Tribbles” exploring all paths simultaneously — constructs classical limitations as fundamental rather than contingent, delegitimising alternatives whilst positioning quantum computing as the only path to escaping them.

The imaginary’s temporal construction combines inevitability with protective vagueness, following Selin’s (2008) analysis of future-making. Freedman’s statement that “we’re going to be leaving behind this more primitive method of processing information” positions quantum transition as inevitable progression rather than contingent possibility. Yet the research is simultaneously compared to “a boomerang” thrown “with as much force as they can, but there’s no guarantee of when or how it will come back.” This combination — directional

confidence without temporal commitment — is characteristic of frontier-stage imaginaries. Acknowledging uncertainty enhances credibility whilst preserving the transformative claim: Freedman’s admission that “we don’t know yet” coexists with confidence that quantum computing “could change the world.” The imaginary is expansive precisely because it has not yet been tested against material reality.

4.4.3 Performing the Vanguard into Existence

Drawing on Austin’s (1962) speech act theory as developed by Callon (1998) and MacKenzie (2006), the report functions not merely as description of Microsoft’s quantum position but as performative utterance constituting that position. When Warnick declares that Station Q operates “at the cutting edge of an emerging scientific field,” this statement does not represent a pre-existing reality awaiting documentation — it actively constitutes that reality through circulation and uptake. The report’s public performance, circulating quantum imaginaries to audiences who lack independent access to the laboratory, exemplifies what Callon (1998) terms the performative dimension of descriptions: they do not merely represent realities but enact them. The topological approach is declared “unique”; uniqueness is thereby performed into significance, potentially shaping research priorities, investment decisions, and hiring patterns in ways that could make the declaration self-fulfilling.

Performativity at the frontier stage operates, however, under characteristic constraints. Without demonstrated results, performative declarations lack material anchoring and must compete against rival claims. The report’s acknowledgement that “physicists don’t think we’re crazy people anymore” reveals this tension: the declaration of mainstream acceptance simultaneously acknowledges prior rejection, suggesting that earlier performative attempts met resistance. The boundary work analysed above — the cultural positioning, the charismatic figuration, the expertise hierarchies — functions partly as a response to this constraint, seeking to establish authority through means other than material demonstration. Knorr Cetina’s (1999, 2010) analysis of epistemic cultures further illuminates Station Q’s distinctive knowledge-making environment: it combines the contemplative (Freedman’s “long walks on the beach”), the large-scale collaborative (international conferences and researcher networks), and the hands-on experimental (Kouwenhoven’s iterative adjustments to detect Majorana particles). Crucially, external observers access this epistemic culture only through Warnick’s narrative mediation — a condition of narrative dependency that the report itself both exemplifies and, through its accessible language and vivid metaphors, partially compensates for. This condition is not merely epistemological but political: in the absence of

How Quantum Computing Escapes the Maze

any independent account of Station Q's work, Microsoft holds exclusive authority to determine what quantum computing is understood to be, what progress looks like, and whose assessment of both is legitimate.

4.5 SYNTHESIS: THE FRONTIER IMAGINARY

Step 5: Brief
Synthesis

Warnick's 2014 report constructs what I term the "frontier imaginary" of quantum computing—a vision of scientific exploration at knowledge's edges, conducted by exceptional individuals in unconventional environments, pursuing fundamental research that might revolutionise computing and human understanding of reality itself. This frontier imaginary represents the first stage in Microsoft's vanguard evolution, whose theoretical significance the Discussion chapter develops.

Key Findings

The Station Q report reveals how technological readiness is narratively constituted at early developmental stages. Microsoft constructs readiness not through demonstrated quantum computers but through narrative strategies establishing that the *conditions* for eventual readiness exist: exceptional researchers, appropriate institutional arrangements, promising theoretical frameworks, and preliminary experimental results. The modal constructions pervading the document ("could tackle," "could allow," "could change") position quantum computing as ready-to-become-ready—possessing validated potential despite the absence of working devices.

Authority at this stage derives primarily from charismatic figuration. Freedman's portrayal emphasises extraordinary individual qualities—Fields Medal at 36, heroic comparisons—rather than institutional validation or demonstrated results. This charismatic authority enables vanguard positioning when material evidence remains limited but creates vulnerabilities that subsequent stages must address.

The sociotechnical imaginary exhibits maximum expansion characteristic of frontier conditions. Applications span "wildest imagination-type" possibilities across medicine, cryptography, materials science, and machine learning. This expansive scope enrolls diverse stakeholder communities whilst the absence of constraining results permits ambitious claims. Simultaneously, temporal vagueness protects against falsification—the "boomerang" metaphor acknowledges uncertain return trajectories.

The report's boundary work establishes Station Q's distinctiveness through cultural positioning (California casualness combined with intellectual intensity), disciplinary hybridity (mathematics,

physics, computer science), and institutional liminality (neither purely academic nor purely corporate). These boundaries differentiate Microsoft's approach without requiring comparison to competitors' results—collaborative framing prevails over competitive positioning.

The Frontier Imaginary's Functions

The frontier imaginary performs multiple functions: legitimating Microsoft's quantum investment through heroic narratives of scientific breakthrough; differentiating Microsoft's approach (topological, patient, fundamental) from competitors through the tortoise-and-hare framing; recruiting scientific talent by portraying Station Q as intellectually stimulating environment; balancing enthusiasm with acknowledged uncertainty; and mobilising support from investors, policymakers, and the scientific community by constructing quantum computing as simultaneously fascinating and economically significant.

Transition Pressures

The frontier imaginary contains tensions that subsequent stages must address. Charismatic authority cannot indefinitely substitute for demonstrated results. Indefinite temporal vagueness eventually generates stakeholder impatience. As competitors demonstrate quantum capabilities, collaborative framing becomes less tenable. Expansive imaginaries risk generating expectations that material progress cannot satisfy. These pressures drive the transition from frontier exploration to pragmatic navigation evident in Troyer's 2021 colloquium.

These internal tensions are precisely what the next chapter examines. Chapter 5 analyses how Troyer's 2021 colloquium addresses the pressures generated by the frontier imaginary — replacing charismatic authority with expert assessment, expansive possibility with systematic demarcation, and protective vagueness with calibrated milestones. The maze, as metaphor, makes its first explicit appearance in that document; it is there that this thesis's navigational device originates.

5 MANAGING EXPECTATIONS: NAVIGATING THE QUANTUM MAZE (2021)

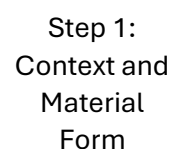
The frontier imaginary analysed in Chapter 4 contains an internal tension: charismatic authority and expansive promises cannot indefinitely substitute for demonstrated results. By 2021, that tension had become acute. This chapter analyses Troyer’s “Towards Practical Quantum Advantage” colloquium, in which Microsoft’s narrative undergoes its most significant recalibration — what I term the “pragmatic turn.” Where Warnick constructed possibility, Troyer systematically dismantles hype. The chapter shows how that dismantling is itself a narrative act with strategic functions and introduces the maze as Troyer’s own framing device — one I appropriate reflexively throughout this thesis. The chapter primarily addresses SQ1 (how Microsoft’s narrative changes under credibility pressure), SQ2 (discursive strategies of demarcation and expectation management), SQ3 (how the projected futures for quantum computing shift in this document), and SQ4 (the political and epistemological conditions of narrative dependency under technical complexity).

5.1 CONTEXT AND MATERIAL FORM

Approximately five years after Jennifer Warnick’s ethnographic portrait of Station Q, the articulation of Microsoft’s quantum computing vision shifted to a new figurehead whose perspective would significantly reshape the company’s technological narrative. Whilst Freedman’s Station Q represented the pioneering scientific vanguard, Matthias Troyer embodied a strategic pivot towards practical implementation and corporate integration of quantum computing technologies.

As a former professor of computational physics at ETH Zürich with an illustrious academic career, Troyer’s 2019 transition to become Microsoft’s “distinguished scientist” and corporate vice president for quantum computing signalled a significant evolution in how the company positioned its quantum efforts. This transition from academia to corporate leadership represented more than a personnel change; it marked a fundamental shift in how Microsoft constructed and communicated its quantum computing imaginary to both scientific communities and potential commercial stakeholders.

Publication Context: Troyer’s (2021) colloquium titled “Towards Practical Quantum Advantage” was delivered on 16 March 2021 as part of UC Berkeley’s Simons Institute quantum computing colloquium series. Due to COVID-19 restrictions, the event was livestreamed and subsequently made available on YouTube. The full presentation spans 65 minutes, with Professor Umesh Vazirani of UC Berkeley’s Electrical Engineering department serving as host. Following the main



Step 1:
Context and
Material
Form

presentation, a question session commenced at minute 58, during which two individuals posed questions.

Material Presentation: Troyer's presentation comprises 37 slides formatted in UC Berkeley's institutional style. The slides' layout alternates between solid blue backgrounds and white backgrounds with dark-grey footers. Troyer uses a clear sans-serif font throughout. Each slide features the UC Berkeley logo in the bottom-right corner, establishing institutional authority. Throughout the presentation, a small window appears in the upper-right corner displaying a half-length portrait of Troyer, who stands against a backdrop of sunny blue skies and trees—a domestic setting reflecting COVID-19's impact on academic communication.

Target Audience: The colloquium targets scholars possessing sophisticated knowledge in disciplines such as general physics, quantum physics, chemistry, materials science, and electrical engineering. The audience's expertise is particularly high-level, a point underscored by Peter Lee, Microsoft's research director, who characterises such presentations as "notoriously complex that after a few minutes into one even he starts to feel a bit thick [and] small" (Warnick, 2014).

Temporal Positioning: The colloquium occurs at a critical juncture in quantum computing's development. By March 2021, NISQ (Noisy Intermediate-Scale Quantum) devices had been demonstrated by multiple corporations. Google had announced "quantum supremacy" in 2019, claiming their "quantum processor performed a calculation that would be practically impossible for classical computers" (Arute et al., 2019). This created both excitement and scepticism—was quantum advantage imminent, or were these demonstrations misleading? Troyer's colloquium directly addresses this tension, explicitly framing it against what he terms the "quantum wishing well" of unrealistic expectations. Crucially, the colloquium was delivered in the same year that Zhang et al.'s Nature paper (2018) — which had claimed evidence for Majorana zero modes, the quasiparticle central to Microsoft's entire topological strategy — was formally retracted (2021) following the identification of data inconsistencies by Frolov and Mourik (2021). Read against this context, Troyer's systematic effort to rebuild credibility through rigorous demarcation of the viable from the non-viable is analytically legible not only as forward-looking expectation management but as a communicative response to a credibility crisis. The colloquium performs the work of authority reconstruction that the retraction had made urgent.

5.2 CRITICAL ASSESSMENT OF ISSUES

Following Asdal and Reinertsen's (2022) methodology, I critically assess what Troyer's colloquium makes present and absent:

Step 2:
Critical
Assessment
of Issues

Issue 1: Microsoft-Centric Perspective

All examples, estimates, and future projections reflect Microsoft's topological approach. Troyer dismisses NISQ devices (the approach pursued by competitors like IBM, Google, and Rigetti) as combining "the disadvantages of analogue being noisy and digital being restricted" (Troyer, 2021, 59:59). This creates a strategic framing where Microsoft's approach appears uniquely viable.

[+] Strategic Presence: Troyer establishes Microsoft's topological approach as the only pathway to scalable quantum computing, justifying Microsoft's investment and slower progress.

[-] Strategic Absence: The colloquium does not seriously engage with competitors' perspectives or acknowledge that other approaches might achieve quantum advantage through different pathways. Alternative viewpoints are dismissed rather than debated.

Issue 2: Optimistic Timeline Assumptions

Troyer states: "With 100,000 physical qubits and fault-tolerant simulations, we can solve quantum problems that are classically hard. That's not so far out; I think we can achieve that in the next few years" (30:31). This timeline proves significant—"next few years" from March 2021 would be 2023-2025. As of the Majorana 1 announcement (February 2025), Microsoft has not yet achieved 100,000 physical qubits.

[+] Strategic Presence: The optimistic timeline generates enthusiasm and justifies continued investment whilst appearing more realistic than wild speculation.

[-] Strategic Absence: The colloquium does not acknowledge the extensive technical challenges remaining or provide contingency timelines if assumptions prove incorrect.

Issue 3: Classical Computing as Static Benchmark

Troyer's analysis compares quantum computers against current classical hardware: "a single today's classical chip, the most advanced NVIDIA chip with 54 billion transistors" (6:12). However, classical computing continues advancing. Quantum-inspired algorithms—which Troyer himself discusses—improve classical performance for problems once thought to require quantum hardware.

[+] Strategic Presence: Using current classical hardware as benchmark makes quantum advantage appear more achievable.

[-] Strategic Absence: The colloquium does not extensively address the moving target problem—classical computing improvements might outpace quantum developments, perpetually delaying practical quantum advantage.

Issue 4: Hype Critique as Positioning Strategy

Troyer explicitly critiques quantum computing hype, dismissing applications like "big data, cure cancer, design new drugs, proteins, predict the stock market" as "the hope for a quantum wishing well" (4:02). This critique performs strategic work—it positions Troyer as realistic and credible whilst delegitimising competing approaches and applications.

[+] Strategic Presence: The hype critique enhances Troyer's credibility, suggesting he provides honest assessment rather than marketing.

[-] Strategic Absence: The colloquium does not acknowledge that Microsoft itself has contributed to quantum hype through previous communications. The critique positions Microsoft as correcting others' errors rather than acknowledging shared responsibility for inflated expectations.

Issue 5: Limited Discussion of Fundamental Challenges

Whilst Troyer discusses technical requirements (error correction, qubit numbers), he provides limited detail about fundamental obstacles. For instance, the challenge of maintaining quantum coherence at scale, the difficulty of fabricating topological qubits reliably, and the theoretical debates about whether Majorana particles have actually been observed receive minimal attention.

[+] Strategic Presence: Focusing on quantifiable requirements (qubit numbers, error rates) makes challenges appear tractable—essentially engineering problems requiring resources and time.

[-] Strategic Absence: The colloquium minimises fundamental scientific uncertainties that might make quantum computing impractical regardless of engineering investment.

Issue 6: Ambiguity About "Practical Quantum Advantage"

Troyer defines practical quantum advantage as "really solving a useful problem that is useful either academically or for industry—faster and better on a quantum computer than on any known classical computer" (3:04). This definition is deliberately broad—"useful either academically or

How Quantum Computing Escapes the Maze


for industry" encompasses virtually any computational problem. What counts as "useful"? Who decides?

[+] Strategic Presence: The broad definition maximises potential applications, maintaining enthusiasm about quantum computing's scope.

[-] Strategic Absence: The definition avoids specifying concrete benchmarks or success criteria, protecting against falsification. If quantum computers eventually solve obscure academic problems faster than classical computers, this technically satisfies the definition even if commercial impact remains minimal.

5.3 NARRATIVE ANALYSIS: AT THE ENTRANCE TO THE MAZE

Troyer's colloquium employs sophisticated narrative and rhetorical strategies to navigate between quantum computing's promise and its hype. Unlike Warnick's ethnographic storytelling, Troyer employs technical argumentation, systematic comparison, and metaphorical frameworks to construct his vision of quantum computing's viable pathway.



Step 3:
Narrative
Analysis

5.3.1 Opening: Establishing Motivation and Identity

Troyer begins by explaining his transition from academia to Microsoft:

I got to the point where I said to really make progress, I would love to have quantum hardware because I know how to solve these problems really easily once we have a quantum computer. So, that's why I joined Microsoft. (Troyer, 2021, 1:25)

This opening performs important identity work. Troyer positions himself as scientist whose research needs drove corporate engagement rather than corporate employee promoting products. The phrase "I know how to solve these problems really easily" suggests confidence whilst the conditional "once we have a quantum computer" acknowledges that hardware remains unavailable. This establishes Troyer as pragmatic visionary—someone who understands both quantum computing's potential and its current limitations.

He then defines his role: "systems architect" (2:05), applying familiar architectural metaphors to quantum system development. This framing positions him as designer concerned with overall system structure rather than narrow technical specialist. Systems architects balance competing requirements, make trade-offs, and envision how components integrate—precisely the work Troyer performs in this colloquium.

5.3.2 Defining the Quest: "Practical Quantum Advantage"

Troyer articulates his central concern:

What I'm interested in is what I call practical quantum advantage. I'm interested in really solving a useful problem that is useful either academically or for industry—faster and better on a quantum computer than on any known classical computer. (3:04)



Figure 4: Matthias Troyer defines his research interests

This definition carefully balances multiple considerations:

- "Practical": Distinguishes from theoretical or asymptotic advantage. Troyer cares about actually solving problems in reasonable timeframes, not merely proving quantum advantage exists in principle.
- "Useful problem": Excludes contrived demonstrations. Random circuit sampling, which Google used for quantum supremacy claims (Arute et al., 2019), doesn't qualify because it serves no purpose beyond demonstrating quantum capability.
- "Either academically or for industry": Keeps definition broad enough to include diverse applications whilst suggesting commercial relevance matters.
- "Faster and better": Requires not just speed but quality—producing more accurate results or enabling previously impossible calculations.
- "Any known classical computer": Establishes high bar—quantum computers must beat not just typical classical hardware but optimised supercomputers and specialised algorithms.

Troyer immediately reinforces this definition negatively:

How Quantum Computing Escapes the Maze

I'm not interested in whether some random circuit can be done better on quantum hardware than classical hardware. I really want to solve a problem that interests somebody and do it better than on any classical computer. (Troyer, 2021, 3:22)

This explicit dismissal of random circuit demonstrations positions Troyer against competitors (particularly Google) who claimed quantum supremacy through such demonstrations. The phrase "interests somebody" adds subjective criterion—problems must have constituency caring about solutions, not merely technical achievement.

5.3.3 Introducing the Maze: Quantum Applications Landscape

Troyer presents a slide titled "A maze of proposed applications" showing numerous claimed quantum computing applications:

When you listen to people, there's a maze of proposed applications. One reads a lot in the media about how quantum computers will help with big data, cure cancer, design new drugs, proteins, predict the stock market, and so on. There are loads of things where people claim that quantum will solve all these problems. (Troyer, 2021, 4:02)

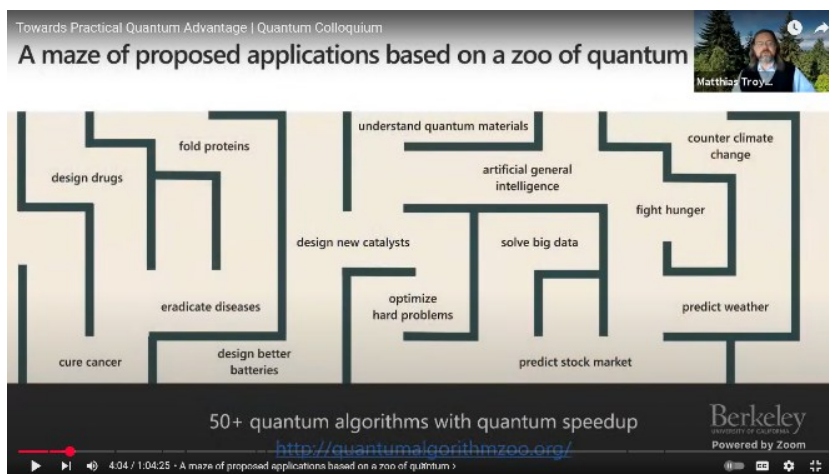


Figure 5: Matthias Troyer introducing the maze of quantum computing

The "maze" metaphor is central to Troyer's presentation. Mazes contain multiple pathways, most leading to dead ends, with only some reaching the centre. This metaphor constructs the quantum computing landscape as:

- Complex and Confusing: Multiple proposed applications make it difficult to discern viable paths from false leads.
- Requiring Navigation: Success requires systematic exploration and evaluation, distinguishing productive paths from dead ends.

- Containing Hidden Structure: Whilst the maze appears chaotic, underlying principles (computational complexity, physical constraints) determine which paths succeed.
- Potentially Solvable: Despite complexity, the maze has solutions—specific applications where quantum advantage is achievable.

Troyer immediately follows with dismissive framing:

But looking at that, for each of these, there's always one known quantum algorithm that can somehow be applied. But this looks more to me like the hope for a quantum wishing well. (Troyer, 2021, 4:02)

The "quantum wishing well" metaphor is devastating. Wishing wells are places where people throw coins hoping desires materialise magically. The metaphor suggests quantum computing has become a repository for wishful thinking—people attribute problem-solving powers to quantum computers based on desire rather than rigorous analysis. This framing positions much quantum discourse as magical thinking rather than scientific reasoning.

5.3.4 Establishing the Methodology: The Race Framework

Troyer proposes a systematic comparison methodology:

For that, I now want to compare a classical system with a quantum system. I want to totally bias it towards the quantum system. I want to say, let me put in the race on the classical side a single today's classical chip, the most advanced NVIDIA chip with 54 billion transistors. (Troyer, 2021, 6:12)

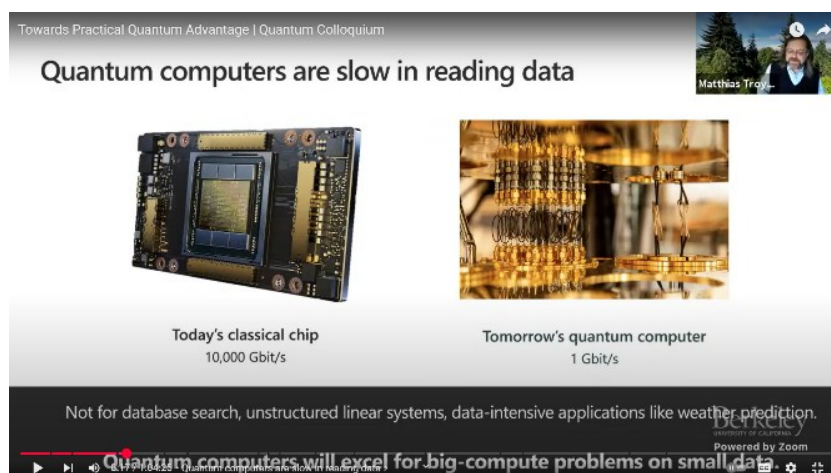


Figure 6: Troyer is putting on a race between classical and quantum computers

The "race" metaphor frames quantum versus classical computing as competition, with Troyer as referee establishing rules. The phrase "totally bias it towards the quantum system" is rhetorically

How Quantum Computing Escapes the Maze

clever—by acknowledging bias, Troyer suggests that even under optimistic assumptions favouring quantum computing, many applications remain unviable. If quantum computing can't win a biased race, it certainly can't win a fair one.

The specific comparison parameters perform important work:

- Classical side: "a single today's classical chip"—not a supercomputer or optimised cluster, just one consumer chip available now.
- Quantum side: "10,000 logical qubits"—a future quantum computer far beyond current capabilities, with "a logical cycle time of a few microseconds" and "all-to-all connectivity"—ideal conditions that don't currently exist.

This asymmetric comparison (current classical versus ideal future quantum) makes quantum's challenges more visible. If future quantum barely beats current classical, quantum advantage remains distant.

5.3.5 Debunking Big Data: The I/O Problem

Troyer's first major argument targets quantum computing for big data:

I want to start and debunk some hype. One thing one often hears is that quantum computers will solve the big data problem. (Troyer, 2021, 7:19)

He provides concrete numerical comparisons:

The classical chip can read 10 terabits a second. The quantum computer with 10,000 qubits at the clock speed of 2 to 10 microseconds can read about one gigabit a second. So that future quantum machine is about a factor of 10,000 slower for classical I/O than today's classical computer. (Troyer, 2021, 7:50)

This numerical comparison is devastating for big data claims. The orders-of-magnitude difference makes quantum computing for big data obviously impractical. Troyer reinforces this conclusion:

Thus, it is very clear to all, I think, just looking at that, it's not for big classical data. Quantum computers will excel on big compute problems on small data or maybe on problems on quantum data that we have to find. (Troyer, 2021, 8:12)

This establishes a crucial constraint: quantum advantage requires problems with small data but extensive computation. This eliminates most machine learning, data analytics, and similar applications dominating commercial computing.

5.3.6 Addressing QRAM: Dismissing Theoretical Solutions

Troyer anticipates a counterargument—quantum RAM (QRAM) might overcome I/O limitations. He explains QRAM's theoretical promise:

The idea of this QRAM oracle is, in a simplified way, you take a hard disk, and you can read all the data at once by putting the read head of the hard drive in a quantum superposition of all the locations on the disk. (Troyer, 2021, 9:00)

This vivid metaphor—a read head simultaneously positioned everywhere—captures QRAM's theoretical elegance. However, Troyer immediately dismisses practical viability:

It's a beautiful theoretical idea. But really implementing it will need quantum error correction, and then the scaling goes back to being proportional to the data size. (Troyer, 2021, 9:32)

This dismissal performs important boundary work. Troyer acknowledges QRAM as theoretically interesting whilst rejecting it as practical solution. The phrase "beautiful theoretical idea" patronises QRAM advocates—their idea is aesthetically pleasing but practically irrelevant.

5.3.7 The Constant Slowdown Problem

Troyer establishes quantum computing's fundamental speed disadvantage:

Just comparing a single classical chip with a future 10,000 logical qubit quantum machine, the quantum machine is 10 to 12 orders of magnitude slower than the classical one. (Troyer, 2021, 12:45)

This "constant slowdown" becomes central to Troyer's argument. Quantum computers must overcome enormous inherent disadvantage before achieving practical advantage. The numbers are stark:

- Logical operations: Classical performs ~5 petaflops (5×10^{15} operations) per second; quantum performs ~few million ($\sim 10^6$) operations per second. Quantum is $\sim 10^9$ times slower.
- Floating-point operations: Classical performs ~5 petaflops; quantum performs ~few thousand operations per second. Quantum is $\sim 10^{12}$ times slower.

These comparisons demonstrate that quantum speedup must be extraordinary—not merely polynomial but exponential—to overcome inherent slowdown.

5.3.8 Eliminating Quadratic Speedups: Grover's Algorithm Dismissed

Troyer analyses whether quadratic speedups (like Grover's algorithm) suffice:

If the constant slowdown is about 10 to the 12, then I have to make 10 to the 12 fewer calls. That means with quadratic speedup, I start winning if quantumly I make 10 to the 12 calls versus classically 10 to the 24. (Troyer, 2021, 13:43)

This calculation shows quadratic speedup requires problems so large (10^{24} classical operations) that even at petaflop speeds, classical computers need "years or months at least." Such problems rarely occur in practice.

Troyer concludes decisively:

Really, we want exponential speedups; that's where we have the best chance. (Troyer, 2021, 15:18)

This conclusion eliminates vast categories of proposed quantum applications. Machine learning (often relying on Grover-type speedups), optimisation problems, database searching, and many other applications become non-viable.

5.3.9 Clearing the Maze: Dropping Out Applications

Troyer systematically eliminates proposed applications:

From that whole list of problems that one often hears, I like to drop out all of the problems that are big data problems, machine learning problems, weather prediction, stock market predictions. (Troyer, 2021, 15:28)

This "dropping out" language is significant—Troyer acts as arbiter determining which applications deserve consideration. The passive voice ("problems that one often hears") distances these claims from specific advocates, framing them as general misconceptions rather than competing visions.

He continues eliminating applications:

This is all tied to the data, and the proposals for problems like protein folding and drug design are also often tied to Grover's speedup. So those I'd like to drop out. (Troyer, 2021, 15:40)

The phrase "I'd like to drop out" reveals this as rhetorical performance—Troyer enacts authority to dismiss applications. The casual tone ("I'd like to") masks the significant implications—entire research programmes are declared unviable.

5.3.10 The Leaf in Spring: Quantum Simulation as Viable Path

Having eliminated numerous applications, Troyer identifies where quantum advantage might exist:

One example is a leaf in spring. When you look outside in the spring and the leaves come up, you see the green everywhere. Nature knows the leaf is green, but calculating the colour of chlorophyll is hard classically. Why is that? It is because nature is quantum and classical computers are not. (Troyer, 2021, 18:03)




Figure 7: anthropomorphic of nature

This metaphor performs crucial work. The leaf—ordinary, visible, unproblematic in nature—becomes computationally intractable for classical computers. This dramatises classical computing's limitations whilst suggesting quantum computers might naturally solve such problems.

The anthropomorphisation is striking: "Nature knows the leaf is green." This personification attributes knowledge to nature itself, suggesting nature operates according to quantum principles that quantum computers might replicate. The metaphor constructs quantum computing not as artificial technology but as aligning with nature's fundamental operations.

Troyer invokes Feynman's famous argument (Feynman, 1982):

That's why Feynman, decades back, said to simulate quantum physics, to simulate quantum problems, we need quantum hardware. (Troyer, 2021, 18:31)

This appeal to authority (Feynman—Nobel laureate and respected physicist) legitimises quantum simulation as quantum computing's primary application. If Feynman identified this decades ago, pursuing other applications ignores foundational insight.

5.3.11 From Analogue to Digital: The Scalability Argument

Troyer distinguishes between analogue quantum simulators (already demonstrated) and digital quantum computers (future):

Just like classical machines that are analogue, it is limited... But digitally, we can. The same has to happen quantumly. (Troyer, 2021, 24:25)

This analogy to classical computing's history—from analogue devices to digital computers—suggests quantum computing must follow similar trajectory. Analogue quantum simulators are interesting but limited (like analogue classical computers). Only digital quantum computers with error correction can scale to solving practically important problems.

This framing justifies Microsoft's focus on fault-tolerant quantum computing rather than NISQ devices. Current demonstrations are analogous to analogue computers—interesting proofs of concept but ultimately dead ends. True quantum advantage requires the quantum equivalent of digital computing's revolution.

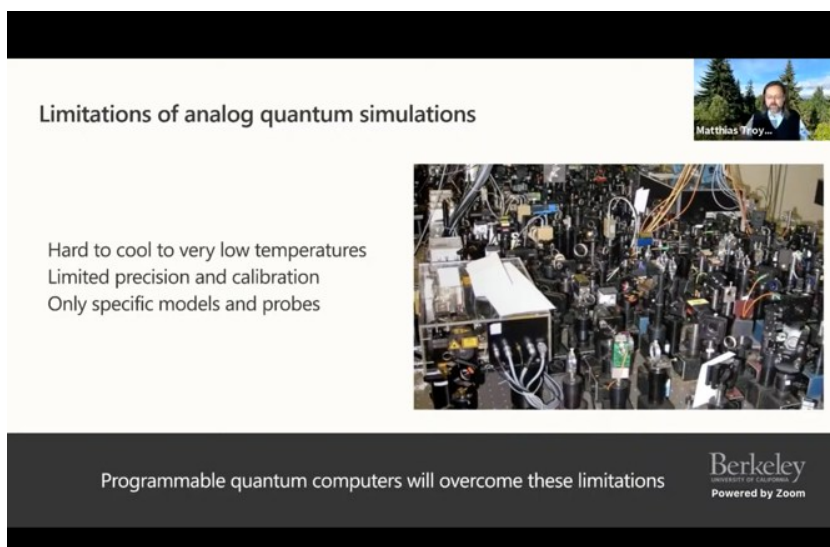


Figure 8: moving from analogue quantum simulations to quantum computing

5.3.12 Chemical Reactions: Barrier and Pathway Metaphors

Troyer explains quantum chemistry applications using vivid spatial metaphors:

As they bind and react, there's a barrier, and the rate for the reaction is exponential in the energy height of the barrier divided by the temperature. In order to know in which direction things react, I have to look at all possible paths, calculate the energies of these barriers, and then find out which barrier is lowest. (Troyer, 2021, 33:21)

The "barrier and pathway" metaphor constructs chemical reactions as navigation problems. Molecules must overcome energy thresholds (activation energies) to transform. Calculating these barriers classically is computationally intractable—testing all possible pathways requires examining exponentially growing state spaces.

Troyer provides dramatic temporal contrast:

Calculating chemical reactions using classical computers would take a billion years, which is far longer than the timeframe of "a few weeks" that researchers aim for. With quantum computers, we estimate that about 5,000 logical qubits could reduce this to roughly a month. (Troyer, 2021, 40:37)

The "billion years versus roughly a month" comparison dramatizes quantum advantage. This isn't marginal improvement but revolutionary transformation—making previously impossible calculations practical.

5.3.13 Quantum Tunnelling: Optimisation Metaphor

When discussing optimisation problems, Troyer employs quantum tunnelling as metaphor:

Instead of climbing over the barrier, one can just quantum tunnel through it. (Troyer, 2021, 48:04)

This metaphor captures quantum mechanics' counterintuitive feature—particles can pass through energy barriers that would be insurmountable classically. Applied to optimisation, quantum tunnelling suggests quantum computers might escape local minima that trap classical algorithms.

However, Troyer quickly complicates this hopeful metaphor:

What we realised over the last decade, looking at quantum annealers, is that you can mimic it classically. You can simulate it efficiently classically, and thus you can mimic the quantum tunnelling on a classical computer. (Troyer, 2021, 48:56)

This revelation—that quantum tunnelling's advantage can be captured by quantum-inspired classical algorithms—demonstrates the moving target problem. What initially appeared as quantum advantage becomes achievable classically through algorithmic innovation. This sobering realisation undermines optimisation as quantum application whilst suggesting quantum computing research benefits classical computing even before quantum hardware succeeds.

5.3.14 Temporal Framing: Multiple Future Horizons

Throughout the colloquium, Troyer employs sophisticated temporal framing, articulating multiple future horizons:

- Near-term (few years): "With 100,000 physical qubits and fault-tolerant simulations, we can solve quantum problems that are classically hard. That's not so far out; I think we can achieve that in the next few years." (30:31)
- Medium-term (next decades): "To really get quantum advantage for a practical problem in the next decades, we need to focus on a problem that has small data but big compute problems." (51:06)
- Long-term (millions of qubits): "Even the simplest application problems will need fault tolerance and hundreds of thousands of qubits to beat the best classical algorithms. Once we get to millions, there are loads of chemistry problems that will solve interesting ones like carbon fixation and more." (51:25)

These staggered timeframes manage expectations whilst maintaining enthusiasm. Near-term achievements (100,000 qubits) appear achievable, justifying continued investment. Medium and long-term visions suggest increasingly significant impacts, maintaining transformative promise without specifying dates enabling falsification.

5.3.15 Addressing NISQ: Dismissing Competitors' Approaches

Troyer explicitly addresses NISQ devices—the approach pursued by many competitors:

NISQ machines don't currently have any roadmap towards that; we need fault tolerance.
(Troyer, 2021, 58:50)

This dismissal is significant. NISQ devices represent billions of dollars of investment by IBM, Google, Rigetti, IonQ, and others. Troyer declares these efforts lack "roadmap towards" practical quantum advantage, positioning Microsoft's fault-tolerant approach as uniquely viable.

He elaborates during the question session:

If you're using an analogue quantum system and now you insist that you only have the discrete gate operations of a digital computer, then you combine the disadvantages of analogue being noisy and digital being restricted. (Troyer, 2021, 59:59)

This characterisation—NISQ combines both approaches' disadvantages whilst capturing neither's advantages—delegitimises competitors' strategies. The phrase "you are handicapping the device" suggests NISQ advocates foolishly constrain their hardware.

Troyer offers alternative framing:

What you should actually do is take the analogue machine and use it as an analogue machine, and then you can do a lot. (Troyer, 2021, 1:01:01)

This suggests NISQ devices should abandon digital quantum computing ambitions and accept status as analogue simulators—useful but limited. This reframes competitors' achievements as interesting physics experiments rather than progress towards universal quantum computing.

5.3.16 Benefits Along the Journey: Classical Computing Improvements

Troyer acknowledges that pursuing quantum computing advances classical methods:

Even on the way there, in the decade it takes us to get there, there will be breakthroughs in classical methods just from tackling the classical chemistry problems we need to solve to make the quantum hardware useful. (Troyer, 2021, 52:28)

This framing performs important work. It suggests quantum computing research justifies itself even if quantum hardware never achieves practical advantage. The journey towards quantum computing—developing better algorithms, understanding quantum problems more deeply, creating quantum-inspired classical methods—produces valuable knowledge regardless of hardware success.

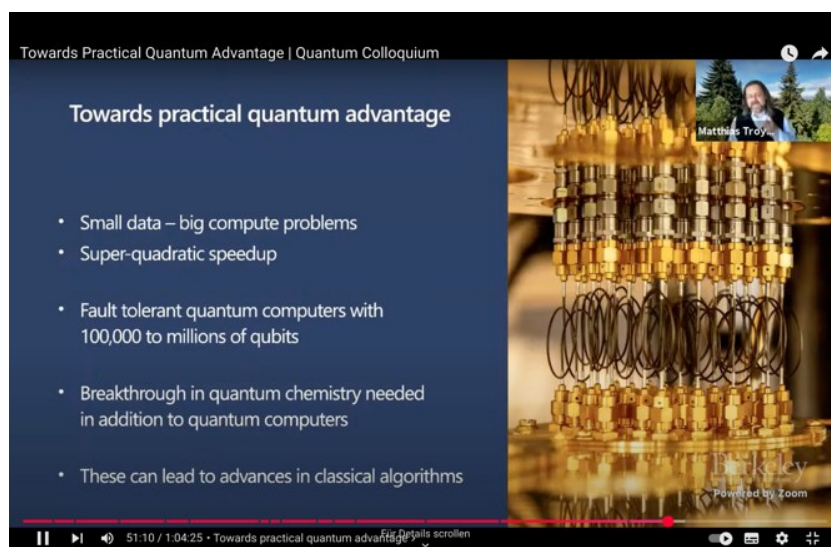


Figure 9: Troyer summarizes prerequisites for quantum advantage

How Quantum Computing Escapes the Maze

This protects against criticism that quantum computing investment is wasted. Even if quantum computers ultimately prove impractical, the research programme yields benefits. This framing echoes space programme defences—whether or not humans colonise Mars, space research produces valuable technologies and knowledge.

5.4 STS INTERPRETATION: ESCAPING THE QUANTUM WISHING WELL

Troyer’s 2021 colloquium marks a fundamental shift in Microsoft’s quantum narrative — not merely an update but a strategic recalibration of the entire vanguard position. Where Warnick’s 2014 report constructed authority through charismatic figuration and expansive imaginaries, Troyer constructs authority through systematic technical analysis and deliberate imaginary constraint. This shift responds to changed circumstances: the post-hype environment following Google’s 2019 quantum supremacy announcement (Arute et al., 2019), accumulated expectations that had not been met, and competitive pressure from multiple corporations now demonstrating quantum capabilities. Three analytical dimensions illuminate how this recalibration operates. The first concerns how the colloquium reconstructs authority and expertise. The second concerns how imaginaries are strategically narrowed and expectations managed. The third concerns the performative mechanisms through which these moves constitute a new kind of vanguard legitimacy.

Step 4:
Theoretical
Inter-
pretation

5.4.1 Reconstructing Authority Through Expert Assessment

The shift from charismatic to expert authority is the colloquium’s most significant legitimisation move. Troyer’s authority derives not from extraordinary personal qualities — though his ETH Zürich credentials and distinguished scientist designation are invoked — but from systematic technical analysis, comparative benchmarking, and apparent methodological rigour. His self-presentation as “systems architect” emphasises professional expertise and organisational role rather than individual brilliance. The colloquium presents detailed calculations, comparative visualisations of classical versus quantum computational capability, and logical argumentation through which credibility is constructed through apparent objectivity rather than charismatic claim-making. Following Latour and Woolgar’s (1986/1979) framework, the inscriptions Troyer deploys — graphs, numerical comparisons, computational benchmarks — function as argumentative resources rather than reports of experimental results, transforming strategic choices about what to compare into apparently neutral technical assessments.

The colloquium's academic setting performs important additional legitimization work. Delivering the presentation at UC Berkeley's Simons Institute positions Troyer within academic rather than corporate discourse, importing the credibility of peer evaluation and scientific independence. His hype critique gains force from this positioning: Microsoft's assessment of quantum computing's limitations appears more credible when delivered in an academic venue than when issued as corporate communication. This is a sophisticated deployment of what Collins (1992/1985) terms core-set membership: Troyer speaks as scientist who happens to work for Microsoft rather than as Microsoft representative promoting products. Yet the colloquium involves significant strategic choices that this positioning obscures — which comparison parameters to employ, which applications to dismiss, which assumptions underlie the billion-year versus one-month runtime calculations. Narrative analysis can identify these choices; it cannot determine whether they reflect accurate technical assessment, which is precisely the epistemological condition that SQ4 investigates.

The boundary work the colloquium performs is more aggressive than the 2014 report's disciplinary negotiations. Troyer's explicit contrast between "practical quantum advantage" and the "quantum wishing well" demarcates rigorous scientific assessment from marketing speculation — a boundary that positions the colloquium itself on the legitimate side. The "wishing well" metaphor is particularly effective: wishing wells are places where desires are projected without rational basis, suggesting that much quantum discourse involves magical rather than scientific thinking. By positioning himself as debunking hype, Troyer claims the authority of the honest broker. The dismissal of random circuit demonstrations as serving "no purpose beyond demonstrating quantum capability" redirects the boundary between meaningful and meaningless achievement, favouring application-oriented metrics over the supremacy demonstrations that competitors had pursued. This reframing delegitimises Google's 2019 achievement not by contesting its technical validity but by contesting its relevance.

5.4.2 Managing Expectations and Constraining the Imaginary

The colloquium's most visible analytical move is the systematic elimination of quantum computing applications — what Troyer terms "dropping out" and what this thesis analyses as strategic imaginary constraint. Drawing on Borup et al.'s (2006) sociology of expectations, the colloquium performs retrospective adjustment: applications previously presented as promising possibilities — big data analysis, machine learning, weather prediction, protein folding — are reclassified as "quantum wishing well" thinking. Each dismissal is anchored in

How Quantum Computing Escapes the Maze

technical arguments about I/O constraints, constant slowdown problems, or quadratic rather than exponential speedups. These arguments construct the eliminations as objective technical assessment rather than strategic choice, transforming disappointment into learning opportunity: Microsoft appears not as having failed to deliver promised applications but as now understanding quantum computing more precisely.

The constraint responds directly to the hype cycle dynamics that Bakker and Budde (2012) identify. By 2021, the gap between quantum computing's promises and its demonstrated capabilities had become visible enough to damage credibility. Troyer's solution is to narrow the imaginary to a domain — quantum chemistry and materials simulation — where Microsoft's fault-tolerant topological approach is precisely optimised. This narrowing is, as Jasanoff and Kim (2009, 2015) would predict, responsive to institutional circumstances rather than purely to technical reassessment: the imaginary contracts to what the current approach can credibly promise. The constraint also performs strategic competitive positioning. Troyer's distinction between digital and analogue quantum approaches — arguing that “the same has to happen quantumly” as classical computing's transition from analogue to digital — delegitimises NISQ devices, which competitors had been demonstrating, by framing them as historically limited precursors rather than viable paths to quantum advantage.

The frontier stage's temporal vagueness yields to staggered milestones: “next few years” for 100,000 qubits, “next decades” for practical applications, eventual achievement of millions of qubits for broad chemistry problems. This specification follows Selin's (2008) analysis of future-making: it appears more credible than vague revolutionary promises by establishing concrete checkpoints, manages expectations by distinguishing near-term from long-term possibilities, and preserves enthusiasm by ensuring that near-term achievements remain apparently reachable. Yet the staggered structure also creates vulnerabilities. Specific milestones generate falsifiable commitments; the “next few years” for 100,000 qubits, if unmet, becomes evidence against rather than for the approach. The pragmatic navigation stage's credibility-building mechanism — transparency about requirements — simultaneously increases exposure to future accountability.

5.4.3 Performativity, Narrative Dependency, and the Absent Laboratory

Drawing on Austin's (1962) speech act theory as developed by Callon (1998) and MacKenzie (2006), the colloquium performs a new kind of quantum computing legitimacy into existence rather than merely reporting pre-existing conditions. When Troyer declares certain applications belong to “the quantum wishing well,” the declaration does not await community

consensus — it attempts to constitute their illegitimacy through uptake. MacKenzie's (2006) concept of Barnesian performativity is instructive here: if researchers redirect efforts away from the dismissed applications, if investors grow sceptical of big data quantum promises, if policy documents begin distinguishing “hype” from “practical advantage,” then Troyer's declarations become more true as more actors treat them as true. The systematic dismissal of NISQ approaches — combining “the disadvantages of analogue being noisy and digital being restricted” — similarly attempts to perform roadmap absence into existence: if no roadmap is perceived, investment follows other paths regardless of NISQ's actual potential.

The maze metaphors that pervade the colloquium function as both epistemic and strategic resources, following Knorr Cetina's (1995) analysis of metaphors in scientific practice. The maze constructs the quantum computing landscape as complex, confusing, and requiring expert guidance — principles that Troyer claims to understand and that position him as the guide who has explored the paths, identified the dead ends, and knows the viable routes. The “race” between classical and quantum computing naturalises a competitive framing that obscures possibilities of complementary development or application-specific advantage. The anthropomorphisation that “Nature knows the leaf is green” performs ontological legitimation — quantum computing as aligned with reality's fundamental structure rather than as artificial technology. These metaphors collectively construct the pragmatic navigation stage's characteristic authority: not the frontier's inspirational vision but the systematic expert assessment of a terrain that only specialists can accurately map.

The colloquium's deepening of narrative dependency deserves particular attention. Unlike Warnick's 2014 report, which at least showed researchers at chalkboards and described laboratory environments, Troyer's presentation contains no laboratory imagery and little discussion of Microsoft's actual hardware progress. Quantum hardware remains black-boxed: qubit numbers and error rates are discussed abstractly without explaining the physical mechanisms involved. This absence is analytically significant. The colloquium performs expert knowledge without demonstrating material capability — audiences must trust Troyer's calculations and comparisons without independent means of verification. Under conditions of narrative dependency, external observers cannot distinguish accurate technical assessment from strategically motivated framing. This inability is not merely a scholarly inconvenience. It is a political condition: the authority to define which quantum applications are viable and which belong to “the wishing well” — an authority that shapes investment flows, research priorities, and policy frameworks — rests exclusively with insiders whose accounts cannot be independently verified. The maze, as Troyer presents it, can only be navigated by

How Quantum Computing Escapes the Maze

following the guide's directions; and the guide's institutional interests are inseparable from the map he provides.

5.5 SYNTHESIS: THE PRAGMATIC TURN

Step 5: Brief
Synthesis

Troyer's 2021 colloquium represents what I term the "pragmatic turn" in Microsoft's quantum computing narrative—a fundamental reconfiguration of vanguard strategy responding to changed circumstances, accumulated expectations, and competitive pressures. The colloquium navigates the post-hype environment through systematic constraint of imaginaries, recalibration of authority bases, and implicit competitive positioning masked by academic framing.

Key Findings

The colloquium demonstrates technological readiness being narratively reconstituted as circumstances change. Rather than celebrating what quantum computing might achieve, Troyer specifies what quantum computing *requires*: fault tolerance, 100,000 to millions of qubits, microsecond cycle times, exponential rather than quadratic speedups. This requirements-based framing constructs readiness as engineering challenge with identifiable milestones rather than speculative frontier with uncertain horizons.

Authority shifts from charismatic genius to expert assessment. Troyer's systematic technical analysis, comparative benchmarking, and methodological rigour construct credibility through apparent objectivity. His academic positioning—ETH Zürich professor delivering colloquium at UC Berkeley's Simons Institute—imports academic legitimacy into corporate communication, enabling hype critique that enhances credibility whilst strategically targeting competitors' approaches.

The sociotechnical imaginary undergoes dramatic contraction. The systematic "dropping out" of applications—big data, machine learning, weather prediction, protein folding—narrows quantum computing's viable scope to quantum chemistry and materials science. This constraint responds to hype cycle dynamics: expansive promises proving unsustainable require narrowing to defensible domains. The contraction positions Microsoft's fault-tolerant approach as addressing quantum computing's actual possibilities.

Temporal framing shifts from protective vagueness to staggered milestones: "next few years" for 100,000 qubits, "next decades" for practical applications, eventual millions of qubits for broad

chemistry problems. This specification appears more credible whilst creating flexibility—if near-term goals aren't achieved, medium and long-term visions remain intact.

Boundary work intensifies and redirects. Troyer's contrast between "practical quantum advantage" and the "quantum wishing well" demarcates rigorous assessment from marketing speculation. The dismissal of NISQ devices as combining "the disadvantages of analogue being noisy and digital being restricted" delegitimises competitors' approaches whilst maintaining academic politeness.

The Pragmatic Turn's Functions

The pragmatic turn performs multiple functions: restoring credibility by acknowledging limitations whilst preserving promise within narrower bounds; differentiating Microsoft's approach by delegitimising competitors' achievements as dead ends; recalibrating expectations through explicit promise-requirement cycles; justifying continued investment in Microsoft's longer-timeline approach as pursuing the only viable path; and appealing to scientific talent who prefer rigorous analysis over marketing enthusiasm.

Tensions and Transition Pressures

The pragmatic turn creates new vulnerabilities. The dramatic shift from frontier enthusiasm creates narrative discontinuity—applications dismissed as "quantum wishing well" had appeared in earlier communications as promising possibilities. Specific milestones ("next few years") create falsifiable checkpoints. The aggressive competitive positioning invites counter-claims if competitors achieve practical quantum advantage through dismissed approaches.

These tensions generate pressures toward breakthrough claims. Timeline commitments intensify demonstration requirements. Competitor achievements accumulate, challenging Microsoft's positioning. Commercial stakeholders expect movement toward deployable products. These pressures drive the transition from pragmatic navigation to breakthrough achievement evident in the 2025 Majorana 1 announcement.

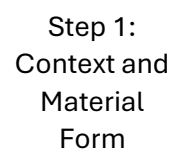
These pressures set the conditions for the final document in this analysis. The maze has been navigated as far as systematic expert assessment can take it; now a different narrative move is required. Chapter 6 analyses how Microsoft's 2025 Majorana 1 announcement responds to precisely these pressures — not by refining the map, but by claiming to have found the exit.

6 CLAIMING BREAKTHROUGH: MAJORANA 1 (2025)

The pressures identified at the close of Chapter 5 — falsifiable milestone commitments, accumulated competitor achievements, and commercial stakeholder expectations — culminate in the document analysed here. The 2025 Majorana 1 announcement represents Microsoft's most consequential quantum narrative act: a declaration that the maze has been solved, that the long-sought path has been found. This chapter examines how that claim is constructed, what discursive strategies give it force, and why its contested reception in the scientific community reveals the limits of narrative authority. It primarily addresses SQ1 (how the breakthrough declaration represents a further shift in Microsoft's narrative), SQ2 (how performative declarations construct and contest authority), SQ3 (how the futures projected for quantum computing re-expand in relation to Microsoft's strategic positioning), and SQ4 (the epistemological and political stakes of narrative dependency at the moment of claimed breakthrough).

6.1 CONTEXT AND MATERIAL FORM

Approximately eleven years after Jennifer Warnick's ethnographic portrait of Station Q, and four years after Matthias Troyer's pragmatic colloquium, Microsoft's quantum computing narrative underwent its most significant transformation. Where Warnick's 2014 report emphasised pioneering scientific exploration and Troyer's 2021 presentation carefully navigated between hype and possibility, Microsoft's February 2025 announcement of the Majorana 1 chip signals a decisive shift towards claims of practical realisation and commercial readiness.



Step 1:
Context and
Material
Form

Publication Context: The announcement, authored by Catherine Bolgar and published on 19 February 2025 on Microsoft's "Source" platform, represents the company's most significant quantum computing communication since Station Q's establishment. The timing is critical—occurring precisely as industry predictions suggest quantum computers will begin leaving laboratories for real-world deployment (Swayne, 2024). The announcement coincides with publication of peer-reviewed research in *Nature* (Aghaee et al., 2025), lending scientific credibility to Microsoft's claims.

Material Presentation: The webpage employs sophisticated visual rhetoric. It opens with a striking hero image showing a close-up of the Majorana 1 quantum chip against brass laboratory equipment. The photograph functions as what Latour (1986) would term an "inscription device"—transforming abstract quantum phenomena into visible, tangible objects that can be circulated

and scrutinised. The image establishes material credibility through visual evidence of the technology's physical existence.




Figure 10: A hero image of the Majorana 1 quantum chip (photo by John Brecher for Microsoft)

The document structure follows clear hierarchy: a bold headline proclaiming Microsoft has "carved a new path," followed by introductory paragraphs establishing the breakthrough's significance, then sections featuring direct quotes from three key Microsoft figureheads—Chetan Nayak, Matthias Troyer, and Krysta Svore—interspersed with technical explanations and application scenarios. Three portrait photographs punctuate the text, showing each researcher in laboratory settings, positioned as authoritative voices who embody Microsoft's quantum computing vision.




Figure 11: The Palm of the Hand—Visualising Quantum Scalability (photo by John Brecher for Microsoft)

Notably, the document features a photograph that will become central to Microsoft's rhetorical strategy: an image of the Majorana 1 chip held in a human hand. This "palm of the hand" image

How Quantum Computing Escapes the Maze

performs crucial work visualising quantum computing's scalability and accessibility, which I analyse in detail below.

Target Audience: The document addresses multiple audiences simultaneously:

- Technology journalists seeking breakthrough stories
- Investors evaluating Microsoft's quantum computing viability
- Potential customers considering quantum computing adoption
- Policymakers making decisions about quantum technology support
- Scientific community assessing Microsoft's technical claims
- General public interested in technological innovation

This multiplicity requires balancing technical credibility with accessibility, managing enthusiasm whilst appearing realistic, and addressing specialists' concerns whilst engaging broader audiences.

6.2 CRITICAL ASSESSMENT OF ISSUES

Following Asdal and Reinertsen's (2022) methodology, I critically assess what this document makes present and absent:

Step 2:
Critical
Assessment
of Issues

Issue 1: Gap Between Peer-Reviewed Claims and Commercial Promises

The announcement references "a new paper published Wednesday in Nature" as validation. However, examining the relationship between peer-reviewed findings and commercial claims reveals significant gaps.

[+] Strategic Presence: The Nature publication (Aghaee et al., 2025) provides scientific legitimacy, suggesting independent validation of Microsoft's work. The phrase "peer-reviewed confirmation" positions Microsoft's claims within established scientific authority structures.

[-] Strategic Absence: The peer-reviewed paper validates measurement capabilities and observation of Majorana signatures—it does not validate the commercial viability claims or scalability assertions dominating the announcement. By positioning the Nature publication at the narrative centre, the document imports scientific authority into commercial and engineering claims extending far beyond what peer review actually validated.

Issue 2: Controversial Scientific Status of Majorana Observations

The announcement treats Majorana particle observation as settled fact. However, ongoing scientific controversy surrounds these claims. Nature published an article noting that whilst

Microsoft claims quantum-computing breakthrough, "some physicists are sceptical" (Castelvecchi, 2025). The controversy involves:

- Questions about whether observed signals definitively indicate Majorana zero modes
- Alternative explanations for experimental observations
- Reproducibility concerns
- Interpretation debates within the physics community

[+] Strategic Presence: The announcement presents Majorana observation as accomplished achievement, establishing foundation for scalability claims.

[-] Strategic Absence: The document does not acknowledge ongoing scientific debates or alternative interpretations of experimental data. This silence transforms contested claims into settled facts, potentially overstating scientific consensus.

Issue 3: The Eight Qubit to One Million Qubit Gap

The announcement repeatedly emphasises the "path to a million qubits." However, the current Majorana 1 chip contains eight qubits. The document must bridge this six-orders-of-magnitude gap between current reality and promised future.

[+] Strategic Presence: The document articulates a "clear path" and describes scalability as solved problem: "Majorana 1 demonstrates a path to scale up to the million qubits needed to achieve commercial quantum advantage."

[-] Strategic Absence: The document provides minimal technical detail about specific challenges remaining. The phrase "a path to scale" suggests solved problem rather than monumental engineering challenge requiring overcoming numerous obstacles. The document does not specify timelines for moving from eight to millions of qubits, protecting against falsification whilst maintaining imminent breakthrough rhetoric.

Issue 4: Lack of Independent Verification

All quoted sources are Microsoft employees—technical fellows Chetan Nayak, Matthias Troyer, and Krysta Svore. No independent scientists, external validators, or critical voices appear.

[+] Strategic Presence: Microsoft's technical fellows speak with authority about their work, establishing expertise and credibility.

[-] Strategic Absence: Independent scientific assessment remains absent, creating what Collins (1992/1985) would recognise as a "core-set" problem. The boundary between insider expertise and independent assessment becomes blurred when only corporate voices define technological

How Quantum Computing Escapes the Maze

readiness. External perspectives might provide more cautious assessments of scalability challenges, timelines, and commercial viability.

Issue 5: Temporal Framing: "Years, Not Decades"

The announcement's most striking temporal claim appears twice: quantum computers capable of solving "meaningful, industrial-scale problems in years, not decades."

[+] Strategic Presence: This framing generates urgency and excitement. "Years" suggests imminent commercialisation justifying immediate investment and attention. The contrast with "decades" implicitly criticises competitors pursuing slower approaches.

[-] Strategic Absence: "Years" remains deliberately imprecise. Two years? Five years? Nine years? This imprecision provides flexibility—if commercial quantum computing arrives in 2030 (five years), Microsoft's prediction appears accurate; if it arrives in 2033 (eight years), "years, not decades" still technically holds. The framing generates enthusiasm whilst avoiding specific falsifiable predictions.

Issue 6: Applications Catalogue Without Implementation Detail

The document presents compelling application scenarios:


- Self-healing materials repairing "cracks in bridges or airplane parts, shattered phone screens or scratched car doors"
- Breaking down microplastics or "tackling carbon pollution"
- Calculating enzyme behaviour to "eradicate global hunger" through boosting soil fertility

[+] Strategic Presence: These applications connect quantum computing to recognised societal challenges (infrastructure decay, plastic pollution, food security), enrolling diverse stakeholder communities in quantum computing's promise.

[-] Strategic Absence: The document provides no technical detail about how quantum computers would actually enable these applications. The leap from "calculating molecular properties accurately" to "self-healing phone screens" or "eradicating global hunger" involves numerous unstated assumptions and implementation challenges. The applications function rhetorically rather than technically—they construct desirable futures without specifying pathways to achievement.

6.3 NARRATIVE ANALYSIS: FROM MAZE TO MILESTONE

The Majorana 1 announcement employs sophisticated narrative architecture transforming Microsoft's quantum story from exploration (2014) and navigation (2021) to achievement (2025). This section analyses the document's narrative construction.



6.3.1 Opening: The Revolutionary Invention

The Semiconductor Analogy

The announcement's opening sentence establishes the document's master metaphor:

In the same way that the invention of semiconductors made today's smartphones, computers and electronics possible, topoconductors and the new type of chip they enable offer a path to developing quantum systems. (Bolgar, 2025)

This analogy performs extraordinary rhetorical work. By comparing topoconductors to semiconductors, the document:

- Establishes Historical Parallel: Semiconductors enabled the digital revolution, transforming twentieth-century society. The analogy suggests topoconductors will similarly revolutionise twenty-first-century computing.
- Claims Fundamental Innovation: The comparison is not to incremental improvement (a better processor) but to entirely new material category enabling new technological paradigm.
- Implies Inevitability: Semiconductors' success wasn't apparent initially but retrospectively appears inevitable given their fundamental advantages. The analogy suggests topoconductors' success will similarly prove inevitable.
- Positions Microsoft as Inventor: Just as specific companies (Fairchild Semiconductor, Intel) captured value from inventing and commercialising semiconductors, Microsoft positions itself as inventor-commercialiser of topoconductors.

The Transistor Metaphor

Chetan Nayak's quote reinforces the semiconductor analogy:

We took a step back and said 'OK, let's invent the transistor for the quantum age.' (Bolgar, 2025)

This metaphor is even more audacious than the semiconductor analogy. The transistor is arguably the twentieth century's most important invention. By claiming to have "invented the transistor for the quantum age," Nayak positions Microsoft's achievement as civilisational-scale innovation. The phrase "took a step back" suggests strategic wisdom—rather than pursuing incremental improvements, Microsoft identified fundamental requirements and invented the essential foundations. This narrative portrays Microsoft as visionary whilst implicitly criticising competitors pursuing less fundamental approaches.

The Million Qubit Threshold

The opening section establishes "a million qubits" as critical benchmark:

Whatever you're doing in the quantum space needs to have a path to a million qubits. If it doesn't, you're going to hit a wall. (Bolgar, 2025)

This statement performs multiple functions:

- Establishes Success Criterion: The million-qubit threshold becomes the standard against which all quantum computing approaches are measured.
- Delegitimises Alternatives: Approaches lacking million-qubit scalability are framed as inevitably failing ("hit a wall"), regardless of near-term achievements.
- Positions Microsoft as Unique: The document constructs Microsoft's topological approach as the only pathway to millions of qubits, suggesting competitors will inevitably encounter insurmountable obstacles.

6.3.2 Part Two: Scientific Vindication

The Risk-Reward Narrative

The document constructs a narrative of vindication for risky choices:

Building a quantum computer has always been challenging. That difficulty led many companies to pursue approaches that seemed simpler to realize but ultimately weren't

able to scale to the millions of qubits required for commercial quantum computing. By contrast, Microsoft took a different path that, although considered high risk, high reward, is now paying off. (Bolgar, 2025)

This passage constructs quantum computing's history as competition between approaches. Competitors pursued "simpler" paths that "ultimately weren't able to scale," whilst Microsoft pursued "different path" that is "now paying off." The narrative transforms Microsoft's slower progress into strategic superiority—the tortoise vindicated against the hare.

The phrase "high risk, high reward" acknowledges previous uncertainty whilst asserting current success. Microsoft's approach appeared risky but has proven correct, vindicating long-term investment.

The Topoconductor: New State of Matter

The announcement introduces "topoconductor" as fundamental innovation:

The topoconductor creates an entirely new state of matter—not a solid, liquid or gas but a topological state. (Bolgar, 2025)

This claim is remarkable. States of matter are fundamental physical categories. By claiming to create "entirely new state of matter," the document positions Microsoft's work as fundamental physics rather than incremental engineering.

The phrase "not a solid, liquid or gas but a topological state" employs parallel structure emphasising novelty. Topological states join fundamental physical categories, suggesting Microsoft has made discovery-level contribution to physics itself.

DARPA Validation

The document leverages governmental endorsement:

Microsoft was one of two companies to be invited to move to the final phase of the program, known as Underexplored Systems for Utility-Scale Quantum Computing (US2QC), which is part of the U.S. Defense Advanced Research Projects Agency (DARPA). (Bolgar, 2025)

DARPA validation provides crucial external legitimacy. DARPA funded Internet development, GPS, and numerous transformative technologies. Being "one of two companies" selected suggests rigorous evaluation and Microsoft's approach surviving competitive assessment.

How Quantum Computing Escapes the Maze

The connection to national security interests (DARPA is Department of Defense agency) adds gravitas. Quantum computing isn't merely commercial technology but strategic national capability, justifying massive investment and rapid development.

Digital Control: The Simplification Narrative

The document emphasises architectural innovation:

Prior approaches to creating topological qubits required analogue "finetuning" at each step in a computation—making them very difficult to scale. By contrast, Majorana 1 proves that topological qubits can be controlled digitally, redefining and vastly simplifying how quantum computing works. (Bolgar, 2025)

This passage constructs Microsoft's innovation as "simplification"—solving complexity problems that plagued earlier approaches. The contrast between "analogue finetuning" (complex, difficult) and "digital control" (simple, scalable) echoes Troyer's 2021 argument about digital superiority over analogue approaches.

The phrase "redefining and vastly simplifying how quantum computing works" makes extraordinary claim—Microsoft hasn't merely improved quantum computing but fundamentally redefined it. This positions competing approaches as obsolete, operating under old paradigm that Microsoft has superseded.

6.3.3 Part Three: Transformative Applications

The Applications Catalogue

The document's final section articulates sociotechnical imaginaries through specific application scenarios. Each application addresses recognised societal challenge:

- Self-Healing Materials, repairing shattered phone screens or even cracks in bridges or airplane parts or scratched car doors. This application spans scales from infrastructure (bridges, airplanes) to consumer products (phones, cars), suggesting quantum computing will impact both public goods and everyday life. The specificity ("shattered phone screens") makes quantum computing's benefits tangible and relatable.
- Environmental Applications, breaking down microplastics or tackling carbon pollution. These applications connect quantum computing to urgent environmental challenges. Microplastics and carbon pollution are widely recognised problems lacking good solutions. Positioning quantum computing as enabling solutions enrolls environmental

constituencies whilst constructing quantum computing as contributing to planetary sustainability.

- Food Security, calculating enzyme behaviour to eradicate global hunger by boosting soil fertility. This is perhaps the most ambitious application claim. "Eradicate global hunger" is millennium-scale aspiration. By suggesting quantum computing might contribute to solving world hunger, the document constructs quantum computing as addressing humanity's most fundamental challenges.

Troyer's AI-Quantum Integration Vision

Matthias Troyer articulates the most utopian vision:

Any company that makes anything could just design it perfectly the first time out. It would just give you the answer. [...] The quantum computer teaches the AI the language of nature so the AI can just tell you the recipe for what you want to make

This formulation presents remarkably frictionless future where computational power eliminates trial-and-error. The phrase "just give you the answer" constructs quantum-AI systems as oracles that simply provide solutions, bypassing messy iterative development.

This vision echoes science fiction scenarios where superintelligent computers solve problems instantly. The casual language ("just design it perfectly," "just give you the answer") masks the extraordinary nature of these claims—perfect design without iteration would represent revolutionary transformation of engineering and innovation processes.

Maintaining Uncertainty While Claiming Transformation

Despite bold application claims, the document maintains careful hedging:

There's a good case to be made that quantum computing could change the world. (Bolgar, 2025)

The modal "could" acknowledges uncertainty whilst the phrase "change the world" maintains transformative rhetoric. This balance protects credibility—Microsoft acknowledges possibilities rather than guarantees—whilst maintaining enthusiasm generating investment and support.

6.4 STS INTERPRETATION: MATERIALISING THE FUTURE

The 2025 Majorana 1 announcement represents the third and most consequential shift in Microsoft's quantum narrative. Where the 2014 report constructed authority through charismatic figuration and the 2021 colloquium

Step 4:
Theoretical
Inter-
pretation

How Quantum Computing Escapes the Maze

through expert assessment, the 2025 announcement constructs authority through vindication narrative — the claim that a long-contested approach has been validated. This shift reorganises every dimension of the analytical framework: authority sources, imaginary scope, temporal framing, boundary work, and performative mechanisms all undergo simultaneous reconfiguration in response to the breakthrough stage's specific strategic requirements. Three analytical dimensions structure the interpretation that follows. The first concerns how breakthrough status is constituted through literary and authority-compounding mechanisms. The second concerns the strategic re-expansion of sociotechnical imaginaries accompanying breakthrough claims. The third concerns the performative dynamics of the announcement and the scientific controversy that followed it — a controversy that illuminates the limits of narrative authority under conditions of narrative dependency.

6.4.1 Constituting Breakthrough: Literary Technologies and Authority Accumulation

The announcement performs dramatic modalization work, moving claims along the continuum from contested speculation toward established fact. In 2014, Majorana particles were “compelling evidence” requiring careful qualification; in 2021, Troyer discussed quantum chemistry requirements abstractly; in 2025, the announcement states that topological qubits “are real” and that Microsoft has “carved a new path.” Drawing on Latour and Woolgar's (1986/1979) framework, this modalization reversal transforms contested experimental interpretations into the foundation for commercial promises. The Nature publication functions as crucial inscription device: by stating “a new paper published Wednesday in Nature marks peer-reviewed confirmation,” the announcement imports scientific authority into claims extending far beyond what peer review validated. The photographs function similarly — the hero image of the Majorana 1 chip against brass laboratory equipment transforms abstract quantum phenomena into visible, tangible objects that appear to warrant the announcement's claims.

The “palm of the hand” photograph warrants particular attention as an inscription device. The anonymous hand — no jewellery, no visible face — transforms an individual gesture into a universal human scale, inviting viewers to imagine themselves holding this artefact. The chip's visual ordinariness performs important work: quantum computing need not be exotic or incomprehensible; it can be held, examined, possessed. The photograph directly contradicts Troyer's 2021 warning that competing approaches require “an impractical computer the size of an airplane hangar,” condensing Microsoft's scalability argument into a single image. Yet what the photograph conceals is equally significant: the chip cannot function in a human

hand. It requires a dilution refrigerator cooler than outer space and elaborate control systems. The image stages a particular moment — chip at room temperature, inactive, displayable — that bears little resemblance to the operational state. As Knorr Cetina (1999) distinguishes, this is an “object to think with” rather than an “object to work with.” The photograph functions as what Latour (1986) terms an immutable mobile: stable, reproducible, easily circulating across press releases, investor presentations, and policy documents whilst maintaining its rhetorical impact.

Authority is further constructed through compounding mechanisms that create an appearance of comprehensive validation whilst all sources ultimately converge on Microsoft. DARPA selection provides governmental validation; Nature publication provides scientific endorsement; technical fellow credentials establish internal expertise; the vindication narrative — “high risk, high reward” approach having “paid off” — transforms Microsoft’s slower progress from potential liability into demonstrated wisdom. Each source reinforces the others, but as Collins (1992/1985) would note, this is authority accumulation rather than independent verification: DARPA evaluates Microsoft’s submissions, Nature reviews Microsoft’s paper, technical fellows are Microsoft employees. The competitive framing intensifies this logic: “many companies” pursued approaches that “ultimately weren’t able to scale,” whilst Microsoft’s “different path is now paying off.” Authority at the breakthrough stage derives from claimed achievement validating prior vision — the tortoise has reached the finish line — rather than from the objective technical assessment that characterised the pragmatic navigation stage.

6.4.2 Strategic Imaginary Re-Expansion

The announcement’s most analytically striking feature is the reversal of the imaginary contraction that characterised Troyer’s 2021 colloquium. Applications dismissed as “quantum wishing well” thinking — broad societal transformation, AI integration, universal problem-solving — return in expanded form: self-healing materials, environmental remediation, food security, any company being able to “design it perfectly the first time out.” This re-expansion reveals, as Jasanoff and Kim (2009, 2015) would predict, that imaginaries respond to strategic imperatives rather than purely to technical reassessment. Having declared fundamental problems solved, maintaining constrained imaginaries would appear unnecessarily modest. The imaginary re-expands because the breakthrough claim creates rhetorical space for it, not because the underlying technical situation has changed in ways that would justify revisiting the applications Troyer systematically eliminated.

How Quantum Computing Escapes the Maze

The temporal framing undergoes parallel reconfiguration. Troyer’s “next decades” for practical applications compresses to “years, not decades,” creating urgency that justifies immediate investment and attention. Yet “years” remains deliberately imprecise — two years? Nine years? — preserving flexibility whilst generating enthusiasm. The announcement trades the pragmatic stage’s credibility-protecting specificity for the breakthrough stage’s mobilising urgency. Following Selin’s (2008) analysis of future-making, this temporal compression performs imminence: quantum computing’s transformative impacts are no longer distant possibilities requiring patient investment but near-term realities justifying immediate action. The phrase “just give you the answer” — describing quantum-AI systems as oracles eliminating trial-and-error — constructs a remarkably frictionless future that reveals the re-expanded imaginary’s distance from the careful hedging that characterised Troyer’s pragmatic assessment.

The unaddressed tension between the re-expanded imaginary and Troyer’s prior dismissals is analytically significant. Troyer himself appears in the Majorana 1 announcement, now articulating AI-quantum integration visions that go considerably beyond the quantum chemistry applications he identified as the only viable domain in 2021. This discontinuity is not acknowledged. The announcement presents the breakthrough as vindicating Microsoft’s long-term strategy, but the imaginary re-expansion quietly reclaims territory that the pragmatic turn had explicitly abandoned. This dynamic confirms that imaginary evolution in Microsoft’s quantum narratives responds primarily to positioning needs: the 2021 constraint enhanced credibility when hype acknowledgement was strategically important; the 2025 re-expansion mobilises stakeholders when breakthrough claims make enthusiasm newly credible. The imaginaries are strategically deployed, not straightforwardly reflective of technical assessments.

6.4.3 Performativity, Scientific Controversy, and the Limits of Narrative Authority

Drawing on the speech act theory as developed by Austin (1962), Callon (1998) and MacKenzie (2006), the announcement functions as a performative utterance attempting to constitute breakthrough status rather than report independently verified achievement. When Microsoft declares “Microsoft has carved a new path to quantum computing,” this statement generates material consequences through uptake — reshaping competitive dynamics, influencing investment decisions, redirecting policy attention — regardless of whether the underlying technical claims are fully validated. The shift from the pragmatic stage’s conditional formulations (“could calculate,” “might enable”) to declarative assertions (“Microsoft has

carved,” “creates an entirely new state of matter,” “proves that topological qubits”) reflects the breakthrough stage’s characteristic mode: constituting new realities through authoritative announcement rather than awaiting comprehensive validation.

The scientific controversy that followed the announcement, however, reveals the limits of performative authority under conditions where expert communities can mount resistance. Nature reported that independent physicists maintained “evidence falls short of being definitive” and that subsequent theoretical work suggested Microsoft’s observations could be explained by mechanisms other than Majorana zero modes (Castelvecchi, 2025; Garisto, 2025). This controversy is analytically significant in multiple ways. It reveals the gap between corporate communication — emphasising “peer-reviewed confirmation” — and scientific consensus, demonstrating how selective citation of authority sources can construct certainty where significant uncertainty persists. It exemplifies the narrative dependency problem from an unexpected angle: not only external lay observers but this thesis itself cannot adjudicate the controversy, lacking both laboratory access and specialist physics expertise. And it demonstrates, following MacKenzie’s (2006) analysis of performativity, that declarations require uptake to succeed — when quantum physicists resist Microsoft’s breakthrough framing, the performative force of the announcement is contested rather than constituted.

Microsoft’s response to the controversy exemplifies the boundary work that Gieryn (1983) analyses. External validation — DARPA selection, Nature publication — is emphasised to establish legitimacy; critics are positioned as either insufficiently expert or overly cautious; the company’s long-term commitment to the approach is invoked as evidence of serious scientific engagement rather than commercial opportunism. The controversy thus becomes a struggle over which actors constitute the relevant expert community whose assessments determine what counts as achievement — a boundary contest, in Gieryn’s terms, between competing claims to scientific authority. What systematic narrative analysis reveals is how these claims are constructed, how the announcement’s literary technologies, authority-compounding mechanisms, and performative declarations function strategically. The analysis cannot resolve whether the underlying technical claims are valid. Under conditions of narrative dependency, that determination requires access — to the laboratory, to the experimental data, to the specialised expertise needed to evaluate them — that remains unavailable to external observers, and that is precisely the epistemological condition this thesis was designed to theorise. But it is also a political condition: the contest between Microsoft’s breakthrough declaration and the scepticism of independent physicists is a contest over interpretive authority — over whose account of what happened in the laboratory constitutes

the publicly recognised reality. That contest is not resolved by scientific evidence alone; it is also resolved, or left unresolved, by the institutional, rhetorical, and communicative resources that competing actors can bring to bear. The narrative dependency condition ensures that corporate actors begin that contest with a structural advantage: they hold the laboratory, the data, and the primary narrative.

6.5 SYNTHESIS: THE BREAKTHROUGH IMAGINARY

Step 5: Brief
Synthesis

Microsoft's 2025 Majorana 1 announcement constructs what I term the "breakthrough imaginary"—a vision of quantum computing as having crossed decisive thresholds from research to technology, from speculation to achievement, from laboratory confinement to imminent commercial deployment. This breakthrough imaginary represents the third stage in Microsoft's vanguard evolution, transforming strategies employed during both frontier exploration and pragmatic navigation.

Key Findings

The announcement demonstrates breakthrough status being narratively constituted through performative declaration rather than demonstrated capability. Microsoft constructs breakthrough through modalization reversal (transforming contested observations into established facts), authority accumulation (compounding technical credentials, peer review, and governmental validation), visual rhetoric (materialising scalability through the palm-of-hand image), and temporal compression ("years, not decades").

Authority shifts from expert assessment to vindication narrative. Microsoft's "high risk, high reward" approach has "paid off." The long-term investment in topological qubits, questioned whilst competitors pursued faster approaches, proves correct. Authority derives from claimed achievement validating past strategic decisions rather than ongoing technical assessment.

The sociotechnical imaginary re-expands dramatically. Applications dismissed during pragmatic navigation return: self-healing materials, environmental remediation, food security. Having declared fundamental problems solved, expansive imaginaries become newly credible. Troyer's AI-quantum integration vision—"Any company that makes anything could just design it perfectly"—returns to revolutionary rhetoric grounded in claimed breakthrough.

Temporal framing compresses to imminent deployment. The pragmatic stage's "next decades" becomes "years, not decades." This compression creates urgency justifying immediate investment whilst maintaining flexibility through deliberate imprecision. Competitor framing

shifts from implicit critique to explicit triumph—others pursued approaches that "weren't able to scale"; Microsoft's path is "now paying off."

The scientific controversy documented in *Nature* (Castelvecchi, 2025; Garisto, 2025) reveals the gap between corporate communication and scientific consensus. Microsoft declares "peer-reviewed confirmation" whilst independent physicists maintain "evidence falls short of being definitive" (p. 872). This gap illuminates how technological readiness is discursively constructed through strategic narrative, potentially diverging from scientific evaluation.

The Breakthrough Imaginary's Functions

The breakthrough imaginary performs multiple functions: retrospectively vindicating Microsoft's long-term investment; positioning Microsoft as having won quantum computing's foundational competition; creating urgency through compressed timelines justifying immediate investment; preparing potential customers for quantum computing integration; enrolling diverse stakeholders through compelling application scenarios; and maintaining credibility positioning through authority accumulation despite expansive claims.

Tensions and Vulnerabilities

The breakthrough imaginary creates significant vulnerabilities. Scientific controversy threatens its empirical foundation—if Microsoft's experimental interpretations prove incorrect, the entire narrative architecture collapses. Specific claims create falsifiable checkpoints. The re-expansion to applications previously dismissed creates narrative discontinuity for attentive observers. The gap between claimed breakthrough and demonstrated capability (eight qubits versus promised millions) creates ongoing exposure as stakeholders await material progress matching rhetorical claims.

With this third passage complete, the full trajectory of Microsoft's quantum narrative maze is now visible: from the frontier imaginary's open-ended promise, through the pragmatic turn's systematic narrowing, to the breakthrough imaginary's bold declaration of arrival. Chapter 7 draws together the findings from all three chapters to address the thesis's primary and subsidiary research questions, developing the theoretical contributions that this diachronic journey makes possible.

7 DISCUSSION: THE EVOLUTION OF QUANTUM NARRATIVES AND THE LIMITS OF KNOWLEDGE

Having completed the passage through all three sections of Microsoft’s quantum narrative maze, this chapter steps back to survey the terrain as a whole. Where Chapters 4, 5, and 6 followed each document through the analytical framework in sequence, this chapter reads across them — asking what the full trajectory reveals about how corporate narratives construct technological readiness, how theoretical concepts developed in Chapter 2 illuminate the patterns found, and what the exercise of navigating by insider-produced signposts alone reveals about the epistemological and political conditions of studying emerging technologies. The chapter addresses all four subsidiary questions and develops the thesis’s principal theoretical contributions.

7.1 SYNTHESISING NARRATIVE EVOLUTION

The comparative analysis reveals quantum computing's sociotechnical trajectory as Microsoft constructs it—not as linear progress toward predetermined goals but as strategic adaptation to changing circumstances, competitive pressures, and expectation dynamics. Table 2 synthesises the evolution across the three documents, organised according to the theoretical framework's analytical dimensions.

Table 2: Evolution of Microsoft's Quantum Computing Narratives (2014–2025)

Theoretical Dimension	Frontier Exploration (2014)	Pragmatic Navigation (2021)	Breakthrough Achievement (2025)
Corporate and Laboratory Narratives (Czarniawska; Latour & Woolgar)			
Modalization	Conditional possibilities ("could tackle," "could allow")	Systematic elimination of claims; hedged promises for viable domains	Declarative assertions; contested claims presented as established facts
Literary Inscription	Ethnographic storytelling; chalkboard photographs; Science cover as validation	Graphs, calculations, numerical comparisons; academic presentation format	Nature publication; chip photographs; "palm of hand" image as scalability inscription
Cycles of Credit	Charismatic researchers attracting collaborators; academic prestige networks	Academic credentials importing legitimacy; Simons Institute positioning	Compounding authority sources: technical fellows, DARPA, Nature, laboratory imagery

Theoretical Dimension	Frontier Exploration (2014)	Pragmatic Navigation (2021)	Breakthrough Achievement (2025)
Vanguard Evolution (Hilgartner; Borup et al.; van Lente & Rip)			
Authority Basis	Charismatic genius (Freedman's Fields Medal, exceptional individuals)	Expert assessment (systematic analysis, technical competence, methodological rigour)	Vindication ("high risk, high reward" paying off; long-term strategy proven correct)
Vision-Building Strategy	Celebrating frontier exploration; constructing pioneering identity	Navigating between hype and possibility; managing post-hype disappointment	Claiming achievement; declaring fundamental problems solved
Expectation Management	Generating enthusiasm whilst acknowledging uncertainty openly	Retrospective adjustment of promises; articulating promise-requirement cycles	Compressing time-lines; re-inflating expectations based on claimed breakthrough
Temporal Framing	Protective vagueness ("boomerang" with uncertain return)	Staggered milestones ("next few years," "next decades," eventual millions)	Compressed urgency ("years, not decades") with deliberate imprecision
Discursive Strategies (Gieryn; Collins; Austin; Callon; MacKenzie)			
Boundary Work	Quantum vs. classical computing; main-stream vs. fringe science	Science vs. hype; practical vs. theoretical; digital vs. analogue; Microsoft vs. competitors	Topological vs. all alternatives; scalable vs. limited; laboratory vs. deployment threshold
Expertise Construction	Heroic figuration; exceptional individuals transcending ordinary capability	Core-set membership through technical demonstration; questioning others' expertise	Vindicated experts; authority through demonstrated correctness of strategic choices
Performativity	Constituting Microsoft as quantum vanguard through declaration	Performing alternatives' illegitimacy; constituting viable pathways through elimination	Performing breakthrough status; declaring achievement generating material consequences
Competitor Framing	Collaborative scientific community: friendly rivalry ("share recipes")	Implicit dismissal (NISQ as "combining disadvantages"; applications as "wishing well")	Explicit triumph (competitors "hit a wall"; Microsoft's path "paying off")

Theoretical Dimension	Frontier Exploration (2014)	Pragmatic Navigation (2021)	Breakthrough Achievement (2025)
Sociotechnical Imaginaries (Jasanoff & Kim; Selin)			
Application Scope	Expansive ("wildest imagination-type applications" across multiple domains)	Constrained (systematic elimination, narrowing to quantum chemistry and materials science)	Re-expanded (self-healing materials, environmental remediation, food security)
Societal Benefit Claims	Universal transformation ("understand and control the very building blocks of the universe")	Focused transformation (nitrogen fixation, carbon capture within chemistry bounds)	Broad societal transformation (infrastructure, environment, hunger eradication)
Future Construction	Open-ended possibility; revolutionary but temporally indefinite	Structured pathway; achievable through specified requirements and milestones	Imminent realisation; transformative impacts within "years"
Imaginary Function	Stakeholder enrolment; investment justification; talent recruitment	Credibility restoration; competitive differentiation; expectation recalibration	Resource mobilisation; market creation; competitive triumph
Narrative Dependency (Knorr Cetina; this thesis)			
Laboratory Visibility	Ethnographic glimpses (researchers at chalkboards, laboratory at the beach)	Absent laboratories; knowledge claims through calculation rather than demonstration	Staged visibility (chip photographs); infrastructure black-boxed
Verification Possibility	Limited (Majorana detection as "compelling evidence, not definitive proof")	Impossible for external observers; reliance on expert authority	Impossible; scientific controversy reveals gap between claim and consensus
Strategic Mediation	Corporate communication presenting Microsoft favourably	Academic framing masking competitive positioning	Comprehensive authority accumulation obscuring contested scientific status

Patterns Revealed by the Synthesis

Table 2 reveals several significant patterns in Microsoft's narrative evolution:

Coordinated shifts across dimensions. The theoretical dimensions do not evolve independently but shift in coordinated fashion. When authority basis moves from charismatic (2014) to expert (2021) to vindication (2025), corresponding changes occur in boundary work, imaginary scope, and temporal framing. This coordination suggests vanguard evolution operates as integrated strategic repositioning rather than piecemeal adjustment.

Strategic responsiveness. Each stage's configuration responds to specific circumstances. Frontier-stage expansive imaginaries and charismatic authority suit conditions where demonstrated results remain limited but enthusiasm must be generated. Pragmatic-stage constrained imaginaries and expert authority suit post-hype conditions requiring credibility restoration. Breakthrough-stage re-expanded imaginaries and vindication authority suit conditions where claimed achievement justifies renewed enthusiasm.

Imaginary dynamics as strategic resource. The application scope row reveals imaginaries as strategically deployed rather than straightforwardly reporting technical possibilities. Expansion (2014), constraint (2021), and re-expansion (2025) follow strategic logic—mobilising stakeholders, restoring credibility, capturing breakthrough momentum—rather than reflecting linear technical development.

Increasing performative stakes. The performativity row reveals escalating stakes across stages. The frontier stage performs vanguard status—relatively modest claim. The pragmatic stage performs alternatives' illegitimacy—more aggressive but protected by academic framing. The breakthrough stage performs achievement itself—highest stakes, most vulnerable to falsification if claims prove incorrect.

Narrative dependency persistence. The bottom rows reveal narrative dependency persisting across all stages despite different visibility strategies. Whether through ethnographic storytelling (2014), academic presentation (2021), or authority accumulation (2025), external observers remain dependent on Microsoft-produced narratives, unable to independently verify underlying claims.

7.2 ADDRESSING THE PRIMARY RESEARCH QUESTION

How do corporate narratives construct and legitimise claims of technological readiness for emerging technologies confined to inaccessible laboratories?

The analysis demonstrates that technological readiness is narratively constituted rather than objectively determined. Across all three documents, Microsoft constructs quantum computing's readiness through narrative strategies rather than demonstrated capability—and the specific strategies evolve as circumstances change.

In 2014, readiness is constructed through establishing that *conditions* for eventual achievement exist: exceptional researchers, appropriate institutional arrangements, promising theoretical frameworks, preliminary experimental results. The modal constructions ("could tackle," "could

How Quantum Computing Escapes the Maze

allow") position quantum computing as possessing validated potential despite the absence of working devices.

In 2021, readiness is reconstructed through requirements articulation—specifying what quantum computing *needs* (fault tolerance, qubit thresholds, cycle times) to achieve practical advantage. This requirements-based framing transforms readiness from speculative potential to engineering challenge with identifiable milestones.

In 2025, readiness is constituted through performative declaration of breakthrough. Modalization reversal transforms contested observations into established facts. Authority accumulation compounds validation sources. Temporal compression creates urgency. The announcement performs breakthrough status regardless of underlying technical realities—generating material consequences through uptake.

The scientific controversy crystallises this finding. Microsoft declares "peer-reviewed confirmation" whilst independent physicists maintain evidence is not definitive. This gap reveals that technological readiness is discursively constructed through strategic narrative, and that construction may diverge significantly from scientific consensus. Readiness is not a property awaiting discovery but an achievement requiring continuous narrative work.

7.3 THEORETICAL CONTRIBUTIONS

7.3.1 Contribution 1: Narrative Dependency as Epistemological and Political Concept

SQ4: What does the reliance on corporate narratives — in the absence of direct laboratory observation — reveal about the epistemological conditions for studying technologies under access constraints, and what are the methodological implications for STS?

This thesis develops narrative dependency as a concept characterising the epistemological and political conditions for studying emerging technologies confined to inaccessible laboratories. Narrative dependency is not solely an epistemological condition but also a political one: when physical and epistemic access is foreclosed, the authority to determine what a technology is understood to be, what progress looks like, and whose assessment counts as credible rests exclusively with those inside the laboratory. The analysis of Microsoft's three documents demonstrates this political dimension concretely: across all three stages, corporate actors hold the primary narrative, and the struggle over interpretive

authority — visible most clearly in the 2025 Majorana controversy — is a contest over who gets to define what a breakthrough is and what evidence is sufficient to establish one.

As an epistemological concept, narrative dependency arises when four circumstances coincide: laboratory confinement preventing external observation; technical specialisation creating comprehension barriers such that limited access compounds into limited comprehension; corporate control limiting transparency; and public interest generating substantial stakes in accurate understanding. Under these conditions, external observers — including STS researchers, policymakers, investors, and publics — must rely predominantly on insider-produced narratives, not as one option among others but as the primary available epistemic resource.

Narrative dependency differs from related STS concepts. Unlike black-boxing (Latour, 1994/1987), which results from successful stabilisation, narrative dependency results from premature inaccessibility—technology remains unstabilised yet already inaccessible. Unlike the core-set problem (Collins, 1992/1985), which concerns evaluation competence, narrative dependency concerns access constraints preventing even competent observers from direct verification.

The concept has methodological implications. Under narrative dependency, systematic narrative analysis becomes methodologically necessary rather than merely useful. Documents cannot be treated as transparent windows onto development but must be analysed as strategic performances. The five-step analytical approach—context and material form, critical assessment of issues, narrative analysis, theoretical interpretation, synthesis—provides methodology appropriate to these conditions.

7.3.2 Contribution 2: Diachronic Vanguard Analysis

SQ1: How have Microsoft's quantum computing narratives changed between 2014 and 2025, and what do these changes reveal about the dynamics of corporate vision-building in emerging technology contexts?

This thesis extends Hilgartner's (2015) sociotechnical vanguards framework from synchronic to diachronic analysis, identifying three distinct stages with characteristic vision-building strategies:

Frontier Exploration (2014) relies on charismatic authority derived from exceptional individuals, employs expansive imaginaries unconstrained by demonstrated results, maintains protective temporal vagueness, and frames competitors collaboratively as fellow scientific community members.

How Quantum Computing Escapes the Maze

Pragmatic Navigation (2021) shifts to expert authority derived from systematic technical analysis, constrains imaginaries to defensible application domains, specifies staggered temporal milestones, and positions competitors critically through implicit dismissal of their approaches. The 2021 Nature retraction provides important context for this stage: the colloquium's systematic demarcation of viable from non-viable applications functions not only as forward-looking expectation management but as credibility reconstruction following a public scientific failure.

Breakthrough Achievement (2025) claims vindication authority derived from strategic choices proving correct, re-expands imaginaries to broad societal transformation, compresses timelines to create urgency, and frames competitors as having definitively lost the foundational competition.

The diachronic extension reveals that strategies effective at one stage may prove counterproductive at another. Frontier enthusiasm generates stakeholder excitement but creates hype vulnerabilities. Pragmatic constraint restores credibility but limits mobilisation potential. Breakthrough declarations generate urgency but create falsification risks. For vanguard evolution to succeed, it is vital to recognise when transitions should occur and to implement them without catastrophic credibility loss.

Integrating sociology of expectations (Borup et al., 2006; van Lente & Rip, 1998a, 1998b) explains transition pressures. Promise-requirement cycles generate expectations that material progress must eventually address. Hype disappointment necessitates pragmatic recalibration. Accumulated authority and peer-reviewed results enable breakthrough claims. The vanguard evolves not through autonomous choice but through navigation of expectation dynamics that it has itself helped to create.

7.3.3 Contribution 3: Integrated Discursive Legitimation Analysis

SQ2: What discursive strategies do these narratives employ to construct authority and demarcate technological approaches, and how do these strategies relate to one another?

This thesis demonstrates how boundary work (Gieryn, 1983), expectation management (Borup et al., 2006), and performativity (Austin, 1962; Callon, 1998; MacKenzie, 2006) operate as integrated legitimation strategies rather than isolated mechanisms. Microsoft's narratives simultaneously perform these functions, their interrelation providing comprehensive legitimation machinery.

Boundary work shifts across stages: from quantum versus classical computing (2014), to science versus hype and practical versus theoretical (2021), to topological versus all alternatives and scalable versus limited (2025). Each configuration demarcates legitimate approaches whilst positioning Microsoft favourably.

Expectation management evolves from generating enthusiasm whilst acknowledging uncertainty (2014), through retrospective adjustment of promises and articulation of promise-requirement cycles (2021), to compressing timelines and re-inflating expectations based on claimed breakthrough (2025).

Performative declarations escalate from constituting Microsoft as quantum vanguard (2014), through performing alternatives' illegitimacy (2021), to performing breakthrough status itself (2025). The scientific controversy reveals performativity's contested character—declarations require uptake, and resistance from scientific communities limits performative force.

Understanding these strategies' interrelation provides richer analytical purchase than examining any single mechanism independently. Boundary work creates the categories within which expectations are managed; expectation management establishes the temporal frameworks within which performative declarations gain plausibility; performative declarations materialise the boundaries and expectations through their uptake.

7.3.4 Contribution 4: Strategic Imaginary Dynamics

SQ3: What sociotechnical imaginaries do these narratives construct, and how do those imaginaries change over time in relation to Microsoft's strategic positioning?

This thesis traces how sociotechnical imaginaries change across the three documents in response to strategic imperatives, extending Jasanoff and Kim's (2009, 2015) framework to address imaginary evolution over time rather than at a given moment. The analysis reveals a pattern of expansion, constraint, and re-expansion that maps onto the vanguard stages identified above — a finding, rather than a presupposition, of the diachronic analysis.

Imaginary expansion (2014) encompasses “wildest imagination-type applications” across medicine, cryptography, materials science, machine learning, and fundamental understanding of the universe. This expansive scope enrolls diverse stakeholder communities, justifies investment despite uncertain timelines, and attracts scientific talent to frontier exploration.

How Quantum Computing Escapes the Maze

Imaginary constraint (2021) systematically eliminates big data, machine learning, weather prediction, protein folding, and other applications, narrowing viable scope to quantum chemistry and materials science. This constraint restores credibility following hype disappointment whilst positioning Microsoft's fault-tolerant approach as addressing quantum computing's actual possibilities.

Imaginary re-expansion (2025) returns to broad societal transformation: self-healing materials, environmental remediation, food security, AI-quantum integration enabling transformative industrial capability. The unaddressed tension between 2021 dismissals and 2025 returns confirms that imaginary construction responds to strategic positioning needs as much as to technical assessment. The dynamic reveals imaginaries as resources that corporate vanguards deploy strategically rather than straightforward reflections of technical possibility.

7.4 THE BREAKTHROUGH CONTROVERSY REVISITED

The scientific controversy surrounding Microsoft's Majorana claims illuminates this thesis's arguments whilst revealing narrative analysis's inherent limitations.

The controversy demonstrates narrative dependency's operation. External observers cannot adjudicate whether experimental signatures indicate Majorana zero modes or "more mundane physics" (Garisto, 2025, p. 555). The specialised expertise and laboratory access required for evaluation remain confined to small communities whose members themselves disagree. It also demonstrates the political dimension: the contest between Microsoft's breakthrough declaration and the scepticism of independent physicists is a contest over interpretive authority — over who gets to define what the evidence means and what counts as a validated achievement. That contest cannot be resolved by rhetoric alone, but corporate actors begin it with a structural advantage because they control the laboratory, the data, and the primary narrative.

The controversy reveals the gap between corporate communication and scientific consensus. Microsoft's announcement declares "peer-reviewed confirmation" whilst independent physicists maintain "evidence falls short of being definitive" (Castelvecchi, 2025, p. 872). This gap makes visible how corporate narratives perform certainty where significant uncertainty persists — precisely the narrative constitution of technological readiness this thesis analyses.

The controversy also demonstrates performativity's contested character. Whether Microsoft's breakthrough declaration achieves uptake — whether stakeholders treat breakthrough as

accomplished — depends on ongoing negotiation between corporate claims and scientific evaluation. Resistance from the physics community limits performative force regardless of the announcement's rhetorical sophistication.

However, the controversy also reveals narrative analysis's limits. This thesis can analyse how Microsoft constructs breakthrough claims, what discursive strategies the announcement employs, how it functions within vanguard evolution. It cannot determine whether Microsoft actually observed Majorana particles or whether topological qubits will achieve practical quantum advantage. That determination requires laboratory access, experimental data, and specialised expertise unavailable to external observers under conditions of narrative dependency. This is not a methodological failure but a recognition of the epistemological and political situation the thesis theorises.

The discussion arrives at the exit of the maze — or more precisely, at the point where this analysis can determine what the maze itself reveals. Chapter 8 consolidates these insights into a final reflection on the thesis's contributions, its limitations, and the research directions it opens.

7.5 IMPLICATIONS FOR STS METHODOLOGY

This thesis demonstrates methodological innovations with broader implications for studying emerging technologies:

7.5.1 Document Analysis Under Access Constraints

Systematic document analysis (Asdal & Reinertsen, 2022) becomes the chosen methodological approach when ethnographic access is not possible. Documents cannot be treated as transparent windows onto development but must be analysed as strategic performances with audiences, purposes, and effects. The five-step analytical approach—context and material form, critical assessment of issues, narrative analysis, theoretical interpretation, synthesis—provides replicable methodology for future research on technologies exhibiting narrative dependency. The "critical assessment of issues" step proves particularly valuable. Systematically identifying what documents make present and absent, what evidence supports claims, what caveats are omitted, what controversies are acknowledged or silenced—this assessment reveals strategic choices that straightforward reading might miss.

7.5.2 Productive LLM Use

This thesis demonstrates large language models as interpretive aids for engaging with technically complex materials. Troyer's 2021 colloquium presented significant comprehension challenges—presentations that Microsoft's own research director characterises as making him "feel a bit thick... and small" (Warnick, 2014). Rather than limiting STS analysis only to topics where researchers possess deep disciplinary expertise, productive LLM use enabled engagement with specialist content whilst maintaining analytical rigour. The approach requires transparency about LLM use, limited scope focusing on concept explanation rather than analytical interpretation, multiple sources for cross-checking, and critical engagement recognising potential inaccuracies. LLMs functioned as technical dictionaries helping understand *what* Troyer was saying so that STS analysis could examine *how* and *why* he was saying it.

7.5.3 Insider-Outsider Positioning

The thesis models how industry experience can inform rather than merely compromise STS analysis. As Microsoft employee analysing Microsoft documents, my positioning required explicit acknowledgement of potential biases, grounding in established theoretical frameworks, and focus on how narratives work rather than whether claims are true. This insider-outsider status enabled reading documents with both familiarity (understanding corporate communication conventions, recognising rhetorical strategies) and critical distance (applying STS analytical frameworks, identifying strategic absences). The positioning provided analytical leverage when appropriately reflexive—industry experience informs interpretation whilst theoretical frameworks maintain critical distance.

7.6 CLOSING THE DISCUSSION

The analysis of Microsoft's quantum computing narratives across eleven years demonstrates that technological readiness is narratively constituted rather than objectively determined. Corporate narratives construct readiness through evolving strategies—frontier enthusiasm establishing pioneering credentials, pragmatic assessment restoring credibility, breakthrough declaration claiming fundamental problems solved.

The theoretical contributions—narrative dependency as epistemological concept, diachronic vanguard analysis, integrated discursive legitimation strategies, strategic imaginary dynamics—provide analytical resources for understanding emerging technologies under conditions of limited access. The framework and methodology developed here apply beyond quantum computing to

other technologies exhibiting narrative dependency—where laboratory confinement, technical specialisation, corporate control, pre-deployment status, and public interest combine to create situations where external observers must rely predominantly on insider-produced narratives.

The scientific controversy surrounding Microsoft's breakthrough claims reveals both the power and limits of narrative analysis. Systematic analysis reveals how claims are constructed, how their construction evolves strategically, and what this construction reveals about knowledge production. It cannot resolve underlying technical questions—but under narrative dependency, those questions may remain unresolvable through any available method. The thesis's contribution lies in making the narrative construction visible and analytically tractable, providing resources for understanding emerging technologies when direct observation remains impossible.

8 CONCLUSION: TOWARDS BETTER UNDERSTANDING OF EMERGING TECHNOLOGIES

This final chapter steps back from the maze. Having traced Microsoft's three-stage narrative journey from 2014 to 2025 — frontier exploration, pragmatic navigation, and breakthrough declaration — this chapter summarises key findings, acknowledges the limits of the analytical approach, and reflects on what the journey reveals about how emerging technologies are known under conditions of narrative dependency. It also suggests where the next researchers might pick up the thread.

8.1 SUMMARY OF FINDINGS

This thesis addressed the primary research question: How do corporate narratives construct and legitimise claims of technological readiness for emerging technologies confined to inaccessible laboratories?

Through systematic analysis of three Microsoft documents spanning eleven years, the thesis demonstrates that technological readiness is narratively constituted rather than objectively determined. Microsoft constructs quantum computing's readiness through evolving strategies: frontier enthusiasm establishing pioneering credentials (2014), pragmatic assessment restoring credibility through honest-seeming evaluation (2021), breakthrough declaration claiming fundamental problems solved (2025).

The subsidiary questions yielded specific findings:

SQ1: How have Microsoft's quantum computing narratives changed between 2014 and 2025, and what do these changes reveal about the dynamics of corporate vision-building in emerging technology contexts? Microsoft's quantum narratives evolve through three vanguard stages — frontier exploration, pragmatic navigation, breakthrough achievement — each employing distinct vision-building strategies. Authority bases shift from charismatic genius through expert assessment to vindication narrative. These shifts reveal vanguard evolution as strategic adaptation to changing circumstances, competitive pressures, and expectation dynamics rather than linear progress.

SQ2: What discursive strategies do these narratives employ to construct authority and demarcate technological approaches, and how do these strategies relate to one another? Microsoft's narratives employ integrated discursive strategies — boundary work demarcating legitimate approaches, expectation management calibrating promises to circumstances, performative declarations constituting technological status. These strategies operate in concert, their

interrelation providing comprehensive legitimation machinery that is analytically more revealing than any single mechanism examined in isolation.

SQ3: What sociotechnical imaginaries do these narratives construct, and how do those imaginaries change over time in relation to Microsoft's strategic positioning? Sociotechnical imaginaries expand, contract, and re-expand following strategic logic. Application scope responds to positioning needs as much as to technical assessment. The 2025 return of applications Troyer dismissed in 2021 reveals imaginaries as strategically deployed constructions whose evolution tracks corporate imperatives rather than straightforwardly reflecting technical possibilities.

SQ4: What does the reliance on corporate narratives — in the absence of direct laboratory observation — reveal about the epistemological conditions for studying technologies under access constraints, and what are the methodological implications for STS? Narrative dependency — the condition arising from the reliance on insider-produced accounts under conditions of physical inaccessibility, technical specialisation, corporate control, and high public stakes — is both an epistemological and a political condition. External observers cannot verify claims, must navigate strategic mediation, and encounter constructed transparency. The scientific controversy surrounding Majorana claims exemplifies these challenges and illustrates how the struggle over interpretive authority is structural, not merely rhetorical.

8.2 LIMITATIONS

This thesis acknowledges significant limitations. Analysing only Microsoft documents limits generalisability to Microsoft's quantum narratives specifically. As STS scholar without deep quantum physics expertise, my technical evaluation capacity remains limited despite LLM assistance. The thesis analyses emergence in real-time, limiting temporal perspective for assessing significance. Access constraints meant analysing public performances rather than development processes themselves. These limitations are inherent to narrative dependency. The methodology acknowledges that narrative analysis complements rather than replaces technical evaluation and ethnographic observation. The thesis's contribution lies in analysing how claims are constructed, not determining whether they are technically valid.

8.3 FUTURE RESEARCH DIRECTIONS

Several research directions emerge from this analysis:

Comparative Corporate Analysis. Extending this methodology to IBM, Google, IonQ, and other quantum computing corporations would reveal whether the patterns identified here — the three-stage vanguard trajectory, the coordinated shifts across legitimization dimensions, the strategic imaginary dynamics — generalise across the field or reflect Microsoft-specific strategic choices.

Ethnographic Access. Pursuing ethnographic access to quantum computing laboratories would complement narrative analysis, revealing the relationship between public narratives and internal understanding. As quantum computing transitions towards deployment, such access may become less restricted. Exploring the conditions under which corporations might grant independent scholarly observation — and documenting cases where this has occurred — represents a valuable direction that would directly address the narrative dependency condition this thesis theorises.

Retrospective Assessment: Revisiting Microsoft's 2025 predictions in 2030 or 2035 would assess whether breakthrough claims proved accurate, providing evidence about expectation dynamics and narrative-reality relationships.

Cross-Technology Comparison: Examining whether narrative dependency patterns characterise other emerging technologies—synthetic biology, advanced AI, fusion energy—would test the concept's broader applicability.

Corporate Engagement with Ethnographic Research: A significant future direction involves exploring whether corporations like Microsoft might invite ethnographic study of their quantum computing laboratories. Such engagement could benefit corporations in multiple ways whilst advancing STS scholarship.

Independent ethnographic study could provide external validation that corporate communications cannot achieve alone. As this thesis demonstrates, the scientific controversy surrounding Majorana claims reveals gaps between Microsoft's declarations and independent scientists' assessments. Ethnographic observation by credible STS scholars could provide different validation—not of technical claims per se, but of the rigour, openness, and integrity of research practices. This independent witnessing might carry persuasive weight with stakeholders sceptical of corporate self-reporting.

Corporations might recognise that narrative dependency ultimately undermines trust. If external stakeholders must rely exclusively on corporate-produced narratives, sophisticated observers will discount those narratives accordingly. Inviting ethnographic access would partially address this problem, demonstrating confidence in research practices and willingness to submit to external scrutiny. The very act of granting access performs openness, potentially enhancing credibility more than corporate communication could achieve.

Furthermore, ethnographic observers often identify dynamics invisible to insiders. An STS scholar studying quantum laboratories might reveal assumptions, blind spots, or organisational patterns that internal actors take for granted. This external perspective could prove valuable for identifying risks, improving research practices, or understanding how organisational culture shapes technical choices. The sociology of scientific knowledge has repeatedly demonstrated that external observation reveals aspects of knowledge production that practitioners themselves cannot see.

As quantum computing attracts increasing policy attention—export controls, national security concerns, workforce development initiatives—corporate positioning with policymakers becomes strategically important. Ethnographic studies demonstrating responsible, rigorous research practices could support favourable regulatory treatment. Policymakers facing narrative dependency themselves might welcome independent scholarly assessment as input for governance decisions.

Such studies would require careful structuring to address legitimate corporate concerns. Confidentiality provisions could protect proprietary information; embargoes might govern sensitive technical details; review processes could ensure commercially sensitive information remains protected. Security constraints, particularly given DARPA involvement and national security implications, might require restricted access to certain laboratory areas—restrictions ethnographic studies can accommodate whilst still providing valuable insight into accessible aspects of research practice.

The case for ethnographic access rests on recognising that narrative dependency ultimately serves no one well. Corporate claims would carry greater weight if independently witnessed; external observers would understand emerging technologies better if granted access; and STS scholarship would benefit from empirical grounding currently unavailable. Exploring conditions under which such access might become possible—and documenting cases where it has been granted—represents a valuable direction for future research addressing the epistemological challenges this thesis has identified.

8.4 BROADER IMPLICATIONS

This thesis's findings extend beyond quantum computing to emerging technology governance generally:

Narrative-Aware Governance: Policymakers assessing emerging technologies must recognise that corporate communications construct technological status rather than neutrally report it. Governance under narrative dependency requires critical analysis of how claims are constructed, not merely what claims are made.

Transparency Imperatives: Narrative dependency could be reduced through greater transparency and access. However, transparency faces real barriers—proprietary interests, security concerns, technical complexity. Governance must develop approaches for assessing technologies when full transparency proves impossible.

Expectation Management: The strategic deployment of imaginaries revealed in this thesis suggests governance should actively manage expectations rather than passively accepting corporate framings. Regular independent assessments, historical analysis of expectation cycles, and accountability mechanisms for failed predictions deserve consideration.

8.5 FINAL REFLECTION

This thesis began with a puzzle: quantum computing appears simultaneously “already begun” and perpetually “years away.” The maze metaphor that structures this analysis was not imposed from outside but borrowed from within — from Troyer’s own 2021 colloquium, where it describes the bewildering landscape of quantum applications. That the same metaphor describes the epistemological situation of the external observer is not coincidental. Microsoft’s corporate narratives have functioned, across eleven years, as Ariadne’s thread: the insider-produced guide that external observers must follow, because no independent path through the maze exists. This thesis has followed that thread carefully and critically — tracing its construction, noting where it turns, and asking what it forecloses as well as what it opens.

This thesis never set out to adjudicate Microsoft’s claims. Its goal was not to determine whether the exit exists, but to analyse how the maze was built — how authority was constructed, how promises were calibrated, how the boundary between legitimate science and hype was drawn and redrawn across three documents and eleven years. Narrative dependency, the condition that makes document analysis the analytically appropriate method, means that the question of whether Microsoft’s topological path leads to a genuine exit remains open — awaiting experimental verification, engineering development, and the

unforgiving test of deployment. But the analysis has also shown that this openness is not simply a matter of scientific uncertainty: it is a political condition, in which the authority to define what counts as progress, what counts as a breakthrough, and what counts as sufficient evidence is held by those who control the laboratory. That structural asymmetry is what this thesis, in theorising narrative dependency, has sought to make analytically visible.

What this thesis has produced is a map. The three-stage trajectory — frontier imaginary, pragmatic turn, breakthrough declaration — charts how Microsoft constructed the maze for public navigation: how imaginaries expanded, contracted, and re-expanded; how charismatic authority gave way to expert assessment and then to vindication narrative; how temporal framings shifted from protective vagueness to staggered milestones to compressed urgency. That map retains its analytical value regardless of how the material questions are eventually resolved. If Microsoft's claims prove correct, it documents how a genuine breakthrough was narratively constituted before independent verification was possible. If they prove incorrect, it documents how corporate narrative sustained credibility across three developmental stages in the absence of deployable results. The framework and methodology developed here are available for application to quantum computing's continuing development, and to other emerging technologies where similar conditions of narrative dependency obtain.

APPENDIX I – TRANSCRIPT OF TROYER’S COLLOQUIUM

The transcript was created using notegpt.io. Microsoft Copilot translated the spoken language into written language.

Reference:

Troyer, M. (2021). *Towards Practical Quantum Advantage | Quantum Colloquium*.

<https://www.youtube.com/watch?v=WY3htdKUGsA>

The title of Matthias Troyer’s colloquium is “Towards Practical Quantum Applications”, given on 16 March 2021.

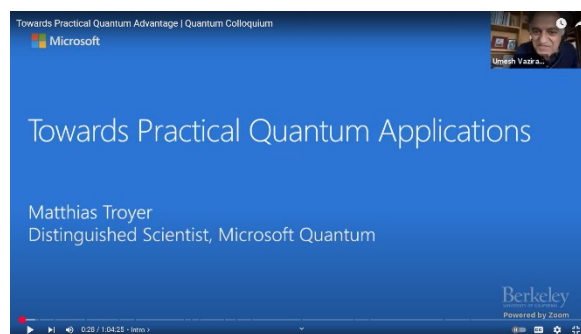
Host: Umesh Vazirani, Professor of Electrical Engineering and Computer Science at the University of California, Berkeley, and the director of the Berkeley Quantum Computation Center.

The text snippets used as direct quotes are highlighted.

00:00:00

[Host:] So, Matthias, did you check if you can share the screen with us? Do you have co-host privileges? [Troyer:] We can try that now. Let me just see if it works. See, and now I think... Okay, fantastic, great. [Host:] Okay, so, well, welcome everybody to the colloquium today. Thanks, Brian, for this wonderful introduction to the field.

It's a great pleasure to have Matthias Troyer here from Microsoft Quantum, where he is a distinguished scientist. Matthias is probably one of the world's leading computational condensed matter physicists. He's been awarded the Rahman Prize and the Hamburg Prize for Theoretical Physics. But more to the point, he has spent many years thinking single-mindedly about how to achieve quantum advantage. So, it's a great privilege to have Matthias talk to us about it. Please, Matthias.



Slide 1 – Towards practical quantum applications

00:00:45

[Troyer:] Yes, thank you. I got to the point where I said to really make progress. I would love to have quantum hardware because I know how to solve these problems really easily once we have a quantum computer. So, that's why I joined Microsoft.

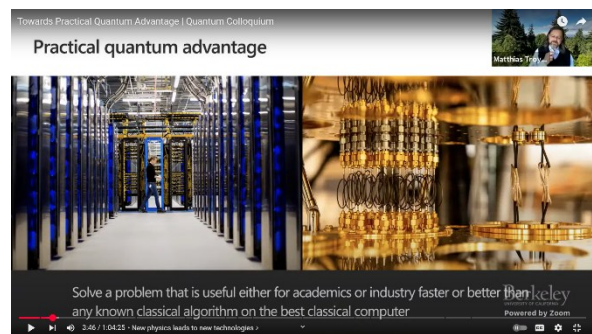
In the meantime, I've learned new jobs, worked on quantum modelling, learned to work with customers, and now my new role is that of the systems architect, thinking about what is needed to really build the quantum computer.

The basic question I'm asking today is: What do we need to really solve practical problems with quantum hardware? The basics are that new physics always lead to new technologies. Some of thermodynamics led to the invention of the steam engine and the industrial revolution. Electromagnetism led to electronics, computers, cell phones, and much more. Quantum physics is now more than a century old, but we're only starting to think about what else we could do with quantum systems. What I'm interested in is what I call practical quantum advantage. I'm interested in really solving a useful problem that is useful either academically or for industry—faster and better on a quantum computer than on any known classical computer. I'm not interested in whether some random circuit can be done better on quantum hardware than

classical hardware. I really want to solve a problem that interests somebody and do it better than on any classical computer. These are the kinds of things where we say, "Hey, we know we can't do it classically, but quantumly we can." That's really the interesting application and that's why it's worth investing in building quantum hardware.



Slide 2 – New physics leads to new technologies

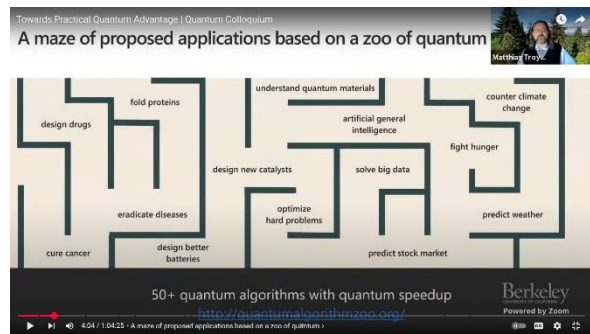


Slide 3 – Practical quantum advantage

How Quantum Computing Escapes the Maze

00:03:21

When you listen to people, there's a maze of proposed applications. One reads a lot in the media about how quantum computers will help with big data, cure cancer, design new drugs, proteins, predict the stock market, and so on. There are loads of things where people claim that quantum will solve all these problems. But looking at that, for each of these, there's always

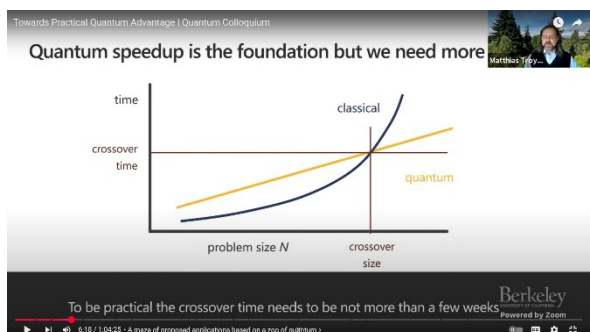


Slide 4 – A maze of proposed applications based on a zoo of quantum

one known quantum algorithm that can somehow be applied. But this looks more to me like the hope for a quantum wishing well. These are all hard problems for which we want a solution, and we need to see now which of those problems can really be tackled by quantum computers. The starting point, of course, is a quantum algorithm with quantum speedup. The most comprehensive list is on quantum algorithms by Stephen Jordan. But that is only the start because most of the work on algorithms so far has been on complexity theory. It is the asymptotic scaling when the problem size gets big enough; the complexity increases slower quantumly than classically. But in order to really realize that, I want to go to a problem size where quantum is faster than classical. For small problems, classical is faster because it's simpler and cheaper to build classical hardware. It's only for the big sizes where quantum will win. I want to look at the cost of a size where the crossover happens and the time because I care about solving a problem in a few weeks and not in a millennium, not asymptotically. So the question is: What are the problems where the crossover time can be short enough that it could be practical in the next decades?

00:05:19

For that, I now want to compare a classical system with a quantum system. I want to totally bias it towards the quantum system. I want to say, let me put in the race on the classical side a single today's classical chip, the most advanced NVIDIA chip with 54 billion transistors. Let's take one of those. Quantumly, I envision that in some time we'll build a machine with 10,000 logical qubits, fast ones, with a logical cycle time of a few microseconds. Let me ignore data movement by assuming that with some matching, we'll manage all-to-all connectivity. I want to compare those two to see what this future quantum computer could do



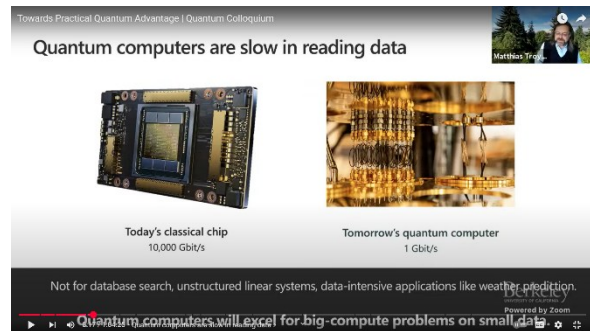
Slide 5 – Quantum speedup is the foundation, but we need more

How Quantum Computing Escapes the Maze

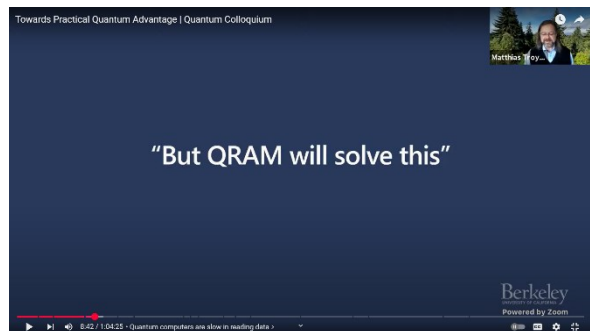
better than a single classical chip that I can just buy today. There, I want to start and debunk some hype. One thing one often hears is that quantum computers will solve the big data problem. But when you look at that, quantum computing in the classical chip can read 10 terabits a second. The quantum computer with 10,000 qubits at the clock speed of 2 to 10 microseconds can read about one gigabit a second. So that future quantum machine is about a factor of 10,000 slower for classical I/O than today's classical computer. Thus, it is very clear to all, I think, just looking at that, it's not for big classical data. Quantum computers will excel on big compute problems on small data or maybe on problems on quantum data that we have to find.

00:08:12

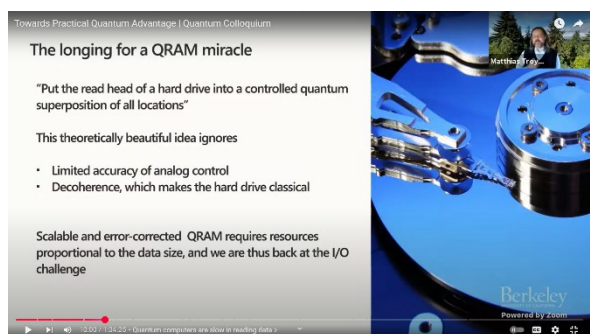
But then mentioning that, one often also hears that quantum RAM (QRAM) will solve this. I've tried to understand that, and the quality helped me understand what QRAM is about, how it speeds things up. The way I like to present it in a simplified way, and that's something we can talk about later, is basically the idea behind QRAM is taking the classical computer with the I/O system and putting everything into a quantum superposition. So, the idea of this QRAM oracle is, in a simplified way, you take a hard disk, and you can read all the data at once by putting the read head of the hard drive in a quantum superposition of all the locations on the disk. To anybody who builds mechanical quantum systems or organic quantum devices, it's clear the dissipation will stop that this is not scalable because it's analogue, there's noise, decoherence, and it is a beautiful theoretical idea. But really implementing it will need quantum error correction, and then the scaling goes back to being proportional to the data size. It would be great if somebody could make that really rigorous and give us a bound on what it takes in a fault-tolerant way. But that's why we're saying I really want to focus on the small data problems that are hard to compute. The neat way of saying it is quantum



Slide 6 – Quantum computers are slow in reading data



Slide 7 – But QRAM will solve this



Slide 8 – The longing for a QRAM miracle

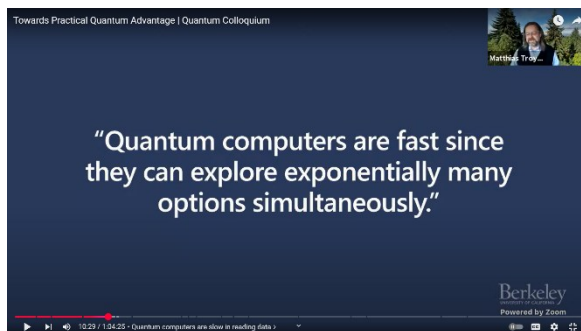
How Quantum Computing Escapes the Maze

computers are faster because they can explore exponentially many inputs at once. But we know it's not that easy because while it can get to compute in superposition, I need more than that because I need to be able to read out the results. When I just look at the answer, I get some random output that won't help me. So, we need to go beyond that. One way of extracting the data would be with Grover's search, Grover's algorithm, with a quadratic speedup over classical computing. So, it may not be exponentially faster, but in many cases, we can get at least a quadratic speedup.


00:11:12

I want to look into how the numbers work out by comparing that classical chip with the future quantum computer. Today's classical chips are amazingly fast. The chip per second numbers are about five petaflops for logical operations. If I built an ASIC with the same pair of transistors, it could be a bit faster. On the quantum computer, if I want to do a logical operation, I can do about a few million operations per second. If you want to do a floating-point operation, assuming I ignore all the cost of data movement and layout, ignoring all the overhead of reversible computation, the peak performance that it could get is a few thousand operations per second. So, we see just comparing a single classical chip with a future 10,000 logical qubit quantum machine, the

quantum machine is 10 to 12 orders of magnitude slower than the classical one. But there's a constant slowdown. The quantum machine will then win with quantum speedup because it scales better as we go to larger systems. But if we now think about it, I can overcome that constant slowdown because I have to make fewer calls to an oracle. But if the constant slowdown is about 10 to the 12, then I have to make 10 to the 12 fewer calls. That means with quadratic speedup, I



Slide 9 – Quantum computers are fast since they can explore exponentially many options simultaneously



Slide 10 – Quantum computers need fewer operations, but each operation is slower by 10-12 orders of magnitude

Towards Practical Quantum Advantage | Quantum Colloquium

Quantum speedup can overcome the constant slowdown

Matthew Troy

Overcoming a 10^{12} slowdown by making 10^{12} times fewer function calls

Speedup	Classical function calls	Quantum function calls
Quadratic	10^{24}	10^{12}
Cubic	10^{18}	10^6
Quartic	10^{16}	10^4

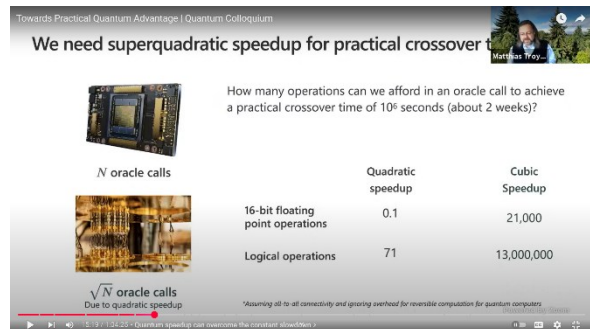
Slide 11 – Quantum speedup can overcome the constant slowdown

How Quantum Computing Escapes the Maze

start winning if quantumly I make 10 to the 12 calls versus classically 10 to the 24. Or with cubic speedup, it looks a bit better; it's a million versus 10 to the 18. But you need to make 10 to the 12 times fewer calls because the machine is a constant about 10 to the 12 slower.

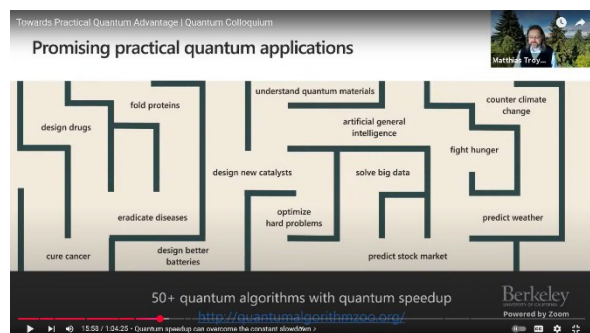
00:13:43

And now you can turn that around and ask: If we have quadratic speedup, what is the time we need to do that? I think one can see it right here. The time to do (10^{24}) operations is years or months at least, so quadratic speedup may not be enough. I turned it around and asked: How many operations can I do in one call to the oracle so that the crossover time is not more



Slide 12 – We need superquadratic speedup for practical crossover

than a few weeks? The answer is simply that with floating-point operations, I can't get there with the assumptions I made on the quantum machine and a single classical one. I just can't have it. But if I have a very simple oracle with just a few up to 972 logical operations, then with those assumptions, the quantum machine might win. We're still very optimistic for quantum, pessimistic for classical. So, I'd say clearly, we need more than quadratic speedups. Quadratic speedups are beautiful in theory, but if we want to build a machine and calculate something faster than classically, then I will need more than quadratic speedups. If we go to cubic speedups, they can do a bit more, 20,000 operations. So, I think that's when it starts getting interesting. I want to look at what is now the real cost of doing that on the quantum machine and whether we could solve the problem by just having 100 GPUs. Really, we want exponential speedups; that's where we have the best chance.



Slide 13 – Promising practical quantum applications

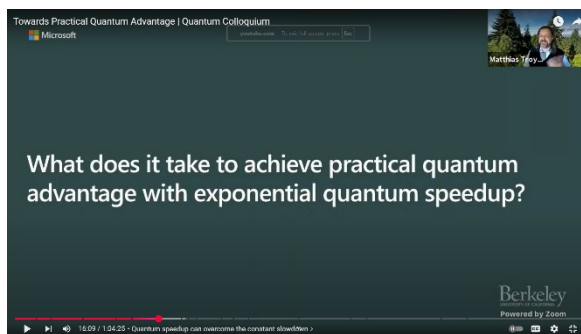
How Quantum Computing Escapes the Maze

00:15:40

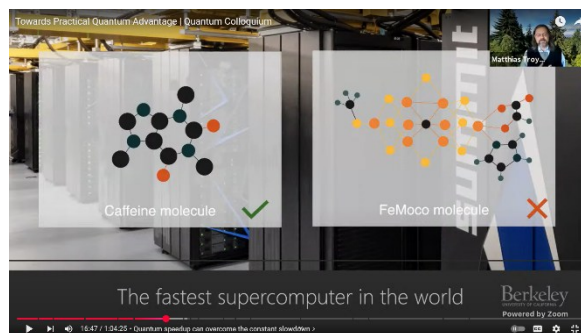
From that whole list of problems that one often hears, I like to drop out all of the problems that are big data problems, machine learning problems, weather prediction, stock market predictions.

This is all tied to the data, and the proposals for problems like protein folding and drug design are also often tied to Grover's speedup. So those I'd like to drop out. I really think we need to focus on the algorithms with exponential speedup. What do we have that we can do there? One problem is factoring. I'm not that interested in factoring numbers; there's not a great application where I solve a real problem.

But there are other problems where I'm getting interested and where there can also be exponential quantum speedup. Those are problems of quantum systems, simulating quantum systems. When I simulate a simple molecule like caffeine, I can't do it exactly brute force. But with the methods you mentioned in the tutorial, we have good enough classical approximations that let me simulate a molecule. I can predict its properties; it works. But there are other molecules, like here I mentioned the nitrogenase molybdenum-iron complex, where we don't have any known classical algorithm that can calculate it well enough. So, there are these problems in chemistry and in material science where we have no way of solving them to the level we want to get to on a classical computer, but where a quantum computer can do it.



Slide 14 – What does it take to achieve practical quantum advantage with exponential quantum speedup?




Slide 15 – The fastest supercomputer in the world


How Quantum Computing Escapes the Maze

00:17:31

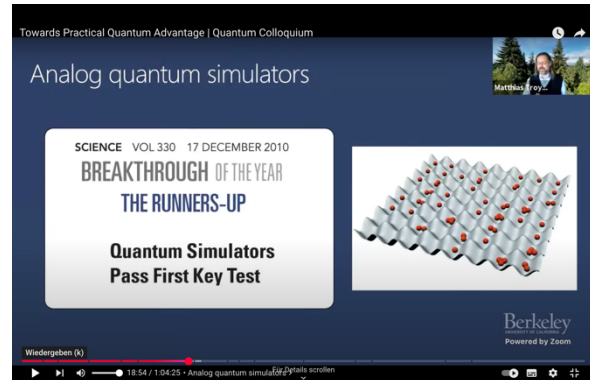
One example is a leaf in spring. When you look outside in the spring and the leaves come up, you see the green everywhere. Nature knows the leaf is green, but calculating the colour of chlorophyll is hard classically. Why is that? It is because nature is quantum and classical computers are not. That's why Feynman, decades back, said to simulate quantum physics, to simulate quantum problems, we need quantum hardware. We should think of a computer that works on quantum mechanical principles, as we all know. They already exist in some way because the first such machines were analogue, just like in the classical world where the first computers were analogue ones. So quantumly, people started building simulators for special-purpose machines that could solve certain quantum models already more than a decade ago. For example, you think of cold atoms, where people take an atom in a trap, shown on the left here, and use lasers to build a special-purpose computer, an analogue one, to really build that model in the lab and measure things on it. With that, one can do models that are hard classically. But how do we know whether it works or not? The first thing done was to take a model that one knows how to solve classically, test it, and validate it. This is something we've done in a lattice. We simulated it classically, we simulated it on the quantum device analogue, and we compared. We wanted to show that it worked. The answer was, of course, at first it didn't agree at all; it was just nonsense. Then we looked closer and realized




Slide 16 - Nature computes using quantum



Slide 17 - Feynman proposed to use quantum computers to simulate quantum physics



Slide 18 - Analog quantum simulators



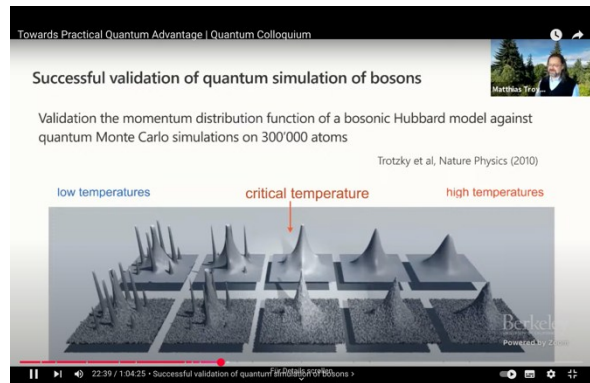
Slide 19 - Ultracold atomic gas quantum simulators

How Quantum Computing Escapes the Maze

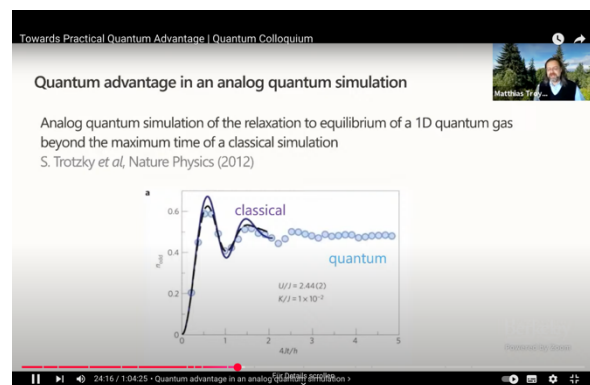
that it did not agree because we didn't really understand what was built. But once we understood how the quantum gear calculator worked, we got to this picture. **The back row shows the classical results; the front row the quantum results.** I have shown the negative distribution function of bosons in the lattice. Left is cold, right is hot. When it's hot, there's no feature. At low temperatures, when you cool down, you see at some point here the secondary peak appears, and that's where the Bose gas forms a so-called Bose-Einstein condensate. The agreement between the quantum mechanical simulation in the front row and the classical one is really great. So those systems work. We can build non-trivial quantum simulators that solve certain quantum models. Then you can go to problems that are hard classically. You can go to fermions or dynamics. Here's an example of a 10-year-old paper that shows a quench in a one-dimensional gas of bosons. The classical calculation here stops after a while, whereas the quantum simulation continues. Here, there's not much feature in there; it equilibrates basically.

00:23:55

But on the other hand, there could be models where dynamics can go to points that are classically hard but quantumly achievable. There is an advantage of quantum hardware over classical hardware that is analogue. **Just like classical machines that are analogue, it is limited.** One limitation is it is hard to cool it down to low temperatures. A second problem of anything analogue is precision, the calibration precision, the measurement precision one can get. You can calibrate it to a percent, a bit better, but not more. That is the same reason why in classical computing we switched from analogue to digital. Just because when it's analogue, we cannot really calibrate it well enough to scale, we cannot error-correct it. **But digitally, we can.**



Slide 20 – Successful validation of quantum simulation of bosons



Slide 21 - Quantum advantage in an analog quantum simulation

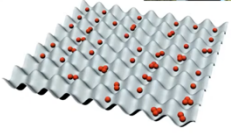
Slide 22 – Limitations of analog quantum simulations

The same has to happen quantumly. To go to a broader class of models, really scale up, and make it more precise, we need a gate model, a digital quantum computer that can be error-corrected. With that, I now want to see what is one of the simplest quantum models that I could simulate where a gate model machine could give me a quantum advantage. How can I do something that, if done in the analogue way, can be surpassed by doing it digitally? One of the simplest problems we've looked at, with the colleagues in Europe especially the groups of Andrew Daley, Peter Zoller, Immanuel Bloch and team members, was to look at how many quantum gates and qubits are needed to get to a point where it's classically hard but could be done analogue and what is needed in a gate model approach. In contrast to previous work, we didn't look at what we need to get the wave function to a certain fidelity. We focused more on properties we can measure. We asked: What do we need to get the measurements of certain local properties to the same point that they could do it in an experiment? We strongly relaxed the conditions, but we found that for it to be hard, I need to do much more than about 50 quantum particles. In this case, we looked at the quantum magnet in a transverse field Ising model. We said we need to look at least at a 10x10 lattice, so about 100 quantum spins, which could be about 100 qubits. That's where it gets hard classically. 64 qubits, I could still do classically approximately. 9x9 might get hard, so we could choose 9 or 10. I've chosen 10 here. Then I need to propagate it for a certain time, and the time has to be at least as long as the linear extent of the lattice, or the system would not even see the size. That's why I'm choosing a two-dimensional lattice and not a one-dimensional chain. In a 1D chain, I would have to propagate for a time of 100 to see the system size. Here, I'm choosing a 10x10 system and then propagating it for a time of 10. To describe it, I need just 100 qubits. To accurately simulate it with high fidelity, I would need a very small-time step. But what we found is if I relax the condition and say I just want to simulate it well enough that it is comparable to experiments and compared to the precision they could get from measurements in experiments, given an unavoidable calibration error, you can actually take pretty large time steps. We found that one can do the simulation with 100 qubits in just a few 10,000 gate operations. This is great news because this is far smaller than what was thought before, that with 30,000 gate operations we could get to the point where I'm doing something that is classically hard. But doing 30,000 gate operations still means the error on every single gate operation has to be less than one in a million, or the result will just be random noise. This is still way beyond the current state where I really need to do fault tolerance. We


What do we need to get to the simplest practical quantum advantage?

Quantum simulation of a simple classically intractable quantum spin model to the same precision as analog simulations

	Transverse field Ising model
Lattice sites	10 x 10
Propagation time	10
Qubits	100
Total number of gates	30,000



N. Pearson (ETHZ + IBM)
 G.H. Low, M. Troyer (Microsoft)
 S. Bravyi, A. Shor (IBM)
 A. J. Daley (Oxford)
 P. Zoller (Innsbruck)
 I. Bloch (MPQ Munich)



Powered by Zoom

Slide 23 – What do we need to get to the simplest practical quantum advantage?

How Quantum Computing Escapes the Maze

estimate we need about 100,000 physical qubits to do that. With 100,000 physical qubits and fault-tolerant simulations, we can solve quantum problems that are classically hard. That's not so far out; I think that can be achieved within the next years.

00:30:34

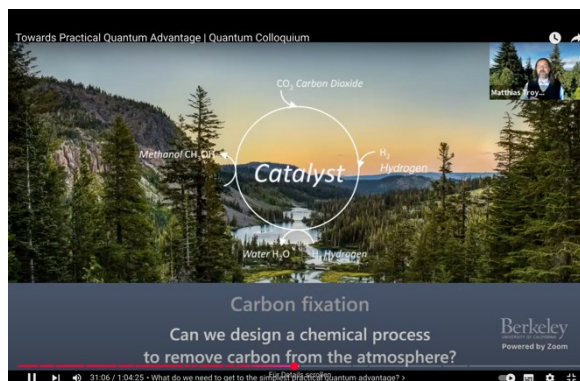
That can be achieved in the next few years. So let me now go from this simplest problem to a really interesting problem that is commercially impactful and see what is needed there. The problem is in chemistry, and I want to look at the problem of negative catalysis. We did a case study on a catalyst for negative carbon fixation. We want to, for example, find a material that helps us take CO₂ from the atmosphere and turn it into methanol.

00:31:32


In order to do that, one can either try out many different catalysts or we can try to predict how they work through calculation. For that, we have to go through the reaction cycle. On the left here, I have the catalyst reacting with the CO₂ molecule. It binds, goes through two steps, and then the catalyst binds again with the next molecule. For that to work, this reaction cycle has to be the dominant one. We don't want it to just get stuck with the catalyst plugging it up; we really want this cycle to happen. We want to be able to predict and calculate it before we start making the catalyst.

00:32:46

In order to do that, I have to look at all possible directions in which things could go, all possible product states, and all possible reaction pathways. Then I look at the energy along the reaction. As they bind and react, there's a barrier, and the rate for the reaction is exponential in the energy height of the barrier divided by the temperature. In order to know in which direction things react, I have to look at all possible paths, calculate the energies of these barriers, and then find out which barrier is lowest. Then I'm calculating the reaction rates, and I hope that the barriers on the path we want to achieve are the lowest.



Slide 24 - Can we design a chemical process to remove carbon from the atmosphere?



Slide 25 - Catalysis: carbon fixation

00:33:29

To do that, I need to know these energy differences. I need to know the barrier heights during the reactions, and they need to be accurate. They need to be accurate compared to the temperature because if the accuracy is much less than the temperature, then we really just don't know where it goes. It has to be at least the temperature. In this context, chemical accuracy, as mentioned in the tutorial, means I

Towards Practical Quantum Advantage | Quantum Colloquium

Determining eigenvalues through quantum phase estimation

Goal: Extract energy levels
 $\hat{H}|\psi_n\rangle = E_n|\psi_n\rangle$

Solving Schrödinger's equation for real time evolution
 $|\psi(t)\rangle = e^{-i\hat{H}t}|\psi(0)\rangle$

Use quantum computer to simulate time evolution
 Eigenstates accumulate phase $e^{-iE_n t}$
 Quantum phase estimation (QPE) for eigenvalues.

Energy
 — Excited states
 — Ground state

$|\psi_n\rangle$
 $\frac{|0\rangle + |1\rangle}{\sqrt{2}}$
 $e^{-iE_n t}$
 $\frac{|0\rangle + e^{-iE_n t}|1\rangle}{\sqrt{2}}$
 $|\psi_n\rangle$

Slide 26 - estimating eigenvalues through quantum phase estimation

have to be able to calculate the energy better than the temperature. The challenge here is that the temperature is an energy scale that's typically a million times smaller than the total energy of a molecule. To really be able to calculate the energy to six, seven, or eight digits of precision, many of the classical methods mentioned in the tutorial are not good enough.

00:34:42

That's where quantum computers come in. They come into the workflow where I'm exploring possible structures. For these structures, I want to solve the problem on quantum hardware, or I want to solve it in some way, but classically it's not good enough. Quantumly, I can't do it brute force because the molecule is too big. What one then does is solve that problem on quantum hardware and then feed it back to the classical calculation, look at pathways, find new candidates, and iterate. We want to do many calculations of accurate energies of the quantum system. To do that, we're using quantum phase estimation. We'd like to know the spectrum, especially the ground state. All I can do on the quantum hardware is simulate the dynamics of a quantum system.

Towards Practical Quantum Advantage | Quantum Colloquium

Quantum computing enhanced computational catalysis

Explore catalyst structures

Generate model parameters within active space

Obtain accurate energies using scalable quantum computer

Assemble free energy

Kinetic analysis on reaction pathways

New insights to modify structures

Slide 27 - Quantum computing enhanced computational catalysis

00:36:51

I want to implement the time evolution of a quantum system. If you combine that with phase estimation, by measuring the phase the wave function picks up, I can get the energy because the phase picked up by an eigenstate is proportional to the energy of the state at a time. Thus, by measuring the phase with phase estimation, I can get the energy of the state. It boils down

How Quantum Computing Escapes the Maze

ultimately to the time evolution of a quantum state under the Hamiltonian, the energy function that describes the molecule.

00:37:37

I've shown before how for the quantum magnet that was super easy. In the chemistry problem, it's a bit harder because the model is more complex. The reason why it's more complex is that the Coulomb interaction between electrons is long-range. When I have two electrons and write down the model for them, the number of terms in the model is (n^4) . From a complexity theory viewpoint, it's polynomial, but the scaling with (n^4) terms scares me, even classically. The question is: Is it really efficient in a practical sense? Can I really do it well? What will it take to do it, or are the constants so large that it might be hard?

Towards Practical Quantum Advantage | Quantum Colloquium

Simulating time evolution on quantum computers

There are $O(N^4)$ interaction terms in an N -electron system

$$H = \sum_{pq} t_{pq} c_p^\dagger c_q + \sum_{ppqr} V_{ppqr} c_p^\dagger c_q^\dagger c_r c_r \equiv \sum_{m=1}^M H_m \quad e^{-\Delta t H} = \prod_{m=1}^M e^{-\Delta t H_m}$$

Efficient quantum circuits available for each of the N^4 terms

37:55 / 1:04:25 • Simulating time evolution on quantum computers

Slide 28 – Simulating time evolution on quantum computers

00:38:47

When we first looked at that, it indeed looked hard. It looked bad when took an algorithm, we implemented it and estimated the runtime. With the assumptions we have now of a clock speed at the logical level of microseconds, it would be about a billion years. A billion years is beyond the few weeks that we want to wait; we can't wait that long. But then, with the advancements that we've made, we improved the implementations and brought the cost down. The reason why four years back we said it could be a month is that we assumed optimistically a future machine might be fast and might have a clock speed of nanoseconds. We now realize it's more than a microsecond.

00:40:36

So we say the runtime of those algorithms was about a century, much better than a billion years, but still too slow. In the last three years, we've brought it down further. Now we think that with

Towards Practical Quantum Advantage | Quantum Colloquium

10 orders of magnitude reduction in quantum algorithm runtime for computational catalysis over six years

Estimated runtimes on fast full-scale quantum computer architecture with around 5,000 logical qubits and logical cycle times of $10\mu\text{s}$

	Year	Expected runtime	Reference
Optimized implementation	2014	1 billion years	PRA 90, 022305 (2014)
	2017	100 years	PNAS 114, 7555 (2017)
Optimized algorithms	2020	1 month	arXiv:2007.14460

39:29 / 1:04:25 • 10 orders of magnitude reduction in quantum algorithm runtime for computational catalysis over six years

Slide 29 – 10 orders of magnitude reduction in quantum algorithm runtime for computational catalysis over six years

about 5,000 logical qubits, we can do it in roughly a month. I want to mention two of the advances that led to that. The first advance is rethinking the problem we want to solve. We want to know the energies and the wave functions, so the eigenstates.

00:41:20

For that, they want to time evolve with the model, the Hamiltonian, and that just takes time. One breakthrough here was realizing that if I don't care about the exact dynamics but care about the spectrum, then I can do the calculations for any function of the Hamiltonian. It turns out that the time evolution under the negative arc sine of HH is much more efficient to implement than under HH itself. A big step forward here was realizing that I can propagate with a certain function of HH that makes it more efficient to get the spectral information, and so the runtime comes down per time step.

Towards Practical Quantum Advantage | Quantum Colloquium

Progress by algorithmic breakthroughs

Eigenvalue problem: $H|\psi_n\rangle = E_n|\psi_n\rangle$

Two advances contributed

1. More efficient quantum circuit available for $f(H) = \sin^{-1}(H)$
 $f(H)|\psi_n\rangle = f(E_n)|\psi_n\rangle$ encodes same spectral information!
2. Exploit sparse representations of H for fewer terms and larger time steps

Ruthenium catalyst configuration VIII with 130 spin-orbitals		
Approach	# Time steps	Overall algorithmic speedup
Unfactorized	10,600	1.0x
Single-factorized	42,200	0.4x
Our optimized algorithm	570	18.5x

41:26 / 1:04:25 • 10 orders of magnitude reduction in runtime for computational catalysis over...

Slide 30 – Progress by algorithmic breakthroughs

00:42:09

The next thing we do is we have n^4 terms, but the question is: Do we really need that many? The next breakthroughs over the last years involved looking at sparse representation. Can I look at those terms and write down a sparse representation, a sparse model for it, that has fewer terms but still maintains the same accuracy? That worked. Not only did the number of terms go down, but we also can make the time steps larger because of that. We can do fewer time steps, and each one is faster. That let us bring down the cost.

00:43:01

I think we're slowly reaching limits here because we're really thinking about what is the minimal information needed to describe the chemistry problem, and it's still a small one. There are still more breakthroughs needed here, and I believe the next breakthroughs really have to come from chemistry. Let me look at a second example to illustrate why. I want to look at a problem here again, a catalyst for nitrogen fixation. There are

Towards Practical Quantum Advantage | Quantum Colloquium

Berkeley FemtoChem
Powered by Zoom

44:16 / 1:04:25 • 10 orders of magnitude reduction in runtime for computational catalysis over...

Slide 31

certain enzymes in microbes in the soil that can fix nitrogen, which is the basis for fertilizers. To understand how it works, you would, in principle, have to look at that big enzyme with about a

How Quantum Computing Escapes the Maze

thousand electrons. That problem, with the scaling, is just far too big to ever do it brute force quantumly. It's super big, even classically.

00:44:48

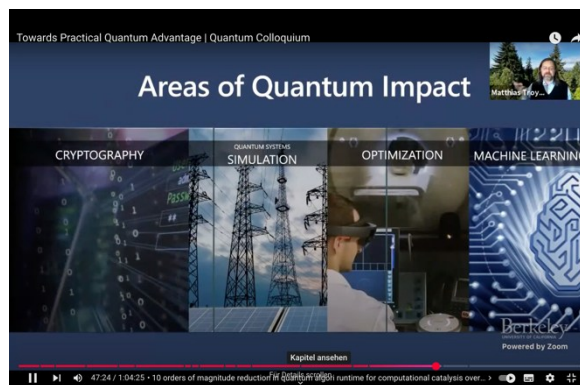
What one then does is look at the active centre, the core of it. There is this cofactor that I mentioned earlier. **We want to focus just on the centre of it and just add a small active space of a few interesting orbitals.** As we do that, we have a problem because we are looking at the big molecule and then making a model just for this tiny active centre, really truncating the model to just a tiny bit of the molecule and throwing out the rest. The problem one has there is that this restriction to the active space, to the centre, incurs an uncontrolled error.

00:45:34

We are using the quantum computer to precisely and accurately calculate the energy of a model that itself has an uncontrolled error that might be bigger than the temperature. The problem here is really how do I do this embedding? How do I simulate this big molecule accurately if I can only look at a tiny piece of it, do that quantumly exactly, but how do I now make sure that the truncation error of getting to that problem does not make the exactness of the quantum simulation irrelevant? We have to solve this problem with interactions with the rest of the system.

00:46:13

That is a problem that we need to start solving now because unless we solve it, the exact energy from quantum hardware will not really take us all the way. The good news is being motivated by this and starting to work on solving that problem in new ways because of the motivation from quantum can help us classically even before we have the thousands of qubits needed to solve the chemistry problem. I want to end by showing one area where the same thing happened. When we talk about quantum and quantum impact, cryptography is one simulation, but you also, for a long time, talked about optimization.

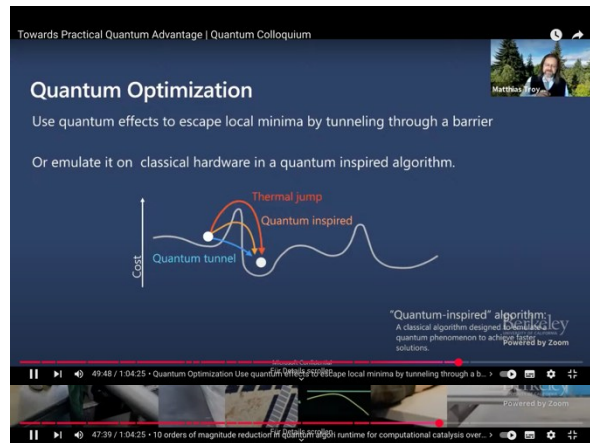


Slide 32 - Areas of quantum impact

How Quantum Computing Escapes the Maze

00:47:07

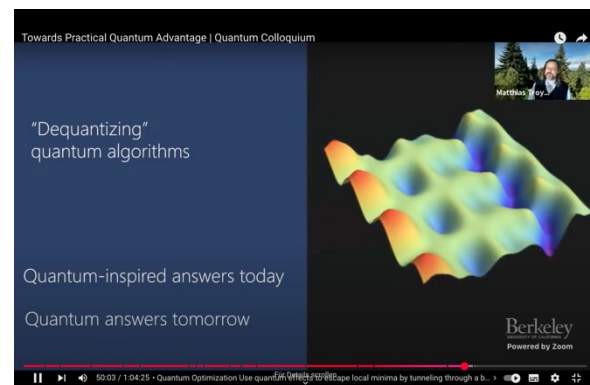
There are many different hard problems in optimization: oil and gas, material discovery, power grids, signal processing, machine learning, scheduling, logistics, traffic, and more. There are loads of problems where we want to solve a hard optimization problem. The hope was that using quantum effects; one can escape from a local minimum. Instead of climbing over the barrier, one can just quantum tunnel through it. That was the motivation for quantum annealing. There have been cases where quantum tunnelling is exponentially better than climbing over the barrier.



Slide 33 – Quantum Optimization

00:48:56

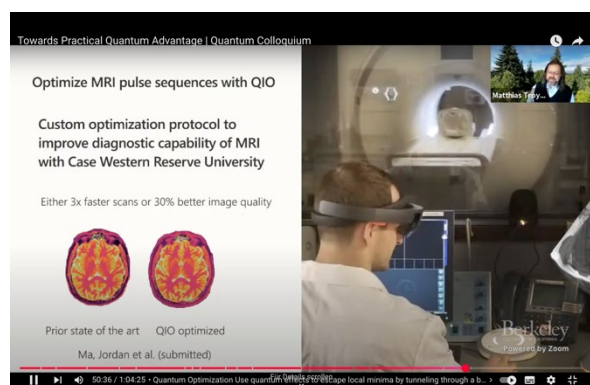
What we realized over the last decade, looking at quantum annealers, is that you can mimic it classically. You can simulate it efficiently classically, and thus you can mimic the quantum tunnelling on a classical computer. What that led to is a quantum-inspired algorithm where basically one takes the ideas of the quantum algorithm, dequantizes it, and turns the quantum algorithm into a related classical algorithm. With Monte Carlo sampling, it gives us much of the speedup from the quantum algorithm on classical hardware.



Slide 34 – “Dequantizing” quantum algorithms

00:49:49

One example of thinking about quantum algorithms and how we use them gave us new ideas. One example was sequences that reduce the scan time by a factor of three, which is a big breakthrough. Similarly, I predict that the breakthroughs we make in chemistry by thinking about quantum will lead to breakthroughs there even before we have the quantum hardware scaled up. To end now, what I've been trying to show today is that to really get quantum advantage



Slide 35 – Optimize MRI pulse sequences with QIO

How Quantum Computing Escapes the Maze

for a practical problem in the next decades, we need to focus on a problem that has small data but big compute problems. To really beat the classical hardware we have, we need to look for problems with more than quadratic speedup. Even the simplest application problems will need fault tolerance and hundreds of thousands of qubits to beat the best classical algorithms.

Once we get to millions, there are loads of chemistry problems that will solve interesting ones like carbon fixation and more. That's where they can help. Even on the way there, in the decade it takes us to get there, there will be breakthroughs in classical methods just from tackling the classical chemistry problems we need to solve to make the quantum hardware useful. So, on the way to building the million-qubit machine, we will have lots of progress from quantum computing on classical hardware. Thank you.

00:52:05

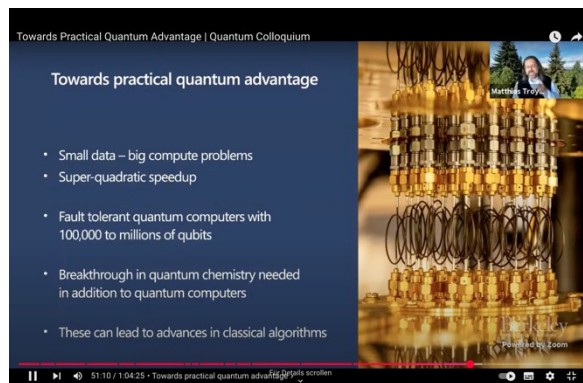
[Host:] Great, thanks Matthias for a great talk. Actually, it was very interesting to see that your very pragmatic view about quantum advantage in quantum algorithms coincided with the complexity theoretic view where we always look for super-polynomial or exponential speedups rather than polynomial speedups, which we treat as being somewhat suspicious. It's actually very interesting that the two ends of this spectrum agree more than the middle. So, let's open it up to some questions before we get to the panel.

00:53:30

Let's see if any of the moderators have questions. I actually had a question in the meantime. At the beginning of your talk, you had these calculations saying that the clock speed of a quantum computer would necessarily be many orders of magnitude slower than a classical computer. Is that a practical consideration currently, or is there something inherently true about it even looking somewhat far into the future?

00:54:24

[Troyer:] I think this can likely be improved with breakthroughs in error-correcting codes, but I think it is fundamental. The reason is that at the lowest level, I have physical qubits, and I have to control the physical qubits. I have to do the physical gate operations, and that needs a classical computer to control and manipulate the quantum state. If you then on top of that do the syndrome measurements, the main cycles, things slow down. For a surface code, I need a code distance of



Slide 36 – Towards practical quantum advantage

maybe about 20 or 30, and each cycle is about 10 gates to implement. So, I have a few hundred gate operational cycles that I do, and thus I am at least a factor of a few hundred slower than the classical computer that controls the physical gate operations. Then I also have the space overhead of quantum error correction. If I take a code distance of 20 in a surface code, it means I need about a thousand physical qubits for a logical qubit.

00:55:45

Okay, it might have more physical qubits than logical qubits, but with that, I'm about a factor of a hundred thousand. Then the qubit is just much more complex than a transistor, so the classical logic needed to actually control the qubit is a factor of 100,000 bigger. So, then I have my next factor of 100. Just from simple applications, I see that that's really illuminating. Then there's the size on that chip. I can really squeeze tiny transistors; I don't know if you can make the qubits so small and fit that many on a single wafer. That's the next factor of a thousand, which is kind of the density, or the size of the physical qubit compared to the size of the transistor. I think if the breakthroughs really come on the classical side with classical control, then it will profit classical computing as well. The breakthrough has to come in either much better qubits or in new quantum error correction codes that could maybe give us a breakthrough, or both. That's how we can shave off that factor.

00:57:19

[Questioner named Steve:] Great. Can I ask a question? Yes, please. Just to make sure I understood your discussion of the transverse field Ising model, I want to make sure I got the basics right. You were saying that a 10x10 system with 100 qubits and 30,000 gate operations—did I understand correctly that you thought that was within reach on the scale of a year? [Troyer:] No, no, no. Okay, so I said this is something that is just beyond what I can do classically.

00:57:59

I don't think I can do it accurately to that precision of about a percent for the prediction over that time on a classical computer. Any classical algorithm that can get us there? [Questioner named Steve:] No. My question was on a NISQ device. [Troyer:] On a NISQ device, if I need to do 30,000 rotation gates accurately and I want to have results accurate to 99 percent, then I need a fidelity on my qubits that is about seven nines on the gate operations. Because I want to do about 50,000 gate operations, the error rate should be less than one percent. So, I need an error per gate operation of less than one in a few million. I'm saying that NISQ machines don't currently have any roadmap towards that; we need fault tolerance.

How Quantum Computing Escapes the Maze

00:58:40

[Questioner named Steve:] Okay, maybe I should have asked the question more stupidly. What do you think might be possible with NISQ devices and the transverse field IC model? Let me ask the question that way. [Troyer:] NISQ devices on the transverse field IC model on a size that I can't simulate classically. Yes, I know you love the model. Let's take 100 spins. Then per time step, you need to have at least one rotation gate per term. So, per time step, we need about 300 or about a thousand gate operations for one time step. With three nines fidelity, that means I can do one time step and things start to randomize. One time step we can still do classically, but when you go for ten-time steps and you can choose a big-time step, then it becomes hard classically. So, I'd say on NISQ devices, I want to change my answer and say if you insist on using it like a digital machine, if you're using an analogue quantum system and now you insist that you only have the discrete gate operations of a digital computer, then you combine the disadvantages of analogue being noisy and digital being restricted. What you should actually do is take the analogue machine and use it as an analogue machine, and then you can do loads. If you use the qubit and just view it as an analogue quantum simulator, then there are many people here in the audience who've done that for decades, and there you can do a lot. That was even done recently by D-Wave on their systems. When they started to look at their machine as an analogue simulator for an Ising model, they can do interesting work there. But when you look at the NISQ machine, that's my point of view now, as an analogue machine and use it like a gate model machine, then you are handicapping the device. The discrete gates become useful once you have fault tolerance; you need it for that.

01:01:24

NISQ devices have already for a long time solved nice IC models in an analogue way. You don't need to wait a year; you can have it now. [Host:] Quickly, I think Ike Trung has a question. [Ike Trung:] Thanks, Sesh, and thank you, Matthias, for a wonderful talk. My question is about the second half of your talk. You mentioned having breakthroughs for maybe devices and quantum error correction, but how about breakthroughs based on asking better questions? Specifically, I noticed that you're still focusing on eigenvalue estimation using phase estimation. But in your beautiful work with Microsoft and Guangzhao and so forth on this new singular value transform algorithm, it allows you to ask more sophisticated questions like polynomial transforms of the eigenvalue spectrum, as you know very well. Can you think of better questions that take advantage of that so that you ask a big quantum question and then transform the system so you can answer that all at once?

01:02:40

[Troyer:] That's it. Here, basically, we have transformed the problem using those tricks of

Guangzhao and others to say we can transform the problem to make it easier to solve. But then we still looked at the energy of the spectrum. The bigger question is: Are there any practical applications for quantum algorithms that we have not thought about yet? [Ike Trung:] Or chemically relevant polynomial transformed questions like reaction rates or thresholds or asking about ranges of eigenvalues instead. I know this is an extended question, and I don't want to get into the panel time too long, but maybe if you have a brief answer.

01:03:19

[Troyer:] The brief answer is basically for reaction rates; the simulator often boils down to energies. For energies, you run into the Nyquist limit, so the time you need to do something to get it to a certain level is bounded from below. That's likely the answer. Without that, you could maybe avoid Heisenberg. There might be a way to reformulate that in the question: Can we prove a lower bound on that? Or can we, by breaking through lower bounds in some smart way, have ideas on how to get the energies better?

APPENDIX II – QUESTIONS TO LARGE LANGUAGE MODELS ON QUANTUM COMPUTING

This appendix documents the types of questions posed to large language models (Microsoft 365 Copilot and Anthropic's Claude Sonnet) to facilitate comprehension of technical concepts encountered during the analysis of Microsoft's quantum computing narratives. These questions were designed to help a social scientist without extensive background in quantum physics understand the theoretical foundations underlying the documents analysed in this thesis. The transcript of Troyer's colloquium (Appendix I) was uploaded to the LLMs for reference.

Note on Usage: These prompts were used iteratively, with follow-up questions seeking clarification, requesting references, and cross-checking explanations between different LLM systems. The LLMs functioned as technical dictionaries helping understand *what* the documents were saying so that STS analysis could examine *how* and *why* they were saying it. All explanations were treated as preliminary interpretations rather than definitive accounts, and key technical claims were verified against published scientific literature where possible.

The complete transcript of interactions with 365 Copilot and Claude Sonnet is available upon request from the author.

Prompt 1: Quantum Superposition

"I am an STS scholar analysing corporate narratives about quantum computing. Please explain the concept of quantum superposition in terms accessible to a social scientist. What does it mean for a qubit to be in multiple states simultaneously, and why is this significant for computation? Please provide references to foundational physics literature."

Prompt 2: Quantum Entanglement

"For my thesis on quantum computing narratives, I need to understand quantum entanglement. Please explain what entanglement means physically, why Einstein called it 'spooky action at a distance,' and how it enables quantum computing capabilities that classical computers lack. Use analogies appropriate for someone without physics training."

Prompt 3: Decoherence and Error Correction

"In analysing Matthias Troyer's 2021 colloquium, he discusses the need for 'fault tolerance' in quantum computing. Please explain what quantum decoherence is, why it poses such a fundamental challenge, and what fault-tolerant quantum computing means. Why does this require many more physical qubits than logical qubits?"

Prompt 4: Majorana Particles

"Microsoft's quantum computing strategy centres on 'Majorana particles' or 'Majorana zero modes.' Please explain what these are, why they are called 'quasi-particles,' how they differ from ordinary particles, and why their observation remains scientifically controversial. I need to understand the technical basis for the claims Microsoft makes in their 2025 announcement."

Prompt 5: Topological Protection

"Microsoft claims their approach uses 'topological protection' for quantum information. Please explain what topology means in this context, how the 'donut with a hole' metaphor relates to protecting quantum states, and why this approach might be more robust against environmental interference than other qubit technologies."

Prompt 6: Topoconductors

"In the 2025 Majorana 1 announcement, Microsoft introduces the term 'topoconductor' and claims it creates 'an entirely new state of matter.' Please explain what this means technically, how topoconductors relate to superconductors and semiconductors, and whether this claim of a 'new state of matter' is scientifically conventional or extraordinary."

Prompt 7: NISQ Devices

"Troyer dismisses 'NISQ' (Noisy Intermediate-Scale Quantum) devices as combining 'the disadvantages of analogue being noisy and digital being restricted.' Please explain what NISQ devices are, why companies like IBM and Google have pursued this approach, and what Troyer means by his characterisation. What is the technical debate between NISQ and fault-tolerant approaches?"

Prompt 8: Superconducting Qubits vs Topological Qubits

"Please compare the superconducting qubit approach (used by IBM and Google) with Microsoft's topological qubit approach. What are the technical trade-offs? Why might one approach be 'faster to demonstrate' but potentially 'limited in scalability' as Microsoft claims?"

Prompt 9: Quantum Annealing

"Troyer mentions quantum annealers and 'quantum tunnelling' in the context of optimisation problems. Please explain what quantum annealing is, how it differs from universal quantum computing, and why Troyer argues that classical algorithms can 'mimic' quantum tunnelling effects."

Prompt 10: Grover's Algorithm and Quadratic Speedup

"Troyer systematically dismisses applications relying on 'Grover speedup' as insufficient for practical quantum advantage. Please explain what Grover's algorithm does, what 'quadratic speedup' means mathematically, and why Troyer argues this is insufficient given the 'constant slowdown' of quantum versus classical operations."

Prompt 11: Quantum Chemistry Simulation

"Troyer identifies quantum chemistry as the primary viable application for quantum computers. Please explain why simulating molecular behaviour is computationally difficult for classical computers, what 'the electron correlation problem' means, and why quantum computers might have an inherent advantage for such simulations."

Prompt 12: The I/O Problem

"Troyer argues that quantum computers cannot solve 'big data' problems due to input/output constraints. Please explain the technical basis for this argument. What are the physical limitations on how quickly data can be read into and out of a quantum computer, and why does this matter for machine learning applications?"

Prompt 13: Quantum Measurement

"For my analysis of Microsoft's claims about 'measuring Majorana particles,' please explain how quantum measurement works. What does it mean to 'measure' a quantum state, why does measurement collapse superposition, and what are the challenges in detecting Majorana zero modes specifically?"

Prompt 14: Scientific Controversy Over Majorana Detection

"Nature reports that 'some physicists are sceptical' of Microsoft's Majorana claims and that 'evidence falls short of being definitive.' Please explain what alternative physical mechanisms might produce similar experimental signatures to Majorana zero modes, and why this controversy matters for evaluating Microsoft's breakthrough claims."

Prompt 15: The Million Qubit Threshold

"Microsoft claims that quantum computing requires 'a path to a million qubits' for commercial viability. Please explain why this specific scale is significant. What is the relationship between physical qubits and logical qubits in fault-tolerant systems? Why does Troyer state that even the 'simplest application problems' require 'hundreds of thousands of qubits'?"

REFERENCES

- Aghaee, M., Alcaraz Ramirez, A., Alam, Z., Ali, R., Andrzejczuk, M., Antipov, A., Astafev, M., Barzegar, A., Bauer, B., Becker, J., Bhaskar, U. K., Bocharov, A., Boddapati, S., Bohn, D., Bommer, J., Bourdet, L., Bousquet, A., Boutin, S., Casparis, L., . . . Microsoft Azure Quantum (2025). Interferometric single-shot parity measurement in InAs–Al hybrid devices. *Nature*, 638(8051), 651–655. <https://doi.org/10.1038/s41586-024-08445-2>
- Arute, F., Arya, K., Babbush, R., Bacon, D., Bardin, J. C., Barends, R., Biswas, R., Boixo, S., Brandao, Fernando G. S. L., Buell, D. A., Burkett, B., Chen, Y., Chen, Z., Chiaro, B., Collins, R., Courtney, W., Dunsworth, A., Farhi, E., Foxen, B., . . . Martinis, J. M. (2019). Quantum supremacy using a programmable superconducting processor. *Nature*, 574(7779), 505–510. <https://doi.org/10.1038/s41586-019-1666-5>
- Asdal, K., & Reinertsen, H. (2022). *Doing document analysis: A practice-oriented method*. SAGE Publications Ltd; SAGE Publications Inc.
- Austin, J. L. (1962). *How to do things with words*. Oxford University Press.
- Bakker, S., & Budde, B. (2012). Technological hype and disappointment: lessons from the hydrogen and fuel cell case. *Technology Analysis & Strategic Management*, 24(6), 549–563. <https://doi.org/10.1080/09537325.2012.693662>
- Blunt, N. S., Camps, J., Crawford, O., Izsák, R., Leontica, S., Mirani, A., Moylett, A. E., Scivier, S. A., Sünderhauf, C., Schopf, P., Taylor, J. M., & Holzmann, N. (2022). Perspective on the Current State-of-the-Art of Quantum Computing for Drug Discovery Applications. *J. Chem. Theory Comput*, 18, 7023. <https://doi.org/10.1021/acs.jctc.2c00574>
- Bolgar, C. (2025). *Microsoft's Majorana 1 chip carves new path for quantum computing - Source*. Microsoft Corporation. <https://news.microsoft.com/source/features/innovation/microsofts-majorana-1-chip-carves-new-path-for-quantum-computing/>
- Borup, M., Brown, N., Konrad, K., & van Lente, H. (2006). The sociology of expectations in science and technology. *Technology Analysis & Strategic Management*, 18(3-4), 285–298. <https://doi.org/10.1080/09537320600777002>
- Boyle, A. (2023, February 2). DARPA boosts Microsoft's far-out effort to build topological quantum computer. *GeekWire*. <https://www.geekwire.com/2023/darpa-microsoft-topological-quantum-computer/>

How Quantum Computing Escapes the Maze

- Brown, N., & Michael, M. (2003). A Sociology of Expectations: Retrospecting Prospects and Prospecting Retrospects. *Technology Analysis & Strategic Management*, 15(1), 3–18. <https://doi.org/10.1080/0953732032000046024>
- Callon, M. (1998). The Laws of the Markets. *The Sociological Review*, 46(S1), 1–57. <https://doi.org/10.1111/j.1467-954X.1998.tb03468.x>
- Cao, Y., Romero, J., Olson, J. P., Degroote, M., Johnson, P. D., Kieferová, M., Kivlichan, I. D., Menke, T., Peropadre, B., Sawaya, N. P. D., Sim, S., Veis, L., & Aspuru-Guzik, A. (2019). Quantum Chemistry in the Age of Quantum Computing. *Chem. Rev*, 119, 10915. <https://doi.org/10.1021/acs.chemrev.8b00803>
- Castelvecchi, D. (2021). Evidence of elusive Majorana particle dies - but computing hope lives on. *Nature*, 591(7850), 354–355. <https://doi.org/10.1038/d41586-021-00612-z>
- Castelvecchi, D. (2025). Microsoft claims quantum-computing breakthrough - but some physicists are sceptical. *Nature*, 638(8052), 872. <https://doi.org/10.1038/d41586-025-00527-z>
- Collins, H. (1992/1985). *Changing order: Replication and induction in scientific practice* (1 [print.]). Univ. of Chicago Pr.
- Czarniawska, B. (2004). *Narratives in social science research. Introducing qualitative methods*. Sage Publications.
- Deshpande, A. (2022). Assessing the quantum-computing landscape. *Communications of the ACM*, 65(10), 57–65. <https://doi.org/10.1145/3524109>
- Feynman, R. P. (1982). Simulating physics with computers. *International Journal of Theoretical Physics*, 21(6-7), 467–488. <https://doi.org/10.1007/BF02650179>
- Flick, U. (2014). *An introduction to qualitative research* (Edition 5). Sage.
- Frolov, S., & Mourik, V. (2021). *Data manipulation and omission in 'Quantized Majorana conductance', Zhang et al, Nature 2018*. <https://zenodo.org/records/4587841>
- Garisto, D. (2025). Microsoft quantum computing 'breakthrough' faces fresh challenge. *Nature*, 639(8055), 555–556. <https://doi.org/10.1038/d41586-025-00683-2>
- Gieryn, T. F. (1983). Boundary-Work and the Demarcation of Science from Non-Science: Strains and Interests in Professional Ideologies of Scientists. *American Sociological Review*, 48(6), 781–795.
- Hilgartner, S. (2015). Capturing the imaginary. In S. Hilgartner, C. A. Miller, & R. Hagendijk (Eds.), *Genetics and society. Science and democracy: Making knowledge and making power in the biosciences and beyond*. Routledge.
- Holter, C. ten [C.], Inglesant, P [P.], Pijselman, M., & Jirotko, M [M.]. (2024). *Towards responsible quantum computing*. EY; University of Oxford.

- <https://www.ey.com/content/dam/ey-unified-site/ey-com/en-uk/newsroom/2024/06/ey-oxford-uni-whitepaper-quantum-ethics-05-2024.pdf>
- Inglesant, P [Philip], Holter, C. ten [Carolyn], Jirotko, M [Marina], & Williams, R. (2021). Asleep at the wheel? Responsible Innovation in quantum computing. *Technology Analysis & Strategic Management*, 33(11), 1364–1376.
<https://doi.org/10.1080/09537325.2021.1988557>
- Jasanoff, S. (2005). *Designs on nature: Science and democracy in Europe and the United States*. Princeton Univ. Press.
<http://www.loc.gov/catdir/enhancements/fy0704/2004055296-d.html>
- Jasanoff, S., & Kim, S.-H. (2009). Containing the atom: Sociotechnical imaginaries and nuclear power in the United States and South Korea. *Minerva*.
- Jasanoff, S., & Kim, S.-H. (Eds.). (2015). *Dreamscapes of modernity: Sociotechnical imaginaries and the fabrication of power*. The University of Chicago Press.
- Jong, E. de (2022). Own the Unknown: An Anticipatory Approach to Prepare Society for the Quantum Age. *Digital Society*, 1(2). <https://doi.org/10.1007/s44206-022-00020-4>
- Knorr Cetina, K. (1995). Metaphors in the scientific laboratory: why are they there and what do they do? In Z. Radman (Ed.), *From a metaphorical point of view: A multidisciplinary approach to the cognitive content of metaphor* (pp. 329–349). de Gruyter.
- Knorr Cetina, K. (1999). *Epistemic Cultures: How the science make knowledge*. Harvard University Press.
- Knorr Cetina, K. (2010). The couch, the cathedral, and the laboratory: on the relationship between experiment and laboratory in science. In *Science as practice and culture*. University of Chicago Press. <https://d-nb.info/1103026992/34>
- Kop, M. (2023). Quantum-ELSPI: A Novel Field of Research. *Digital Society*, 2(2).
<https://doi.org/10.1007/s44206-023-00050-6>
- Kop, M., Aboy, M., Jong, E. de, Gasser, U., Minssen, T., Cohen, I. G., Brongersma, M., Quintel, T., Floridi, L., & Laflamme, R. (2024). Ten principles for responsible quantum innovation. *Quantum Science and Technology*, 9(3). <https://doi.org/10.1088/2058-9565/ad3776>
- Latour, B. (1986). Visualisation and Cognition: Drawing Things Together. *Knowledge and Society Studies in the Sociology of Culture Past and Present*, Jai Press, 6, 1–40.
- Latour, B. (1994/1987). *Science in action: How to follow scientists and engineers through society* (6. printing). Harvard Univ. Press.
- Latour, B., & Woolgar, S. (1986/1979). *Laboratory life: The construction of scientific facts*. Princeton University Press.

How Quantum Computing Escapes the Maze

- MacKenzie, D. A. (2006). *An engine, not a camera: how financial models shape markets*. MIT Press.
- Matthews, W. H. (2014/1922). *Mazes and Labyrinths. A general account of their history and development ... With illustrations [and a bibliography.]*. Longmans & Co.
- Microsoft. (2025, November 22). *Quantum Ready*. Microsoft.
<https://quantum.microsoft.com/en-us/quantum-ready/get-started>
- Microsoft Corporation. (2025, February 19). *Introducing Microsoft Majorana 1*.
<https://news.microsoft.com/azure-quantum/>
- Nguyen, A. (2025). Export Controls as Innovation Marketing? Sociotechnical Imaginaries in the Ringfencing of Quantum Technologies. *Law, Technology and Humans*, 7(1), 68–83. <https://doi.org/10.5204/lthj.3777>
- NIST. (2024, December 13). *NIST Releases First 3 Finalized Post-Quantum Encryption Standards*. NIST. <https://www.nist.gov/news-events/news/2024/08/nist-releases-first-3-finalized-post-quantum-encryption-standards>
- O'Shea, D. (2022). *IQT San Diego recap: Stay humble, avoid the quantum hype - Inside Quantum Technology*. <https://www.insidequantumtechnology.com/news-archive/iqt-san-diego-recap-stay-humble-avoid-the-quantum-hype/>
- Owen, R [R.], Macnaghten, P., & Stilgoe, J. (2012). Responsible research and innovation: From science in society to science for society, with society. *Science and Public Policy*, 39(6), 751–760. <https://doi.org/10.1093/scipol/scs093>
- Perrier, E. (2022). The Quantum Governance Stack: Models of Governance for Quantum Information Technologies. *Digital Society*, 1(3). <https://doi.org/10.1007/s44206-022-00019-x>
- Preskill, J. (2018). Quantum Computing in the NISQ era and beyond. *Quantum*, 2, 79. <https://doi.org/10.22331/q-2018-08-06-79>
- René von Schomberg. (2013). A Vision of Responsible Research and Innovation. In *Responsible innovation: Managing the responsible emergence of science and innovation in society*. John Wiley & Sons Inc.
- Rietsche, R., Dremel, C., Bosch, S., Steinacker, L., Meckel, M., & Leimeister, J.-M. (2022). Quantum computing. *Electronic Markets*, 32(4), 2525–2536. <https://doi.org/10.1007/s12525-022-00570-y>
- Roberson, T. M. (2021). On the Social Shaping of Quantum Technologies: An Analysis of Emerging Expectations Through Grant Proposals from 2002–2020. *Minerva*, 59(3), 379–397. <https://doi.org/10.1007/s11024-021-09438-5>

- Selin, C. (2008). The Sociology of the Future: Tracing Stories of Technology and Time. *Sociology Compass*, 2(6), 1878–1895. <https://doi.org/10.1111/j.1751-9020.2008.00147.x>
- Shor, P. W. (1994). Algorithms for quantum computation: Discrete logarithms and factoring. In *Foundations of Computer Science, 35th Symposium on (FOCS '94)* (pp. 124–134). IEEE Comput. Soc. Press.
- Soller, H., Nguyen, D. N., & Gschwendtner, M. (2025). *Quantum technology investment opportunities*. McKinsey & Company. <https://www.mckinsey.com/capabilities/tech-and-ai/our-insights/tech-forward/quantum-technology-investment-hits-a-magic-moment>
- Spielberg, S. (Director). (1975). *Jaws*.
- Swayne, M. (2024, December 31). 2025 Expert Quantum Predictions — Quantum Computing. *The Quantum Insider*. <https://thequantuminsider.com/2024/12/31/2025-expert-quantum-predictions-quantum-computing/>
- Troyer, M. (2021, March 16). Towards Practical Quantum Advantage | Quantum Colloquium, *Simmons Institute*. <https://www.youtube.com/watch?v=WY3htdKUGsA>
- UCSB. (2025, April 7). *Microsoft Research: Station Q at UCSB and the future of computing: Moore did not predict the power of a qubit! | UCSB Engineering Industry Center*. <https://industry.ucsb.edu/content/microsoft-research-station-q-ucsb-and-future-computing-moore-did-not-predict-power-qubit>
- van Lente, H., & Rip, A. (1998a). Chapter 7. Expectations in Technological Developments: An Example of Prospective Structures to be Filled in by Agency. In C. Disco & B. van der Meulen (Eds.), *Getting New Technologies Together*. de Gruyter. <https://doi.org/10.1515/9783110810721.203>
- van Lente, H., & Rip, A. (1998b). The Rise of Membrane Technology: From Rhetorics to Social Reality. *Social Studies of Science*, 28(2), 221–254.
- Warnick, J. (2014). *Station Q*. Microsoft Corporation. <https://news.microsoft.com/stories/stationq/>
- Wimmer, M., & Moraes, T. G. (2022). Quantum Computing, Digital Constitutionalism, and the Right to Encryption: Perspectives from Brazil. *Digital Society*, 1(2). <https://doi.org/10.1007/s44206-022-00012-4>
- Wolf-Bauwens, M., & Mandelbaum, R. (2024, January 16). *Defining responsible quantum computing: IBM Quantum Computing Blog*. IBM. <https://www.ibm.com/quantum/blog/responsible-quantum>

How Quantum Computing Escapes the Maze

Zhang, H., Liu, C.-X., Gazibegovic, S., Xu, D., Logan, J. A., Wang, G., van Loo, N., Bommer, J. D. S., Moor, M. W. A. de, Car, D., Op het Veld, Roy L. M., van Veldhoven, P. J., Koelling, S., Verheijen, M. A., Pendharkar, M., Pennachio, D. J., Shojaei, B., Lee, J. S., Palmstrøm, C. J., . . . Kouwenhoven, L. P. (2018). Quantized Majorana conductance. *Nature*, 556(7699), 74–79.

<https://doi.org/10.1038/nature26142>

Zhang, H., Liu, C.-X., Gazibegovic, S., Xu, D., Logan, J. A., Wang, G., van Loo, N., Bommer, J. D. S., Moor, M. W. A. de, Car, D., Op het Veld, Roy L. M., van Veldhoven, P. J., Koelling, S., Verheijen, M. A., Pendharkar, M., Pennachio, D. J., Shojaei, B., Lee, J. S., Palmstrøm, C. J., . . . Kouwenhoven, L. P. (2021). Retraction Note: Quantized Majorana conductance. *Nature*, 591(7851), E30.

<https://doi.org/10.1038/s41586-021-03373-x>