

---

# TOWARDS A METHODOLOGY FOR PROSPECTIVE GOVERNANCE IN QUANTUM COMPUTING TECHNOLOGIES

---

*Carolyn Ten Holter*

Linacre College

supervised by

*Professor Marina Jirotko*

Department of Computer Science

University of Oxford

*A thesis submitted for the degree of Doctor of Philosophy*

Hilary Term 2023





## Acknowledgements

First and foremost I wish to profoundly thank Professor Marina Jirotko as both supervisor and mentor – for guidance and advice, and for the countless opportunities and encouragement she provided to extend my reach and lean out of my comfort zone. Without her I would not be where I am in my career.

Colleagues, students and friends at 39a St Giles, Oxford, with whom I've shared tea, cake, games nights, and plant-watering rotas. In particular, my thanks to Dr Helena Webb, who recklessly offered to read draft chapters.

The Linacre Ladies that Lift have been a source of strength (both figurative and literal) and friendship throughout this work. I would especially like to express my heartfelt appreciation to Professor Catherine Walter for commentary, critique, and support.

I am also incredibly grateful to Professor Stephen Pinfield at the University of Sheffield, whose willingness to take a punt on me at the beginning of my Masters started me on this path.

Finally, my gratitude goes to my participants in this research – without their readiness to share their time, concerns, and expertise with me, there would be no thesis.



# Contents

Acknowledgements .....	3
------------------------	---

Abstract .....	11
----------------	----

## Chapter 1:

1.1: A growing challenge.....	13
-------------------------------	----

1.1.1: Multiple impacts.....	14
------------------------------	----

1.2: The limitations of current practice .....	16
--	----

1.2.1: Retrospective governance – the difficulties of regulation .....	16
--	----

1.2.2: Prospective governance – current approaches.....	17
---	----

1.2.3: Pre-empting negative consequences .....	18
--	----

1.3: Academia, industry and policy .....	19
--	----

1.4: Research questions and objectives .....	20
--	----

1.4.1: Research question .....	20
--------------------------------	----

1.4.2: Aims and objectives.....	21
---------------------------------	----

1.5: Contribution of the work .....	22
-------------------------------------	----

1.6: Methodology .....	23
------------------------	----

1.6.1: Qualitative methods – rationale .....	23
--	----

1.6.2: Jurisdictional limits.....	23
-----------------------------------	----

1.7: Outline of chapters .....	24
--------------------------------	----

## Chapter 2

2.1: Science, society and responsible innovation .....	27
--	----

2.2: RI and the AREA Framework .....	31
--------------------------------------	----

2.2.1: Anticipation .....	32
---------------------------	----

2.2.2: Reflection .....	34
-------------------------	----

2.2.3: Engagement.....	35
------------------------	----

2.2.4: Responsiveness .....	36
-----------------------------	----

2.3: RI in practice – difficulties and possible approaches .....	37
--	----

2.4: Institutionalisation? Embedding RI in research .....	44
---	----

2.5: RI in quantum computing .....	45
------------------------------------	----

2.5.1: The discourse of RI in quantum computing .....	46
2.5.2: Operationalising RI in quantum computing .....	49
<b>2.6: Conclusions .....</b>	<b>52</b>

## Chapter 3:

<b>3.1: Introduction .....</b>	<b>53</b>
<b>3.2: Background .....</b>	<b>53</b>
3.2.1: Societal effects .....	54
3.2.2: Research question .....	55
<b>3.3: Methodology.....</b>	<b>55</b>
3.3.1: Data collection.....	55
3.3.2: Selection and recruitment of participants .....	56
3.3.3: Limitations .....	57
3.3.4: Interview technique .....	57
3.3.5: Data management.....	58
3.3.6: Data analysis.....	58
<b>3.4: Findings .....</b>	<b>59</b>
3.4.1: Policy: shaping the landscape .....	60
3.4.2: Academia: progressing the science .....	64
3.4.3: Industry: creating a sector.....	70
<b>3.5: Discussion .....</b>	<b>71</b>
3.5.1: Societal impacts.....	72
3.5.2: Connecting policy and academia.....	73
3.5.3: Prestige.....	74
3.5.4: Tensions.....	74
<b>3.6: Conclusions .....</b>	<b>75</b>
3.6.1: Security concerns .....	75
3.6.2: Academic entrepreneurialism .....	76
3.6.3: A quantum ‘bubble’? .....	76
3.6.4: Implications for further case studies.....	77
3.6.5: Implications for the Framework.....	78

Chapter 4

- 4.1: Introduction .....79
  - 4.1.1: Institutionalising RI ..... 80
  - 4.1.2: The ‘learning organization’ and responsible innovation ..... 81
  - 4.1.3: Research questions ..... 84
- 4.2: Research method .....84
  - 4.2.3: Analytical processes ..... 85
- 4.3: Findings.....85
  - 4.3.1: Understanding RI ..... 86
  - 4.3.2: Institutionalisation processes, and embedding responsible innovation..... 89
  - 4.3.3: The importance of assessment ..... 92
  - 4.3.4: Student approaches to responsible innovation ..... 94
  - 4.3.5: Broader familiarity – RI practice in industry..... 95
- 4.4: Discussion .....97
  - 4.4.1: RI in industry and commerce..... 98
  - 4.4.2: Institutionalisation of RI at funder level..... 99
  - 4.4.3: Importance of individuals..... 99
- 4.5: Conclusions .....100
  - 4.5.1: Implications for the Framework ..... 101
  - 4.5.2: Implications for further studies..... 102

Chapter 5

- 5.1: Introduction .....103
- 5.2: Background .....104
  - 5.2.1: Research questions ..... 105
- 5.3: Methodology .....105
  - 5.3.1: Identification of participants ..... 106
  - 5.3.2: Analysis..... 107
- 5.4: Findings.....107
  - 5.4.1: National interests ..... 108
  - 5.4.2: Building a quantum sector ..... 109
  - 5.4.3: Process in policymaking ..... 111

5.4.4: Process gaps .....	116
<b>5.5: Discussion .....</b>	<b>126</b>
5.5.1: ‘Breaking it all into bits’ vs real life .....	126
5.5.2: Organic vs systematic proces .....	128
5.5.3: Individuals and networks .....	129
<b>5.6: Conclusions .....</b>	<b>129</b>
5.6.1: Responsible innovation in policy.....	131
5.6.2: Implications for the Framework.....	131
5.6.3: Implications for the industry study .....	132

## Chapter 6

<b>6.1: Introduction .....</b>	<b>133</b>
<b>6.2: Background .....</b>	<b>134</b>
6.2.1: Availability of investment funding .....	136
6.2.2: The problem of inclusivity in novel technologies.....	136
6.2.3: Responsible innovation in industry and commerce.....	137
6.2.4: Research questions .....	139
<b>6.3: Methodology.....</b>	<b>140</b>
6.3.1: Identification and recruitment of participants.....	140
6.3.2: Analysis.....	142
<b>6.4: Findings.....</b>	<b>142</b>
6.4.1: Growth of investment and impact on development.....	143
6.4.2: Capacity challenges .....	147
6.4.3: Spinning out: academia to industry .....	150
6.4.4: Governance in commercial quantum computing .....	152
6.4.5: National concerns and protectionist approaches .....	157
<b>6.5: Discussion .....</b>	<b>159</b>
6.5.1: Governance gaps.....	159
6.5.2: Promises and realities .....	161
6.5.3: Innovation v risk-aversion .....	162
<b>6.6: Conclusions .....</b>	<b>162</b>
6.6.1: Implications for the Framework.....	163

## Chapter 7

<b>7.1: Introduction</b> .....	<b>165</b>
7.1.1: Factors from the landscape study:.....	165
7.1.2: Factors from the institutionalisation study:.....	166
7.1.3: Factors from the policy study:.....	166
7.1.4: Factors from the industry study:.....	166
7.1.5: Additional factors from the literature review.....	167
7.1.6: Underlying principles.....	167
<b>7.2: What is a framework?</b> .....	<b>168</b>
7.2.1: Process models.....	169
7.2.2: Policy frameworks.....	169
7.2.3: Frameworks in industry.....	170
7.2.4: Academic frameworks.....	171
7.2.5: The Double Diamond.....	171
<b>7.3: Creation of the Beehive Framework</b> .....	<b>172</b>
7.3.1: The Beehive Framework.....	174
<b>7.4: Using the Beehive Framework</b> .....	<b>176</b>
7.4.1: Diverging – anticipation and engagement.....	176
7.4.2: Reflecting – surfacing concerns.....	178
7.4.3: Converging – response.....	179
7.4.4: Recording and reporting.....	181
7.4.5: Iterating – a circular process.....	182
7.4.5: Framing the process.....	183
<b>7.5: Validating the Beehive Framework</b> .....	<b>184</b>
7.5.1: Delphi study.....	184
7.5.2: Feedback.....	185
7.5.3: Field-testing.....	185
<b>7.6: Conclusions</b> .....	<b>187</b>

## Chapter 8

<b>8.1: Introduction</b> .....	<b>189</b>
<b>8.2: Themes from the work</b> .....	<b>193</b>

8.2.1: Granularity: individuals, teams, and institutions .....	193
8.2.2: ‘Messiness’ and procedure .....	194
8.2.3: Governance and adaptability .....	195
8.2.4: Diversity vs homogeneity .....	196
8.2.5: Institutions and institutionalisation .....	196
8.2.6: Linearity and repetition.....	197
<b>8.3: Responsible innovation in practice.....</b>	<b>197</b>
<b>8.4: Contribution of the work.....</b>	<b>198</b>
8.4.1: Contributions to responsible innovation .....	198
8.4.2: Contributions to computer science.....	199
8.4.3: Contributions to quantum computing .....	199
8.4.4: Contributions to policy.....	200
8.4.5: Contributions to funders.....	200
<b>8.5: Recommendations for future work .....</b>	<b>201</b>
<b>8.6: Final thoughts .....</b>	<b>201</b>
<b>References .....</b>	<b>203</b>
<b>Appendix A: Papers and other publications.....</b>	<b>219</b>
<b>Appendix B: Sample coding framework.....</b>	<b>221</b>
<b>Appendix C: Sample transcript .....</b>	<b>223</b>

## Abstract

This thesis investigates the value and effectiveness of anticipatory governance techniques in a new technology – in this case, quantum computing. It contributes a novel framework to encourage more widespread institutionalisation, translational capability, and ease of use for such anticipatory governance methods.

Novel technologies frequently have (and indeed are expected to have) societal impact – quantum computing is one such technology. However, technologies that can have societal impact need to both merit and retain societal trust – such trust is predicated upon reliable and responsible methods of development, which enable publics to be assured of both specific and systemic trustworthiness. These concepts are closely allied to governance – but in the case of fast-developing technologies, there is often not yet any ‘traditional’ form of governance such as regulation. Traditional forms of governance are generally reactive, waiting for harm to occur before being framed. This is contrasted with ‘anticipatory’ governance, which seeks to mitigate potential negative outcomes by involving stakeholders in developmental processes, empowering researchers and developers to reflect on potential outcomes and consequences, and enabling them to respond in ways that can reduce harmful societal impacts.

This thesis has investigated these topics through qualitative case studies and analysis in three key domains for quantum computing – industry, academia, and policy. The work has engaged with leaders and significant figures across these sectors, drawing out themes that are specific to quantum computing as well as those that are commonly encountered across computer science disciplines, and some that are found in all novel technologies. As a result of these investigations, it has developed a Framework that can be used for anticipatory governance purposes across research, industry, and in policymaking. The thesis closes by describing how the Framework has been demonstrated and evaluated in numerous contexts; addressing and elaborating on some of the persistent underlying themes that have emerged through the work; and suggesting avenues for further research.



---

# CHAPTER 1: INTRODUCTION

---

*“Quantum computing could impact on society [in ways that] are so profound it is hard to ignore”*

*National Quantum Technologies Programme (NQTP) Statement of Strategic Intent*

The scientific and social context for this thesis is the accelerating development of quantum computing, which – although perhaps not yet considered to be an established branch of computer science – is already having societal impacts as it develops and becomes commercialised. This impact is evidenced through the levels of public and private investment around the world; the growth in mainstream news stories around quantum computing; and increasing public awareness of developments in the field. The impacts of novel technologies are not always uniformly positive, however, and as pervasive and persistent technologies become ever-more embedded in daily life, negative impacts also become more obvious. Methodologies directed towards ameliorating such negative impacts (ideally before they occur) have therefore also become steadily more mainstream. Developing such pre-emptive – or anticipatory – methodologies for governance around one novel technology (in this case quantum computing) can provide a model for use in other areas of computer science and contribute to greater understanding of good practice in more general technological development and deployment.

Quantum computing was chosen as an indicative technology for this purpose not just because its development and impacts have been little studied from a governance perspective, but because the unfolding of such a technology in real-time allows for the investigation of anticipation, consultation, and responsiveness in action – this is anticipatory governance writ large. As this thesis will show, it is possible to map these activities taking place at the highest strategic levels of policymaking and across government, all the way down to individual decision-making processes within research projects. The thesis also discusses the importance of utilising these different levels of granularity, and the enhanced perspectives afforded by this understanding of a range of units of analysis. It will show that these provide an essential ‘systems thinking’ view in terms of the need for a holistic approach to novel technologies.

## 1.1: A growing challenge

Quantum computing has developed rapidly in the decades since Richard Feynman first spoke of it as a hypothetical (Feynman, 1982), and this pace of change has accelerated in recent years, driven by a

number of closely linked factors including surging investment from both private and public sectors, international competitiveness, and technological breakthroughs. A quantum computer uses ‘quantum bits’, known as qubits, to perform its operations. Qubits are exponentially more powerful than classical bits because of certain key characteristics including, for example, the phenomenon of superposition, which allows for a spectrum of calculations to be carried out simultaneously. This and other characteristics of qubits are likely to significantly reduce the time taken to perform calculations but should also enable different *types* of calculation. It is suggested that a quantum computer will offer significant advantages over a classical computer in solving particular challenges that are extremely difficult and time-consuming with classical methods, such as large-number factoring, true randomisation, and optimisation of multiple variables (Lu, 2019). Possible use-cases may include quantum chemistry (Reiher et al., 2017), machine learning (Biamonte et al., 2017), and drug discovery (Cao et al., 2018), which all bring with them the prospect of significant societal impacts, as well as other affordances that have not yet been theorised or discovered. The advantages of discoveries such as these appear attractive enough that the world’s largest technology companies, as well as many research-focused organisations, are investing significant sums into building one form or another of quantum computer (Gibney, 2019). Gibney (2019) also points out that the field is increasingly populated by start-ups in quantum computing-adjacent technological development (for example crystal manufacture), indicating that the development of quantum computing is moving beyond research phases into commercialisation. Globally, the quantum computing market was estimated to be worth £458m in 2021, with predicted growth to £3,932m by 2029<sup>1</sup>. The scale of venture-capital investment, government commitment, and development potential that this growth represents indicates that quantum computing may impact society in multiple ways across a broad front, becoming an enabling technology in ways that are unlikely to be predictable.

### 1.1.1: Multiple impacts

Novel technologies can affect society in both positive and negative ways, as has been seen with the example of machine-learning on classical computers. Through the enormous growth in compute power and the increased availability of vast datasets seen in recent decades, machine-learning has become able to provide significant advantages in fields that handle large quantities of data and have a focus on pattern-recognition. This facility has enabled major advances in, for example, skin-cancer care, where the trained machine-learning model has proven to be adept at helping to identify cancerous lesions (Hekler et al., 2019). However, this same machine-learning technology, when used

---

<sup>1</sup> <https://www.fortunebusinessinsights.com/quantum-computing-market-104855>

for facial recognition applications, for example, has been shown to both entrench and amplify existing biases against people of colour (Garvie, 2019). The potential for societal harm represented by some machine-learning applications and data misuse is demonstrated by Figure 1, which shows a small selection of negative AI-related news from the first part of 2018.

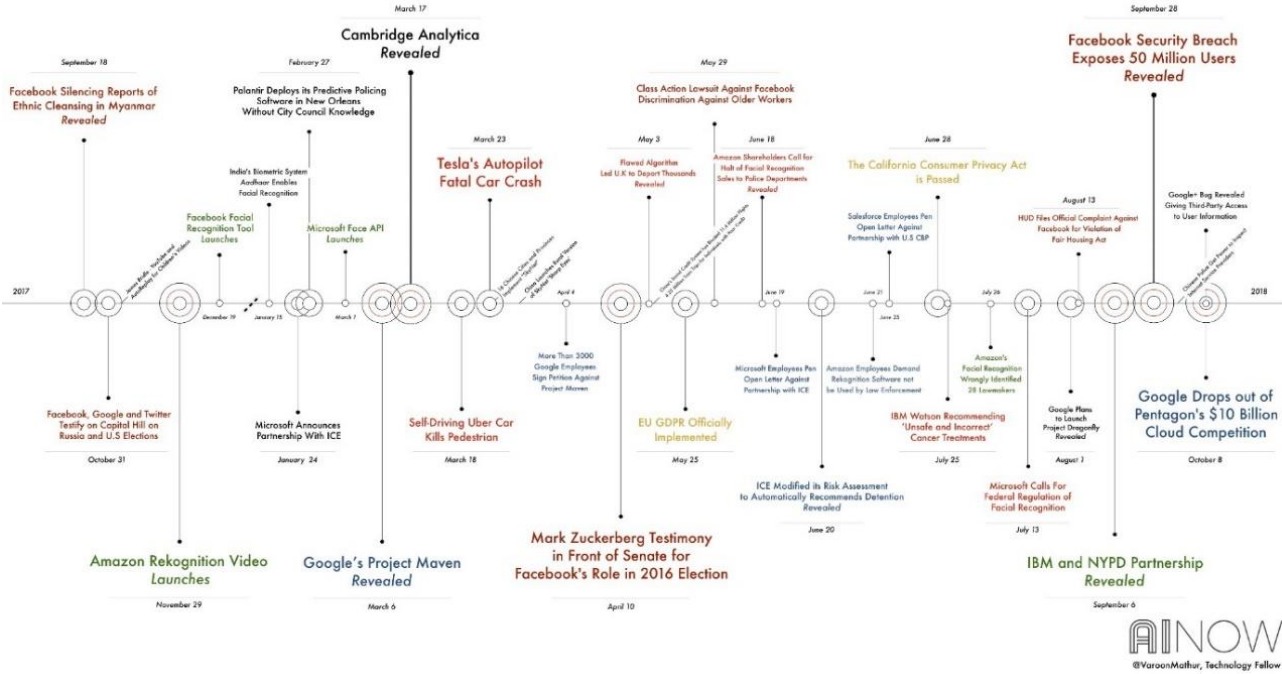


Figure 1: Image taken from a presentation by Meredith Whittaker, president of the AI Now Institute

Negative impacts such as these can and do cause societal pushback and loss of trust – in the case of facial recognition there are vocal campaigns to outlaw its use and IBM, among others, has announced a halt to its facial recognition work (Krishna, 2020), while Microsoft is insisting that users of its facial recognition technology prove they can comply with Microsoft’s ‘responsible AI’ standards (Hern, 2022). It is also possible that the uses of machine-learning that have positive effects may become ‘tainted’ by those that are perceived to have negative effects, giving rise to additional loss of trust. Machine-learning provides an example of both the negative and positive sides of novel technologies, and demonstrates the importance of trying to anticipate and ameliorate the negatives. It also shows the need to commit to revisiting and responding to these challenges as unintended effects become apparent.

As with the example of machine-learning, quantum computing may also have significant – and likely unintended – effects on society. As mentioned above, predictions of the affordances of a quantum computer are varied. For example, if they do indeed allow for high-speed optimisation of large

numbers of variables, this could make them extremely valuable for new drug and materials discovery and other fields that rely on such calculations. They may also in effect create new fields, such as quantum chemistry (Reiher et al., 2017) and quantum machine learning (Biamonte et al., 2017). However, it is widely agreed that one of the capacities of reliable quantum computers is likely to be their ability to break existing cybersecurity protocols through rapid large-number factoring (Johnson, 2019b). The importance of secure communications in an era of embedded and growing Internet dependence is obvious (in addition to the geopolitical implications), and although a cryptographically useful quantum computer that might present society with this particular challenge is some way off, it serves as an example of needing to anticipate and respond to potential negative impacts. In order for the benefits of a novel technology to be realised, these negative outcomes must be minimised to try and ensure that society can and will trust the technology. It is therefore argued that *the ways in which new technologies are driven, shaped, and introduced are critical both for societal acceptance and to ensure that society can receive the benefits while limiting possible negative effects*. These are all aspects of development that could or should incorporate governance, which is a key element in societal acceptability.

## 1.2: The limitations of current practice

As discussed above and in more detail in Chapter 2, the need for governance is a strong theme in quantum, as in many novel technologies – the House of Commons Science & Technology Committee found that publics were likely to be accepting of quantum technologies so long as they were assured of good governance (House of Commons Science and Technology Committee, 2019). Governance types are briefly discussed below and Chapters 2, 5, and 6 examine in more detail how governance-types can form a continuum of granularity but can also broadly be divided into two forms – retrospective and prospective. The retrospective types include regulation and legislation, which (for good and valid reasons) are not usually formulated until harms have already occurred. Prospective forms of governance include responsible innovation and other approaches (such as Real-Time Technology Assessment and Participatory Design) – these generally seek to improve societal outcomes and prevent harms from occurring.

### 1.2.1: Retrospective governance – the difficulties of regulation

The challenge of regulating technological sectors, which are frequently characterised by fast-moving innovation, and driven by speed-to-market, is well expressed by Owen et al., (2013):

*“[regulation] is poorly equipped to govern areas of novel science and technology which are highly uncertain in terms of their current and future impacts, or which,*

*by virtue of their novelty, have no historical precedent. Regulation, put simply, struggles with innovations that it has not encountered before.” (Owen et al., 2013 p32)*

There is currently no UK legislation that specifically applies to quantum technologies, nor any from other countries that this research has discovered to date (with the exception of statutes enabling funding such as the US’s National Quantum Initiative Act 2018). This is not least because quantum computing is at such an early stage that it would not be a productive use of legislative time to enact laws or regulations that may never be needed or may be inapplicable. However, there are not even (at the time of writing) ongoing consultations for a regulatory framework for quantum. For purposes of comparison, there are such consultations in progress for other technologies that are also at an early stage, such as fusion.

It is possible or even likely that quantum computing, at some point in the future, will require formal regulation, including legislation, at both national and international level. However, the process of creating legislation is slow, may be difficult to accomplish, and is generally reactive rather than pre-emptive. Legislation is also framed at high levels that do not seek to address the detail of the numerous potential outcomes. Other governance tools that provide greater degrees of granularity, such as industry and professional standards, require a mature sector with established players and institutional infrastructure. If possible, it is better to address challenges before they arise than to wait for harm to occur, which may undermine societal trust (as in the machine-learning example above).

In the case of quantum computing technologies, therefore, more forward-looking approaches would seem to be indicated as a necessary first step towards the requisite governance.

### 1.2.2: Prospective governance – current approaches

There is some work being carried out to shape elements of a governance ecosystem for quantum computing – for example the IEEE is adopting preliminary informal standards around performance benchmarks and terminology in quantum computing (IEEE, 2019). There are also quantum industry bodies such as the Quantum Economic Development Consortium (QED-C) in the US that seek to proactively create codes of conduct. However, these approaches by and large focus on the creation of high-level agreements that can form a well-regulated market in the quantum computing industrial sector. The World Economic Forum (WEF), by contrast, has outlined a set of general governance principles for the responsible design and deployment of quantum computing (WEF, 2022).

Approaches such as IEEE’s and the WEF’s are helpful, but may be regarded as insufficient in terms of their ability to support public trust, focused as they are on technical standards (such as the meaning of quantum ‘advantage’), actions of the marketplace, or high-level policy. These are certainly necessary, but omit several key factors that, it is argued, also need to be considered. For example, standards-

based approaches do not take into account the societal impacts of these novel technologies, nor do they address questions of innovators' responsibility, or the day-to-day decisions that shape a technology's development.

It was concerns such as these – of societal impact and the responsibility of innovators to consider these questions – that led to the incorporation of responsible innovation specialists in the first phase of the UK's quantum computing programme, the Networked Quantum Information Technologies Hub (the quantum computing Hub was one of four quantum Hubs – the others focused on quantum sensing, imaging and communications (EPSRC, 2020). The requirement for responsible innovation work was considered to be key, and was therefore included in 2019 with the creation of the Quantum Computing and Simulation (QCS) Hub, although no specialist expertise was incorporated. The UK is also not the only jurisdiction to consider prospective governance to be crucial – there are calls in the US for a 'soft law' approach that retains agility while also protecting end-users (Johnson, 2019a) from potential negative consequences.

### 1.2.3: Pre-empting negative consequences

As outlined above, one such pre-emptive approach to governance is Responsible Innovation (RI)<sup>2</sup>. Responsible innovation is particularly apt for a technology in the process of moving from the research to the commercial sector. Although definitions vary, one of the most useful comes from the European RRI-Tools project:

*“RRI is a way to do research that takes a long-term perspective on the type of world in which we want to live... [RRI means] involving society in science and innovation 'very upstream' in the processes of R&I to align its outcomes with the values of society.”*  
([www.rri-tools.eu](http://www.rri-tools.eu))

Responsible innovation methodologies draw together multiple threads and will be discussed in more detail in Chapters 2 and 4. They incorporate consultation with those who may be affected; consideration of both negative and positive outcomes; and responding to both these consultations and reflections – potentially influencing the course of a technology if deemed necessary (Stilgoe, Owen, et al., 2013). Responsible innovation arose partially as a response to some of the scandals involving emerging sciences of the last few decades, such as the BSE crisis, and the complex arguments around genetically modified organisms (GMOs). It has since become embedded within European and UK policy frameworks (de Saille, 2015). It aims to try and ensure that research and innovation are carried out *for* and *with* society in order both to increase acceptance and to improve outcomes for society

---

<sup>2</sup> This is sometimes framed as Responsible *Research and Innovation*, but given the examination of the commercial sector as well as the academic, this work uses the broader term Responsible Innovation

(Owen et al., 2012), and its principles have become adopted across a wide variety of fields – there is work on responsible innovation in nanotechnology (Pandza & Ellwood, 2013), in synthetic biology (Ribeiro & Shapira, 2019) and in climate engineering (Stilgoe, Owen, et al., 2013) – all complex, large-scale novel technologies. This breadth of use makes responsible innovation a useful basis for a method that seeks to pre-empt negative effects in quantum computing. Responsible innovation is not the only such ‘pre-emptive’ approach, as will be discussed, but its embedding in policy frameworks in the UK makes it a valuable exemplar in terms of these anticipatory approaches to novel technologies.

Prospective governance approaches such as responsible innovation that *do* incorporate considerations of societal impact are not without their own challenges, as Chapters 2 and 4 will show.

### 1.3: Academia, industry and policy

The importance of good governance in a novel technology such as quantum computing is not limited to one area of practice. The embedding of new technologies in society requires joined-up approaches that cross boundaries between academic research, commercial marketplaces, and the structuring of new policy – the absence of such holistic approaches may lead to ill-considered or inequitable outcomes for communities or individuals. Without frameworks of governance, societal needs may be only partially considered, if at all.

In particular, such a lack of wider societal consideration may be seen within commercial quantum contexts, where there is no applicable regulation, no industry body with a code of professional standards, and no widespread adoption of pre-emptive techniques. This is especially pertinent because quantum technologies straddle the worlds of academia and industry, with many researchers who are closely engaged in foundational quantum science also involved in start-up businesses in the quantum sector (Gibney, 2019). The literature-review and the fieldwork carried out for this thesis demonstrate that although there is this overlap of personnel and expertise from academia into industry, the very different requirements of the commercial context mean that frameworks of responsibility that may have been learned and applied in academic life are often not brought into the new context. Likewise, the commercial world may have very different requirements in terms of professionalism, legal compliance, and financial drivers.

Any framework for prospective governance must therefore encompass and bridge these different environments in a coherent and structured way. As Chapters 2 and 6 show, although industry has various methods and motivations for achieving ‘responsible’ outcomes, these are not by any means universal, nor are they mandated, and may have serious gaps, as well as potentially failing to recognise the stake that wider society has in the outcomes of their work.

The policy context is also critical to consider as part of any governance framework, because so much work in quantum computing (and quantum technologies in the wider sense) is publicly funded or otherwise supported by government – in the UK over £1bn has either already been invested or is earmarked for investment by government in quantum fields (Reynolds, 2020). The way in which policy is shaped in relation to quantum computing is therefore a highly significant factor when examining the motivations, drivers, potential brakes and possible accelerators for the actions of industrial and academic innovators in quantum computing.

It can thus be seen that any research on prospective governance in quantum needs to address this set of relationships among academia, industry, and policy (known as the triple-helix model) and draw the three strands together in order to create a detailed map of the overall development landscape that can inform the creation of a utilisable framework. This combination of factors has therefore acted as a shaping motivation for the programme of research described in this thesis. The fieldwork has examined ways in which the developers, policymakers, and purveyors of such novel technologies can consciously ameliorate negative effects and improve societal ‘fit’ through the use of prospective (forward-looking) governance.

## 1.4: Research questions and objectives

### 1.4.1: Research question

The aim of this research has been to develop a prospective governance method utilising pre-emptive approaches such as responsible innovation to ensure that the development of quantum computing takes place in accordance with societal needs and is well-governed as it develops. The overall aim is to both secure and retain societal trust and improve the likelihood that these technologies will not be the subject of societal ‘pushback’ such as that seen in the discussion of machine-learning above and also in the older examples such as genetically modified organisms. The possibility of such rejection and mistrust from society has indeed already been considered in relation to quantum computing – the House of Commons Science & Technology Committee found that “support for the development of quantum technologies grew as people’s understanding increased” but also added an important rider “*provided that research was subject to proportionate governance mechanisms.*” (House of Commons Science and Technology Committee, 2019, emphasis mine). It is these crucial questions around governance mechanisms that this research addresses – in particular the adoption of pre-emptive governance methods which seek to create pathways of development that are aligned with societal interests. These concerns therefore give rise to the following overarching research question:

**RQ1:** *How can we use prospective governance methods to ensure that novel technologies, such as quantum computing, are designed and implemented in a way that can fulfil societal needs while minimising negative impacts?*

In addition to this overarching question, the research addresses supplementary questions, including:

RQ 1.1: How do policy development and research mutually influence one another in quantum computing, and how does this manifest in the innovation lifecycle?

RQ 1.2: What is the effect on the innovation landscape in quantum computing of dual-role personae (eg academic researchers who also own or run quantum computing startups)

RQ 1.3: Will one form of prospective governance be appropriate for all phases of the quantum development lifecycle, and are there optimum points in the lifecycle at which to start work on this type of governance?

RQ1.4: Are there unique characteristics of quantum computing that differentiate it from other novel technologies?

## 1.4.2: Aims and objectives

The interconnected and interdependent academic, policy, social and industrial contexts outlined briefly above and detailed more extensively in Chapter 2 could possibly be characterised as a ‘wicked’ problem<sup>3</sup>. The resolution of all these would be beyond the scope of this research. However, this thesis will argue that a focused and theoretically robust use of prospective governance processes could draw together and create more coherence between policy, industry and academic approaches to the development of quantum computing. The aim is to support this through the creation of a structured framework that could help to limit potential negative impacts on society, by providing an easy-to-use methodology and indicators that can both track concerns and provide metrics for success.

The lacunae in literature and practice that give rise to the above research questions are examined in more detail in Chapter 2, and some of the issues arising have been outlined above. In order to address these questions, this research therefore has the following objectives:

1. To understand current ‘responsible’ practice in academia and industry, including levels of stakeholder inclusion.

---

<sup>3</sup> Using Rittel and Webber's 1973 formulation of ‘wicked’ problems: “a problem that is difficult or impossible to solve because of incomplete, contradictory, and changing requirements that are often difficult to recognize ... an idea or problem that cannot be fixed, where there is no single solution to the problem; and "wicked" denotes resistance to resolution”

2. To engage with the academic-industrial context in order to understand the varying drivers and motivations that prevail in the commercial environment and seek to gain greater understanding of the timing of prospective governance approaches.
3. To understand current policy approaches at varying levels of granularity and the impacts of these on various industry sectors for purposes of comparison with the quantum sector, and for deepening the understanding of how policy and regulatory frameworks shape and are shaped by academic work.
4. To create a practical framework based on participatory approaches that can act as a bridge between academic and industrial contexts.
5. To evaluate the utility of the framework from the perspective of academics, policymakers and commercial users through peer-review, through presenting it in academic projects, and through trialling it with a training company.

## 1.5: Contribution of the work

This work will contribute to computer science, to quantum computing, to funding bodies, to policy, and to responsible innovation.

1. **Theoretical contribution to computer science.** Computer science fields are increasingly mandated to incorporate responsible innovation within teaching and research, but there is not widespread understanding of either the substance of responsible innovation or how to apply it in practice. The model offered by this thesis will provide **conceptual additions** to the core AREA Framework that will enable deeper understanding of responsible innovation in ICT fields and a process-based outline for increased usability.
2. **Methodological contribution to computer science.** The work will provide computer science disciplines with an evaluated methodology for engaging stakeholder input at upstream phases of research in order to improve outcomes for both research and applied work.
3. **Theoretical contributions to Responsible Innovation.** The challenges of applying responsible innovation in practice are widely discussed and problematised. This thesis argues for the use of a '**process model**' approach and will contribute such a process model to the field, providing a valuable conceptual addition.
4. **Methodological contribution to quantum computing.** Researchers in quantum computing fields frequently struggle with the connection of their work to end-users or to society. The thesis will contribute a translational bridge between academic quantum disciplines and industrial quantum contexts that will also have wider utility.

5. **Methodological contribution to Responsible Innovation.** Responsible innovation, as will be shown, is notoriously flexible and adaptable – this can make it difficult to apply in practice. The proposed framework will offer methodological innovation in terms of the usability of such anticipatory stakeholder-inclusion approaches.
6. **Empirical contribution to quantum computing.** The work detailed in Chapter 3 has already resulted in the publication of two papers, listed in Appendix A, demonstrating the value of this work to the field.
7. **Empirical contribution to policy.** This thesis identifies lacunae in the practical understanding available to policymakers, including funding bodies. It will provide **substantive knowledge** to policymakers at many levels in the form of a deeper understanding of the oversight possibilities of anticipatory governance and ways in which these can be operationalised.
8. **Empirical contribution to funders.** It will afford funders such as the Engineering and Physical Sciences Research Council (EPSRC) insight into the application of responsible approaches throughout their funding remit. The work described in Chapter 4 has been published and provides a contribution to the evidence-base around the institutionalisation of responsible innovation approaches.

## 1.6: Methodology

### 1.6.1: Qualitative methods – rationale

The research methods used in the research described in the thesis have all been qualitative, for several reasons. Firstly, quantum computing is a very specialised field, with a relatively small community of researchers and innovators (Coenen & Grunwald, 2017) and the research questions were best addressed through drawing on expert knowledge and viewpoints. Secondly, the use of techniques such as semi-structured interviews allows for the investigation of key questions but also permits the surfacing of other possible points for discussion, which can provide fruitful avenues for further research or illumination. Thirdly, the responsible innovation-based approach taken by the fieldwork required engagement with stakeholders in a consultative way in order to draw on different types of expertise and engage with a variety of views.

### 1.6.2: Jurisdictional limits

The companies, research institutions and any policy analysis incorporated into the studies were limited to the UK environment, although comparisons with other jurisdictions have been drawn at various points. This is largely because, as discussed above, the academic field is strongly shaped by policy considerations and the work also crosses seamlessly into industry (frequently with the same

personnel). It is therefore the case that although the work may be generalisable to other industry sectors and/or disciplinary areas within the UK, the very different regulatory context in other jurisdictions would necessitate fresh study.

The case studies have therefore all been based within the UK and sought to address the Objectives (Section 1.4.2) that can help to answer the Research Questions (Section 1.4.1).

## 1.7: Outline of chapters

### Chapter 2: Literature review

This review examines the state of the art with respect to the challenges identified in Chapter 1, namely:

1. Novel technologies such as quantum computing offer many potential benefits to society – however it is important that there is a good ‘fit’ with societal need. Societal trust is also required.
2. For societal trust to be earned and justified, there is a need for technologies that are not only novel but also undergoing rapid change and expansion to be adequately governed.
3. The speed of development for many innovative technologies requires agility of response, leading to the suggestion that anticipatory governance techniques might be useful, as being well suited for rapid adaptation to shifting circumstances.
4. These anticipatory governance techniques form part of a broader continuum of governance from legislation through to personal responsibility.
5. Responsible innovation is reviewed and critiqued as one possible approach to anticipatory governance methods for novel technologies. This includes the challenge of embedding new approaches such as responsible innovation into scientific method and into industry, the absence of process models that may make responsible innovation difficult to grasp and apply in new areas, and conceptual challenges.
6. Responsible innovation is reviewed in relation specifically to quantum computing technologies.
7. Implications for the development of the Framework are discussed.

### Chapter 3: Case Study 1 – the quantum computing landscape

Chapter 3 asks the question *Are there particular features exhibited by the quantum computing landscape that differentiate it from other novel technologies?*

In seeking to map the landscape of quantum computing and interrogate the question of whether it has unique qualities, the study engages with stakeholders across the area, including funders, policymakers, researchers and academics, industry specialists, regulators and entrepreneurs. Outcomes from this study inform the other case studies, and there are also implications for the Framework.

#### Chapter 4: Case Study 2 – RI in practice

Chapter 4 asks the question *How are different fields and disciplines addressing the challenge of incorporating responsible innovation into their work, and what can this tell us about how new norms such as responsible innovation become institutionalised?*

This Chapter picks up the themes identified in the literature review in relation to the institutionalisation of new norms, such as responsible innovation. It investigates this through an analysis of the challenges and opportunities of the rollout of EPSRC-mandated responsible innovation training within UK Centres for Doctoral Training. It also looks at other challenges, such as incorporating anticipatory governance techniques within fields that may feel very distant from societal impacts. Outcomes are discussed, and implications for the Framework.

#### Chapter 5: Case Study 3 – politics and policy

Chapter 5 interrogates the policy landscape, examining its breadth and granularity in order to provide context for the questions around governance. The overriding question for this Chapter is *How is policy developed for novel technologies such as quantum computing?* This incorporates sub-questions such as *'How can policymakers retain the ability to respond to changing circumstances?'* – such as the unforeseen rapidity of developments in quantum computing – and *'How do policymakers address challenges around consultation and governance?'* – especially in technologies where governance is such a priority for citizens.

The Chapter discusses in depth the difficulties faced by policymakers in long-term strategic planning for novel technologies like quantum computing. For example, quantum computing has been a relatively distant goal for most of its development, but in the last few years has accelerated beyond all predictions. Adapting to such quickly-changing circumstances is extremely demanding for bureaucratic processes that are not designed for rapid changes of gear.

#### Chapter 6: Case study – industry

The Chapter 6 case study addresses the third and final strand of the triple-helix model of policy, academia, and industry to ask *'What are the major shaping factors in creating a new sector, and where/how is governance manifested in this process?'* As discussed in the literature review and drawn

out in the preceding case studies, the quantum computing landscape is an example of these three strands being extremely tightly woven, constantly influencing each other. This case study therefore interrogates the final piece of the puzzle, interviewing participants in the commercial quantum world in the UK and the US. It clarifies the business-based perspective on the enablers and constrainers of the quantum computing landscape, and how approaches to ‘responsible’ quantum computing might be viewed and potentially operationalised.

### Chapter 7: the Beehive Framework

This Chapter draws together the work of the case studies to formulate a practical framework for the application of an anticipatory governance methodology to quantum computing. It details the overarching framework and provides indicators for success, but more importantly it also offers a means by which tensions, trade-offs, questions of political power, and value judgements can be incorporated to tailor the framework for specific circumstances. It is this process of tailoring a broad set of outline actions to the needs of an individual project or activity that renders the framework valuable.

The Framework also addresses some of the challenges identified in Chapter 2 for responsible innovation. It provides a ‘process model’ for operationalising responsible innovation that can assist new practitioners in grasping essentials but also provide a scaffold for new projects. Informed by all four case studies as well as the literature review, the Framework has been evaluated across all three strands of the triple-helix model, and in ongoing research projects. The outcomes of these evaluative studies are described here.

### Chapter 8: Conclusions

This chapter interweaves the preceding work to recap the research journey, review the research questions and how they collectively form a picture that has allowed for the building of the Beehive Framework, and discuss ways forward, recommendations, and applicability of both the Framework and the wider research findings for other novel technologies.

---

# CHAPTER 2: LITERATURE REVIEW

---

*“For people that are on the brink of building a world changing technology, you should think seriously about what you’re doing.”*

*Study participant*

This literature review sets out the landscape for responsible innovation as a form of prospective governance that may be useful in attempting to ensure that quantum computing technologies are aligned with societal needs and requirements.

- Section 2.1 examines the background for and emergence of responsible innovation, grounded in the need for society to be able to trust and accept novel technologies.
- Section 2.2 explores responsible innovation in more detail, focusing on the particular framework that has been adopted in the UK and also considering useful insights from other methods and approaches.
- Section 2.3 identifies some of the challenges that have been identified or encountered with responsible innovation in practice, as well as with some of the concepts and assumptions underlying it.
- Section 2.4 discusses how new values and norms such as responsible innovation may become embedded into institutions.
- Section 2.5 assesses responsible innovation as a potentially useful tool in the development of responsible quantum computing and reviews the literature on responsible innovation in quantum computing
- Section 2.6 concludes with a discussion of the requirement and scope for a broad prospective-governance framework for quantum computing technologies.

## 2.1: Science, society and responsible innovation

Commentators including Wynne, (1993, 2006, 2014), Grove-White, Macnaghten, & Wynne, (2000), Stilgoe, Lock, & Wilsdon, (2014) and others have argued for many years that it is vital for the scientific establishment to recognise and engage with the societal context within which it operates. This reflects

what de Saille, (2015) calls the “long shift” away from concepts of science as a ‘neutral’ area – necessarily free from political, social or ethical questions. These commentators and other authors such as Jasanoff (2006) contend that a dearth of such engagement (and the willingness to respond to societal concerns), can lead to mistrust and lack of confidence on the part of society. On the science side of the relationship, a lack of engagement may also lead to an aloof scientific establishment that does not have a good understanding of society’s needs and concerns. The combination of these gaps in understanding and mutual learning may result in societal rejection of potentially useful technologies and a deepening gap between ‘science’ on one side and ‘society’ on the other. Indeed, there was a period when authors argued that the whole of science was undergoing a “crisis of legitimacy”; Farthing et al., (2000), Funtowicz & Ravetz, (1993) and Forsberg et al., (2015) spoke of needing to reform a connection with society that was seen as having been lost. The situation has perhaps improved somewhat since then – more recent research in the UK on public trust in science finds that publics are often well-disposed towards science, and predisposed to trust the scientific establishment (BEIS, 2019). In order for scientific innovations to be accepted it is seen as vital to merit and retain that trust.

### 2.1.1: The importance of societal trust

Part of that societal trust will depend on innovations being well-aligned with societal requirements in both a *product* and a *process* sense – this may mean not just creating innovations that society needs or wants, but carrying out innovation out in a way that enables publics to have confidence in the methods used (Brouneus et al., 2019). Publics need to be able to have confidence and trust in both process and product because technology *affects* society and individuals – and sometimes these effects are not positive. Negative effects can rapidly undermine public trust.

Once trust has been lost in a particular technology, it may be enormously difficult to regain. For example In the 1950s and 60s, public attitudes to nuclear power were almost entirely positive, but repeated scandals, accidents such as Sellafield, and errors, have led to deep mistrust (Dalquist, 2004). Much of society is now extremely sceptical of the value of nuclear power, despite the possibilities it may present for reducing reliance on fossil fuels (Abdulla et al., 2019; Corner et al., 2011). More recently, the rollout of genetically modified organisms is also frequently cited (Burke, 2004) as an example of where public failure of trust resulted in widespread public rejection of the technology.

The importance of this trust-relationship is now widely recognised at policy level, and formed a cornerstone of the EU’s ‘science with and for society’ (SWAFS) theme, expressed partly through the Horizon 2020 programme (European Commission, 2020). In the UK, this line of reasoning is also highlighted in the House of Lords *Science and Society* report (House of Lords Science & Technology Select Committee, 2000), which is regarded as a landmark (per Genus & Stirling, 2018) in the attempt

to improve science-societal relationships. This relationship of trust and co-dependence is seen as particularly crucial in the development of large-scale technologies that could have wide-ranging and significant effects.

### 2.1.2: Realising the benefits of novel technologies

Quantum computing technologies are one such a large-scale novel technology, one in which many governments and private corporations continue to invest very significant sums of money – Germany alone recently announced it was investing EUR2 billion in quantum (Burke, 2020). It is widely accepted that a quantum computer will be significantly more powerful than a classical one, with some suggesting that a quantum computer will be faster than the largest current supercomputers (Johnson, 2019b) – however it is also likely that its affordances will be *different* than those offered by a binary system. A quantum computer may have the facility to, for example, manipulate huge datasets, optimise large numbers of variables, and perform extremely large mathematical calculations (Lu, 2019) – these capabilities have implications for complex system modelling, drug discovery, and market analytics respectively. However, as we saw in Chapter 1 its use-cases may not be wholly positive – it is also clear that one of the earliest possibilities for a quantum computer may be in breaking the current cryptographic formulae on which the security of the world’s communications systems, banking, and finances depend (Bernstein & Lange, 2017; Sharma et al., 2020).

If the potential benefits of these new and distinctive computing capabilities are to be realised and gain public acceptance then it is extremely important to ensure that wider publics are certain of reliable processes and outcomes that can anticipate and minimise harms. The Quantum Technologies report from the House of Commons recognised that the public will require assurances of good governance in quantum technologies if they are to be accepted (House of Commons Science and Technology Committee, 2019).

### 2.1.3: Varieties of governance

Governance can take many forms, can be highly granular and piecemeal, and can take place in many different locales – any system of governance will have many interrelated elements and may, indeed, appear deeply unsystematic (Ozmen Garibay et al., 2023). For most citizens, the most familiar form of governance is legislation. However, technology is notoriously difficult to legislate for, and the example of the Online Safety Bill (proposed in 2017 as the Online Harms White Paper but repeatedly delayed since then) demonstrates that backwards-looking, or retrospective governance, instigated after harms have already occurred, is extraordinarily difficult to frame. As the Paper has progressed towards becoming a Bill, significant alterations and additions have been made to reflect changes that have occurred in technology, in government, and in society since it was drafted. It has taken up many

sessions of Parliamentary time (which is always a scarce resource), numerous interest groups have been involved in the process of scrutiny and amendment of the Bill, and further changes are likely to be required. Many elements are considered to be unworkable in practice<sup>4</sup>, and if the Bill passes, it may well be very difficult to enforce (Aynsley, 2020). In addition, the challenges and harms that the Bill seeks to address are both growing and shifting while it undergoes its legislative passage.

One alternative to this complicated, slow-moving, and potentially ineffective approach is to seek to prevent harms before they occur, using forward-looking forms of governance that aim to improve processes and thereby outcomes. As suggested above, governance can be viewed from a broader perspective as a system with many dimensions. ‘Macro’ forms of governance, which would cover primary and secondary legislation, are both retrospective and relatively distant from the objects of governance. More granular, or ‘micro’ forms of governance, can be prospective and are closer, more immediate, and more individual. Kuhlmann, Stegmaier, & Konrad, (2019) term these forms ‘tentative’ governance, because they are agile, non-prescriptive, and generally flexible. The aim is to use governance that is “dynamic ... prudent and preliminary” – they also note the ‘balancing act’ with more rigid forms of governance.

### ***Prospective governance***

There are several candidates for these forward-looking, non-regulatory types of governance, including standards and codes of conduct – for example some standards bodies such as IEEE are already adopting agreed definitions to ensure international benchmarking for terms such as ‘quantum advantage’ (IEEE, 2019). Standards such as these are frequently framed at a strategic level that largely seeks to ensure a level playing field within the marketplace. This review has not found any codes that are shaped to address societal issues that may arise, nor do these codes usually provide any framework or encouragement to engage with society on a wider basis in order to ensure a strong connection between science and society (as discussed above).

Some forms of prospective governance *do* address these issues of societal concerns and seek to reinforce the trust-relationship discussed above. One of these is Responsible Innovation (RI), a methodology that focusses on the processes shaping the innovation pathway: in addition to the developmental work being carried out on an innovation it draws in anticipation, reflection, engagement and responsiveness to try and ensure that innovations are aligned with societal needs. Concerns about the need for societal acceptance and community engagement led to responsible innovation being directly incorporated into the first funded phase of the UK’s Quantum Hubs

---

<sup>4</sup> <https://www.independent.co.uk/news/uk/bill-the-bill-government-mps-parliament-b2143535.html>

(Inglesant et al., 2016). In the US there are also calls for non-regulatory governance of quantum computing technologies and responsible innovation approaches (Johnson, 2019a). At the global level, the World Economic Forum released its principles for the governance of quantum computing early in 2022 (WEF, 2022), and calls for ‘responsible’ quantum computing. Responsible innovation is not the only form of anticipatory governance – Technology Assessment (in various forms), Participatory Design, and other methods, are also posited as potential approaches to these challenges. In terms of the potential distance between technological development and societal impact, however, responsible innovation takes one of the longest views, as well as explicitly drawing on approaches also found in Applied Ethics and Science and Technology Studies (Grunwald, 2014). This makes it a candidate worthy of investigation in terms of its practical application to the challenge of creating technology that responds to the needs of society.

## 2.2: Responsible innovation and the AREA Framework

The history of the emergence of RI is now well-documented (eg Boenink & Kudina, 2020; Owen, Pansera, et al., 2021). These initiatives emerged over a decade ago, seeking to identify and address uncertainties and risks associated with novel areas of research. They began in nanotechnology (Murphy, 2010), spread to geo-engineering (Stilgoe, 2015) and synthetic biology (Frow & Calvert, 2013), before expanding to include computer science, robotics, informatics, and ICT more generally (Jirotko et al., 2017).

There are various definitions for RI, but in general it is seen as a flexible framework that nevertheless has certain key features. A frequently cited definition is from von Schomberg:

*“Responsible Research and Innovation is a transparent, interactive process by which societal actors and innovators become mutually responsive to each other 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).” (von Schomberg, 2013, p63)*

This shows a clear awareness of the importance of the trust relationship between the scientific establishment and the society within which it operates. It also emphasises the need for ‘transparency’ – an important element in ‘accountability’ aspects of responsible approaches. Within the UK, another framing comes from Stilgoe et al., (2013) who focus more on the relationship between actions taken in the present and their relationship to the future.

*“Responsible innovation means taking care of the future through collective stewardship of science and innovation in the present.” (Stilgoe et al., 2013, p1570)*

These two definitions are useful in understanding RI, both in terms of the centring of the human, or societal element, but also as a time-embedded activity – something that carries decisions made in the

past through into their outcomes in the future. This dual focus, on society and the future, is also seen in the definition from the EU's RRI Tools project:

*“RRI is a way to do research that takes a long-term perspective on the type of world in which we want to live” ([www.rri-tools.eu](http://www.rri-tools.eu))*

Many authors have interrogated the essential meaning of RI – Mertens, (2018), for example, focuses on three assumptions in responsible frameworks: i) that emerging technology needs assessment because of radical novelty and unpredictability, ii) that early intervention is necessary to impact the innovation trajectory and ii) that anticipation of unknowns is necessary to prepare for what cannot be predicted. Other authors focus on the exact meaning of ‘responsibility’ – Pellizzoni (2004) described four aspects, which he divided into two retrospective (liability and accountability) and two prospective forms (responsiveness and care), but he was far from the first to analyse responsibility in this way. Arendt, (1963) and Hart, (1968) were just two of the prominent mid-20<sup>th</sup> century philosophers who argued for a need to shift the tense of ‘responsibility’ from past to future. Responsibility in this sense is *“to exercise foresight and to increase one’s knowledge about the world and how one’s action might interact with and alter it”* (Grinbaum & Groves, 2013, p122) – this is the future-facing form of governance.

This is one of the fundamentals of RI – that it seeks to foreground the relationship with possible futures. It is about the aspects of responsibility that relate to events yet to occur. This sets it apart from some other informal and formal approaches to technological governance (for example Technology Assessment, or regulation), because these are responding to the past, whereas RI engages with the future.

There are many other discussions of definitions and foci, which are often useful for considering the nuances and flexibility of the method, but within the UK, the form of RI adopted by the Engineering and Physical Sciences Research Council is the AREA Framework. This framework for RI is now embedded in all its research calls and RI training is mandated within its Centres for Doctoral Training (CDTs). The Framework consists of Anticipation, Reflection, Engagement and Action/responsiveness (Owen et al., 2013). These will be briefly discussed in turn, before assessing some of the difficulties of operationalising them in Section 2.3.

### 2.2.1: Anticipation

One of the major underpinnings for the concept of prospective governance is that expressed by Collingridge in *The Social Control of Technology*. He discussed a specific type of problem where the

conditions are mutually dependent and/or conflicting. At an early stage of technological development, when changes can be easily made, both the need for changes and the exact changes required cannot be foreseen – by the time they can be foreseen, it is difficult and expensive to change.

*“attempting to control a technology is difficult...because during its early stages, when it can be controlled, not enough can be known about its harmful social consequences to warrant controlling its development; but by the time these consequences are apparent, control has become costly and slow” (Collingridge, 1980, p19).*

It is this mutual interdependence – the tension between the present and the future that will spring from it – that generates the need for the ‘anticipation’ in RI. The requirement for anticipation, therefore, is to prepare, as far as possible, for the unknowns of the future – not least because, as Grinbaum & Groves, (2013) point out – innovation is a “future-creating activity”. These authors take this a step further and argue that in many respects, *anticipation* is already a part of innovation because innovators are considering the future and what can be made that either fits or changes it, therefore the anticipation aspect of RRI could be viewed as simply another perspective on the innovation process.

The act of anticipation is also by definition embedded in the concept of ‘anticipatory governance’. One definition for this is;

*“a broad-based capacity extended through society that can act on a variety of inputs to manage emerging knowledge-based technologies while such management is still possible” (Guston, 2008, vi, cited in Guston, 2014)*

The relationship of anticipatory governance to Collingridge’s dilemma above can clearly be seen. They share not just the requirement to consider the future and society, but also reserve the right to ‘act’ or seek to ‘control’ innovation “*while [it] is still possible*”. This extends anticipatory governance beyond simply preparing for the future and incorporates elements of the ‘responsiveness’ element of RI. The inclusion of ‘a variety of inputs’ also suggests that the stakeholder engagement element is included here. ‘Anticipatory governance’ therefore has significant overlaps with a large part of the RI methodology, not just anticipation.

The most stringent version of anticipatory governance is known as the ‘precautionary principle’, which itself is divisible into a *strong* and a *weak* form. The ‘strong’ form can, in effect, be used to block progress unless and until absence of harm can be proved (essentially an impossibly high bar), whereas the ‘weak’ form is regarded as a more pragmatic approach that permits beneficial innovation even if some harm results (Hemphill, 2020). The ‘weak’ form may have some commonality with RI but may be even more difficult to concretise and will likewise share some of its challenges (see 2.3 below).

### 2.2.2: Reflection

Reflection is one of the lesser-discussed parts of the AREA Framework in the RI literature. Lynch, (2000) describes reflection from a social sciences perspective as *“a central and yet confusing topic”* (Lynch, 2000, p26) listing numerous types and orders of reflexivity. Others including Grimpe, Hartswood, & Jirotko, (2014) adopt a more everyday interpretation that regards activities such as ‘mindfulness’, and ‘taking stock’ as sufficiently valuable in terms of an individual approach to providing a thoughtful background to possibilities and choices.

However, Grimpe et al., (2020) also call for what they term ‘institutional’ reflexivity. They use this term to refer to the day-to-day, collective activities of a given institution, for example a funder, in its considerations of decisions, its responses to feedback, and its relationships with its stakeholders. This develops the term from its use by Wynne, (1993) and others where it signified the general responsiveness of the scientific establishment to societal concerns. Institutional reflexivity is analysed further by Grimpe et al., (2020), who also trace a supplementary concept of ‘collaborative’ reflexivity that emerges from the numerous interpersonal discussions and decisions that occur within the everyday business of the institution. In this understanding, ‘reflection’ for RI purposes takes place in various forms dependent upon the level of granularity, but remains a very practical and routine activity.

Reflection and reflective practice also occupies a significant place in the literature on professions, which is relevant because of the ongoing debate about the professionalisation of ICT-related career roles (for example UK government attempts to create an information security ‘profession’ - Reece & Stahl, 2015). Many professions require reflective work as part of initial qualification or continuous professional certification (Moon, 2004) – in these situations the purpose of reflective work is to create actionable directives for ongoing professional development. These approaches are also seen in Participatory Design which draws on the processes of ‘reflection-in-action’ as discussed in the work of Schön, (1983) to understand how to use a reflective mindset as a part of professional practice. Reflective work such as this is meta-cognitive and designed to encourage deep thinking about motivations, assumptions and directions. Sen, (2010) shows that reflective practice provides *“evidence of learning, self-development, the ability to review issues critically, awareness of their own mental functions, ability to make decisions, and being empowered”* (Sen, 2010, p84). These processes of active, action-oriented reflection are explicitly linked in the ‘professions’ literature to the essential meaning of professionalism (eg Lester, 2015). Within Participatory Design discourse, Ehn, (2002) in particular interprets this as *“a synthesis of constructive action and critical reflection”* (Ehn, 2002, p21). These perspectives from other fields thus have more in common with Grimpe et al (2014)’s understanding of reflexivity as a practical, day-to-day exercise rather than a philosophical construct.

These different levels of reflection and institutional reflexivity all have a role to play in the practice of RI – institutional reflexivity, for example, may also overlap significantly with the requirement to engage with stakeholders (Pereira & Saltelli, 2017).

### 2.2.3: Engagement

Engagement with stakeholders is the third element of the AREA Framework, and it is here that ‘participatory’ approaches such as those seen in Participatory Design are most relevant. The ‘participatory turn’ of the 70s across many disciplines is well-described by Schubotz, (2019) who discusses the increasing number of projects in recent decades that describe their methodology as ‘participatory’. The trend is for increased democratisation of research (Hansen et al., 2019; Nowotny, 2003), with some research methods such as Action Research and Participatory Design also having a specifically ‘political’ understanding that aim to give voice to the under-represented, and frequently an additional ‘transformative’ aspect of empowering social change (Beck, 2002).

It also seems clear that these drives for more inclusive research relate not only to the repeated calls from commentators for a better relationship of trust between science and society as discussed earlier (eg Liberatore & Funtowicz, 2003; Wynne, 1993, 2006, 2014), but also to Collingridge’s dilemma outlined earlier – involving stakeholders and society is seen as one way of trying to intervene at a sufficiently early point to be able to influence an innovation’s trajectory in accordance with societal needs and thereby avoid harms. There is therefore a strong conceptual link between engagement and governance, made explicit in Collingridge’s ‘social control of technology’ theme.

Participatory initiatives have also made their way into high-level policy programmes, although the degrees of actual participation differ greatly between projects. For example the ‘participatory’ element in Participatory Design is envisaged as fully co-created work (Steen, 2013), but at a policy level this may take the form of more consultative exercises, such as the ‘Sciencewise’ programme that carries out extensive public engagement on scientific topics<sup>5</sup>.

It is therefore not surprising to find participation and engagement taking such a prominent role in prospective governance frameworks, as these elements become the guiding principles that prompt a response from developers and technologists.

---

<sup>5</sup> <https://sciencewise.org.uk/about-dialogue/what-is-public-dialogue>

## 2.2.4: Responsiveness

The final element of the AREA framework is usually framed as ‘act’, though understanding it as ‘responding’ allows for a more nuanced approach, particularly as the relevant response may in fact be the *absence* of action. Responsiveness can be seen as an attempt to institutionalise agility.

Pellizzoni, (2004) unpacks the concept further, and discusses responsiveness as ‘responding’ both in the sense of answering, and in the sense of reacting. Nielsen, (2016) agrees that both are important, but suggests a third aspect, which is the relationship between responsiveness and governance. He points to accounts such as those from the European Environment Agency, (2001) which drew attention to the frequency with which warnings were issued on environmental matters but responsiveness was lacking, despite accountability frameworks. These failures of governance were seen as failure to ‘respond’ appropriately to warnings. Nielsen also discusses governance processes of creating socially-robust knowledge (per Nowotny, 2003) by ‘responding’ to citizens through increased democratisation of expertise (Liberatore & Funtowicz, 2003). (This also reflects Collingridge’s concern for the ‘social control’ of technology.)

Responsiveness is also likely to be context-specific – a commercial organisation may be alert to different sets of stimuli than a research organisation. Management literature over the years has moved towards a modality of adaptiveness (Levinthal, 2018). For example, drivers in the market, such as consumer desire for more socially-aware companies, may result in the adoption of frameworks such as Corporate Social Responsibility – defined by the European Commission as ‘the responsibility of enterprises for their impacts on society’ (European Commission, 2011). This type of ‘response’ to markets would usually have no legal consequence if breached, but may very well have an economic consequence if consumers are disappointed with corporate behaviour (Malecki, 2018), making it potentially useful as another type of ‘soft’ governance per Nielsen (2016) above.

As with other elements of the AREA framework, normative questions come to the fore because in any ‘response’ to either internal or external developments there are also questions of what *ought* to happen. Davies & Selin, (2012) describe these questions as “competing normativities” and ask how it is possible to decide whose version of ‘what ought to happen’ takes priority. The political aspects of these normativities therefore also need to be engaged with, as there will be power relations involved in decisions about responses. As well as the political questions, some of the other themes seen in stakeholder engagement also resurface here, such as the willingness to accept challenges to existing assumptions (Forsberg et al., 2015).

## 2.3: RI in practice – difficulties and possible approaches

The application of RI in practice presents some theoretical and methodological challenges.

The flexibility of the RI concept is demonstrated by the breadth of discussion in the literature around definitions of terms, such as the exact meaning of ‘responsibility’, and by the number of indicators of possible ‘good’ RI practice (eg Strand et al., 2015). The discussions above also demonstrate the degree to which the processes of RI – anticipate, reflect, engage and act – overlap and interact with each other. In terms of its description as a ‘framework’, it may perhaps be best understood as a set of ‘frames’ in the social science sense, rather than a model, or scaffold, in the software engineering sense. Ribeiro et al., (2017) describe it as an ‘umbrella’ term, suggesting that RI can be adapted to a variety of different situations without discarding any key elements. *“Its greatest potential may be in its ability to unify and provide political momentum to a wide range of long-articulated ethical and policy issues”* (Ribeiro et al., 2017, p81). This pragmatic approach seeks to leapfrog precision of meaning in pursuit of practical workability.

However, many commentators are less positive about the flexibility of RI. Some criticise the degree to which it has become common in the innovation discourse without being adequately defined – Delgado & Åm, (2018) argue that the indeterminacy around its definition is problematic not just in terms of defining RRI but around conceptualisations of ‘public good’. Da Silva et al., (2019) claim that RRI has “stalled” after its adoption as a process of normative governance, and is without sufficient guidance as to how it is to be done in practice.

Others point to the difficulty of acquiring the additional skillsets that may be required. Conley, Foley, Gorman, Denham, & Coleman, (2017) argue that for RRI to be successful, innovators will need to be experts in not just their field of innovation, but also in how to innovate responsibly – this is known as ‘T-shaped’ expertise (Guest, 1991). There are criticisms of the use of RI as an ‘add-on’, or the potential for it to become a tick-box exercise (eg interviewees for Pansera & Owen, 2020 and Ten Holter et al., 2022).

Even proponents of RI acknowledge its challenges in practice. Per Jirotko et al., (2017), these challenges can include

- a. research and innovation not following predictable paths
- b. the bundling-together of research and innovation hides their differences
- c. pluralistic democracies do not usually have social consensus
- d. stakeholder engagement can be misused
- e. RI is predicated upon particular values, around which there may be disagreement

The pragmatic approach described above does not materially assist with these difficulties of implementing RI in practice. The remainder of this section outlines some of the knottier challenges.

### 2.3.1: Challenges for anticipation

Carrying out anticipatory work in practice presents challenges on many fronts. As von Schomberg points out, *“negative consequences are often neither foreseeable nor intentional”* (von Schomberg, 2013, p59). Epistemologically the future can’t be known, so some commentators ask whether we can subject it to governance processes that run the risk of completely missing the mark (Mittelstadt et al., 2015). Mittelstadt et al., (2015) are among the many commentators that also point to normative challenges – by attempting to steer a technology in one direction or another, choices are being made about what ‘ought’ to be the case.

Some have argued that because it isn’t possible to know with certainty what outcomes any action will have, ‘responsible’ and ‘irresponsible’ innovation are indistinguishable from each other during the innovation process and that this therefore cannot be action-guiding (van de Poel & Sand, 2018). Von Schomberg would disagree, listing behaviours such as *“technology push, neglect of fundamental ethical principles, policy pull, and lack of precautionary measures and technology foresight”* as specifically ‘irresponsible’ innovation (von Schomberg, 2013, p60).

Nordmann, (2014) suggests that even *“perfect knowledge”* of what will happen is insufficient to allow us to react to the future. Likewise, at the other end of the scale from such ‘perfect’ knowledge are entirely unknowable ‘black swans’ – events so unpredictable that it is unlikely that they could have been foreseen in any of their specifics. Nordmann also questions the degree to which we owe a duty of care to the future – what obligations can be ascribed towards our future selves or others’?

There are some responses to these challenges for ‘anticipation’, however.

#### ***Imagining the future***

Envisioning exercises and focus groups are both methods for trying to draw on collective ideas of what lies in store by looking at what might be emerging from academic work, or industry developments, and extrapolating (eg Macnaghten, 2020). The use of envisioning techniques is just one way to ‘prepare’ for the future (Reeves, 2012) and to mitigate uncertainty. Van Oost, Kuhlmann, Ordóñez-Matamoros, & Stegmaier, (2016) discuss the value of scenario analysis as a form of anticipatory governance – these methods are based on understanding that these are not predictive tools, but rather ways of exploring possibilities. Related to this is work on the value of moral imagination as an envisioning tool, which also links to the idea of collective knowledge and public engagement. This

approach recognises that people will make morality-based judgments about what they think *ought* to happen (Lehoux et al., 2020).

### ***Preparedness and foresight***

Nordmann, (2014) argues further from his ‘perfect knowledge’ point that we don’t need to *know* the future, as long as we can be *ready* for it – this requires preparedness and agility, keeping options open and not closing down alternative paths (this links clearly to Collingridge’s dilemma). This is an area where the commercial world has significant experience – the use of ‘foresight’ in industry has been commonplace for many years. Gordon, Ramic, Rohrbeck, & Spaniol, (2020) provide a useful historical analysis of the business literature, showing how business forecasting first used foresight strategies to improve their planning. A trend towards greater and greater contextual detail in order to try and fine-tune these type of predictions was eventually abandoned (Bjorklund, 1988), as businesses realised that “*unforeseeable events were genuinely unforeseeable*” (Gordon et al., 2020, p4). Instead, the literature charts the move to ‘preparedness’, with Rohrbeck & Schwarz, (2013) specifically referring to ‘anticipatory’ approaches. The business management literature has therefore created a plethora of tools with which to ‘anticipate’ – these include relatively well-known approaches such as PESTLE analysis (as generally adapted from Aguilar, 1967) for analysing the macro-business context, or Five Forces diagrams to examine a company’s relationship with their competition (Porter, 1979). It is very possible that some of these tools could be compatible with RI-type approaches.

### ***Ethical anticipation***

Mittelstadt et al., (2015) seek to address the epistemological and normative challenges mentioned above. These can be summarised as the questions of ‘how can we know the future’, and ‘how can we decide what the future ought to be’. Mittelstadt et al suggest that discourse ethics and speculative ethics (per Habermas, 2008, 2011) provide a possible means for grappling with these questions. These ethical frames seek to unpack presuppositions to get at the roots of normative questions, drawing analogies with religious belief to demonstrate the possibility of taking action based on something fundamentally unknowable. They try to connect what *is now*, to what *ought to be* in the future through the use of dialogue and discussion that aims to ensure common points of departure and translation between different contexts (Mittelstadt et al., 2015). The links between this position and the importance of stakeholder engagement within responsible innovation are clear.

### ***Analogies and narratives***

Analogies are a powerful tool that can enable thinking about the future in a manageable way. They are particularly useful for allowing extrapolatory thought, whereby an understanding of a known phenomenon can be extended to a new phenomenon with which it shares characteristics (Schwarz-

Plaschg, 2018). Identifying characteristics that allow for greater comprehension and manageability of new concepts is a well-understood phenomenon, and for example leads to terms such as ‘architecture’ being used to describe software. The potential power of such analogies can be demonstrated by reference to the contention from the earlier part of the 21<sup>st</sup> century that ‘data is the new oil’ – this characterisation of data as something precious and finite may well have been one factor contributing to companies attempting to harvest as much of it as possible. These questions of the metaphorical characteristics of a new concept – for example whether data is indeed ‘the new oil’ – can have significant effects on the ways in which people comprehend and react to discussions around its use or misuse (Coyle, 2018). Framing, including challenges of analogy, and meanings become another challenge for RI in the stakeholder engagement phase, as will be seen.

### 2.3.2: Challenges for engagement

Although frequently discussed as a relatively simple exercise, engagement can take many forms and be difficult to carry out in practice. These difficulties can include questions around existing power-structures (for example the ability to decide *whether or not* publics should be consulted on a given question), process questions (such as *who to include*, and *who is able to decide* who is included), and the management of outcomes (which may be very different than were expected). Some commentators are clear that power-relations should be centred in RI work, because the decisions about when, how and whom to include, as well as to what end, need to be understood in terms of the power structures in play (eg Nielsen, 2016).

#### ***When to consult***

The timing of inclusive approaches may be crucial. Bearing in mind Collingridge (1980) and the need to retain flexibility early on in the process of technological innovation, much RI work calls for ‘upstream’ engagement (eg Banks, 2016; Gerber, 2018; Krabbenborg & Mulder, 2015) – meaning at an early phase of the development process – with da Silva et al., (2019) agreeing that ‘late’ inclusion leads to poor anticipation, because paths are already locked in. From the Participatory Design literature, Beck, (2002) would agree, but points out that “*intervention requires a location*” (Beck, 2002, p80) – in other words it is important to be able to judge when is the correct point at which to intervene. ‘Early’ and ‘late’ are relative terms and give no guidance about precisely when intervention is ideal. It is possible to argue, however, that developers should err on the side of earlier, rather than later intervention and engagement.

#### ***Managing outcomes***

Davies & Selin, (2012) point out that inclusive deliberative processes may appear simple on paper, but in practice are often ‘messy’. This messiness may result in no clear ‘answer’ from a stakeholder

engagement exercise, potentially allowing it to be sidelined altogether (Smallman, 2018). For example, publics may regard ethical and social concerns as inextricably bound up with scientific questions, but the level of nuance and complication that this brings to concrete decision-making for policymakers may be unmanageable. There is frequently an attempt within policymaking to regard social or ethical factors as ‘epiphenomena’ – ie as entirely separate from the ‘science’ question (Smallman, 2019) – which is an area on which they may fundamentally disagree with citizens. In particular, Smallman, (2018) finds that policymakers may tend to regard science as an unquestioned good:- a source of progress, answers to complex questions, and economic growth. There is an obvious tension here with the societal concern that unconstrained science can be problematic, and that industry is a ‘corrupting’ factor. It may be for reasons such as these that Smallman also finds that instances of public engagement exercises actually influencing policy are so rare as to be almost undocumented. Da Silva et al., (2019) adopt a combative position in the face of this messiness, discussing what they term ‘excessive’ inclusion and suggesting that this may lead to lack of resolution, or no clear decision. However, this may be missing the point of inclusion, which is not necessarily to come to a conclusion or achieve consensus but rather to discuss alternatives and possibilities in a collective way.

Another perspective on the possible outcomes of inclusivity relates to the value of ‘alignment’. Blok, (2019) draws attention to this desire to find consensus among stakeholders and cites an in-depth literature analysis from Gray & Stites, (2013) that found a strong tendency to seek alignment and agreement during consultative processes. Blok calls this approach reductive, claiming that it elides or eliminates difference of opinion, and suggests alternative approaches that can acknowledge failure to ‘agree’, failure to ‘decide’ and failure to ‘solve’ these complex problems. Pereira & Saltelli, (2017) agree that when investigating ‘wicked’ problems, you may get ‘wicked’ answers, and that simplistic, or linear models of science and innovation are not useful. In particular it is necessary to recognise that these are often socially and politically contested issues (Jirotko et al., 2017), so consensus may be impossible to achieve. Pilemalm, (2018) recommends seeking out “stroppy” stakeholders who will ask difficult questions and challenge accepted views or the status quo. The question then becomes how to draw useful conclusions from a diversity of views – a possible approach comes from some of the Participatory Design literature, which argues that the important point is to engage with possibilities and alternatives, rather than necessarily focus on an ‘answer’ (Bødker & Kyng, 2018). This will be challenging in practice, and those attempting to innovate ‘responsibly’ may struggle with how they should, or whether they can, integrate diverse viewpoints. These questions will be interrogated in more detail in the fieldwork.

### ***The politics of inclusion***

Forsberg et al., (2015) argue that a realignment of science as part of and co-created by society creates a need for societal input not just into the *how* of innovation, but also into the *why* and the *for whom* and the *what*. This requires decisions to be made about what and who is included or excluded – and these are political decisions, per van Oudheusden, (2011), who is clear in later work that responsible innovation is a normative-political orientation (van Oudheusden, 2014). As Nowotny, (2003) puts it, *“The question of whose knowledge is to be recognised, translated and incorporated into action has been exacerbated under the pressure for democratisation”* (Nowotny, 2003, p152, emphasis mine). Blue & Dale, (2016) agree, arguing that the question of *who* is consulted may be even more exclusionary in areas where technical understanding may be required. Wesselink & Hoppe, (2011) point out that deciding to consult publics in the first place is already a political decision, while others are clear that the way in which questions are framed also has political overtones (Blue & Dale, 2016; Pereira & Saltelli, 2017).

Nielsen, (2016) in particular is emphatic about why interactive or participatory governance of this type is difficult:

- i) The ad hoc nature of the ‘participatory’ process (which may be dictated by political concerns, as above) generally means there is no defined infrastructure or process for inclusion.
- ii) interactive processes are heavily influenced by what may actually be possible
- iii) there can be a lack of transparency about how decisions are taken.

Genus & Iskandarova, (2018) also point out the tension between the concept of RI as a ‘democratic’ method and the reality that innovators are in a position to ‘invite’ participants, undermining claims to inclusivity.

### **2.3.2: Additional challenges**

As well as specific challenges around anticipatory approaches and stakeholder engagement, the practical implementation of RRI as a methodology may face difficulties in terms of its applicability in industry; assumptions around what may be meant by ‘societal alignment’; challenges of repeatability; difficulty of making comparisons between projects; questions of normativity implicit in these approaches; and issues with self-regulation. Some of these are examined below.

### ***Industrial transition***

The challenge of translating responsibility frameworks into industrial contexts is just one part of the complex relationship between academia and industry – a relationship that has shifted in the last thirty

years as governments have recognised the economic value of the knowledge that may be generated in university research labs. De Saille, (2015) provides a valuable analysis of how this evolved at the level of the European Research Area (ERA), as policymakers sought to leverage research knowledge into economic advantage.

There is work that directly addresses the transition of ‘responsibility’ frameworks into industry (eg Dreyer et al., 2017; Timmermans, Yaghmaei, Stahl, & Brem, 2017). These often seek to map the concerns of RI on to prevailing or existing models of ‘responsible’ behaviour in companies, such as Corporate Social Responsibility (CSR) – for example Blok, (2019) discusses how the mitigation of negative consequences, use of stakeholder engagement, and anticipation of future challenges all have utility in the industrial context. Martinuzzi, Blok, Brem, Stahl, & Schönherr, (2018), also make the business case for RI, suggesting that it may increase competitiveness as well as increasing public trust. The point about its value for competitive advantage is also supported by Schönherr, Martinuzzi, & Jarmai, (2019) who argue that responsibility and innovation can be mutually supportive activities, while Chatfield, Iatridis, Stahl, & Paspallis, (2017) point out that many companies are willing to balance economic and non-economic factors to create societal value. Sustainability concerns have also been increasingly evident in industrial sectors, with some commentators examining the value of RI approaches in addressing environmental concerns (Stahl et al., 2019).

### ***Societal alignment***

The phrase ‘societal alignment’ is frequently used in disciplines that seek to foreground the societal context for innovation – this includes disciplines such as Technology Assessment and Science & Technology Studies as well as RI – and is used to describe technology that can work in accordance with society’s needs. However, the literature rarely discusses the point that ‘society’ is not a homogenous mass – it comprises everything from individuals up to federations of countries and all types of collective in between, with many varying requirements (Grimpe et al., 2014). The more granular meanings of the phrase are often not unpacked – as Ribeiro et al (2018) point out, the deeper issues of ‘societal alignment’ – for example social justice, unequal distribution of both positive and negative effects of technology – often remain unaddressed in the RI discourse. Drilling down to the application of these values in practice can result in mundane-seeming but potentially important moves such as the Public Services (Social Value) Act 2013 that allows public sector bodies to take account of social and environmental factors as well as economic ones in procurement exercises (Uyarra et al., 2019). This type of granularity needs to be recognised and included in discussions of ‘societal alignment’ and the RI discourse should recognise that ‘societal alignment’ is not a simple concept. The importance of retaining an approach that operates at different levels of granularity will be revisited in the empirical chapters.

## **Self-regulation**

RI, because of its forward-looking, harm-prevention modality, is situated at the more granular, self-regulatory systems of governance – meaning that it requires individuals, groups and companies to act pre-emptively. This would appear to place the onus on innovators and designers to be the instigators of this type of anticipatory governance. Asking innovators to act in two different modes – for example as both designer and regulator – may create tensions when their different modes have different requirements (such as the *Challenger* shuttle crash where during the decision-making process an engineer was asked to “put his management hat on” (Feynman, 1988), which resulted in a different – fatal – decision). This problem bears comparison with ‘regulatory capture’, where large industry bodies are consulted on legislation that is meant to govern their activities, allowing them in some cases to shape it according to their own wishes – they are being expected to operate in two different modes that may conflict. It is therefore legitimate to question whether self-regulation can be sufficient within ICT where a variety of motivations and drivers will be experienced by innovators, of which ‘responsible innovation’ may be only one. An individual may have multiple perspectives on a decision, some of which may conflict, and a responsible innovation approach may be only one factor to be considered. One possible response to this challenge comes from commentators who stress ‘collective’ reflexivity and responsibility (Grimpe et al., 2014), potentially distributing the onus of self-regulation throughout the institution – which may have its own challenges.

## 2.4: Institutionalisation? Embedding RI in research

There are many aspects to consider when examining how institutions respond and adapt to new modalities such as anticipatory governance, and in practice one of the largest challenges that RI faces is that of becoming embedded into institutions such that it becomes ‘taken for granted’. The body of literature examining ‘institutions’ is complex, global, and vast, with significant theoretical discourses interrogating every aspect of institutional formation, development, and dissolution. However, it is possible to narrow the field of discussion to key works that examine, and seek to explain, the ways in which institutions *change*. In particular the work of Thelen is extremely valuable in this regard. In the *Theory of Gradual Institutional Change*, (Mahoney & Thelen, 2010), the authors argue persuasively that institutions respond to various pressures, both internal and external, and can be observed to shift substantially over time without such shifts ever being explicitly acknowledged.

Randles & Laasch (2016) build on Thelen’s work in their analysis of institutionalisation processes, describing it as the “*stabilisation of norms*” (Randles & Laasch, 2016, p54). They distinguish between formal adoption of such norms into policy and procedure (top-down), and a *de facto* process where new practices are seen in individual and group behaviour irrespective of ‘official’ procedure (bottom-

up). The ‘top-down’ version of this may be seen in the location of the EPSRC as the key point for the institutionalisation of RI in the UK; the Research Council’s promulgation of the AREA Framework; and its inclusion of RI as a mandatory component of funding proposals (for example in doctoral training centres (EPSRC, 2018). Randles & Laasch (2016) also emphasise that for these ‘new norms’ to be embedded, others will likely need to be discarded, a process they term *(de)institutionalisation*. Institutions are by definition ‘stable’ entities – which means that the inertia of their “*incumbent logics and practices*” can be highly resistant to change (Pansera et al 2020). Genus & Iskandarova, (2018) discuss the way in which RI has been institutionalised in the EU and US as “partial”, and describing the motivations for the inclusion of RI as rooted in “*the need to secure desirable outcomes from scientific research and promising innovations*”. Owen, Pansera et al, (2021) also critique the absence of formal methodology in how RI – mandated in funding calls – is then assessed at the proposal stage and during the project’s life. This question of assessment will be seen to be key during the fieldwork.

However, an interesting study from which comparisons can be drawn for quantum technologies is found in the study of RI within a synthetic-biology research centre (Pansera et al., 2020). Like quantum computing, synthetic biology has been seen as a “*promising innovation*” for economic development (per Genus & Iskandarova 2018, above). Pansera et al., (2020) analysed some of the challenges of carrying out the funder-required RI over the course of several years, arguing that with an highly interpretive framework such as RI, it is necessary to examine its use “*in situated practice*” to gain a real understanding of the challenges and successes involved. Pansera et al., (2020) followed a learning arc of RI understanding from a simplistic model linked to older ‘public education’ concepts (focused on public engagement), towards a much more “*deliberative and dialogic*” approach. However, Pansera et al., (2020) are sceptical of whether this type of process can be described as ‘institutionalisation’, as it may not result in changes to wider practice (Owen, Pansera, et al., 2021). The ways in which RI can become embedded within institutions will be discussed further in later chapters, but the experience of an RI “learning curve”, seen in this synthetic biology project, may have ready applicability to RI in quantum technologies.

## 2.5: RI in quantum computing

‘Responsible’ approaches have already been considered as a fundamental part of the development of quantum computing technologies – the UK’s Networked Quantum Information Technologies (NQIT) Hub that was funded in the first phase of EPSRC’s quantum development programme in 2014 included embedded RI specialists as part of its programme (Inglesant et al., 2016; Ten Holter et al., 2021b). However, despite the recommendations from the House of Commons Science & Technology Committee on the public requirement for governance (House of Commons Science and Technology

Committee, 2019) RI specialists were not incorporated into the second phase of quantum computing research, the Quantum Computing & Simulation Hub (although it has designated funding for RI). The next round of funding from the National Quantum Technologies Programme also calls for responsible innovation in all the Hubs<sup>6</sup>, but it remains to be seen how and whether it will be assessed as a key part of the project proposals.

### 2.5.1: The discourse of RI in quantum computing

To date, there is not a large body of research that addresses what may be regarded as the key questions in this field – *what will be the impacts of quantum computing on societies? who will benefit? who will lose?* but it is useful to review some of the discourse around these questions in order to identify gaps and directions for research.

The concept of ‘quantum ethics’ has appeared relatively recently, in *Scientific American*, on the ArXiv and elsewhere online (Holland, 2018; Khan, 2021; Perrier, 2021). These contributions often build on the deep history of the ethical discourses of computing, covered not least in the comprehensive review by Stahl, Timmermans, & Mittelstadt, (2016) of nearly 600 sources. However, it can be argued that consideration of the ethics of quantum computing does not provide sufficient concrete guidance to the researchers and industrial start-ups grappling with the technical and engineering challenges attendant on the day-to-day business of building a quantum computer. This is not to downplay the value or importance of drawing on ethical approaches – in particular Perrier (2021) situates the ethical and moral challenges around quantum computing firmly within the wider discourse of computing ethics – but even ‘applied’ ethical approaches can be extremely difficult to operationalise. Other approaches may also make valuable interdisciplinary contributions to the discourse, such as Michael Frayn’s play, *Copenhagen*, which questions the validity of ethics from the point of view of living in a quantum universe (Ruddick & Heisenberg, 2001), and the novel *Helgoland*, a “meditation on quantum theory” (Rovelli, 2021). Dihal, (2017), also reflects on the ‘narratives’ of quantum. It is important to draw on a multiplicity of voices and viewpoints in order to progress the dialogue around quantum computing impacts. However, the focus of a responsible innovation approach is based in the concrete actions and day to day decisions of those working to create the technology – it is rooted in ethics but has a very practical application.

One of the most focused attempts to date to create a discourse around these critical questions of societal impact was the special issue of the journal *Ethics and Information Technology* in 2017, edited

---

<sup>6</sup> [https://www.ukri.org/opportunity/epsrc-quantum-technology-research-hubs/?utm\\_medium=email&utm\\_source=govdelivery](https://www.ukri.org/opportunity/epsrc-quantum-technology-research-hubs/?utm_medium=email&utm_source=govdelivery)

by Vermaas. Vermaas pointed to the debate on quantum technologies in society that “*is currently not taking place*” (Vermaas, 2017, p241) and sought to generate such a debate. He stressed his concern that quantum theories are often presented as “enigmatic” and essentially incomprehensible, and regarded this as creating a barrier to public understanding, believing that an understanding of the theory is necessary for society to be able to take a view on the impact. As well as calling for more inclusive public debate, he exhorted philosophers of physics to make quantum theory more understandable. (It is interesting to compare this position with the Public Dialogue exercise from the RRI work in NQIT, which found that it was *not* necessary for publics to understand the theory, but that it was necessary to provide considerable explanation of the theoretical applications before impacts could be discussed in any meaningful way - EPSRC, 2017).

One of the ‘philosophers of physics’ so exhorted by Vermaas is Grinbaum, who also contributed to the special issue (Grinbaum, 2017). Grinbaum discusses the ‘narratives’ of quantum computing and how narrative forms can support understanding. This approach has important overlaps with the ‘anticipation’ dimension of RRI, which (as discussed in section 2.2.1) uses techniques such as moral imagination (Lehoux et al., 2020), analogy (Schwarz-Plaschg, 2018) and envisioning (Reeves, 2012) to support anticipatory governance approaches. All these authors point out that framing such innovations in metaphorical or narrative terms allows non-specialists to grasp implications and extrapolate from them without needing to understand the actual phenomena. To some extent Grinbaum supports the pragmatic approach in terms of helping non-specialists gain understanding of quantum technologies, but he also points to the shortcomings of this method, which he argues will not be sufficient to generate the public trust that is necessary.

De Wolf’s contribution to the special issue takes a complementary approach that tackles societal impacts of quantum in practical terms, rather than the ways in which people think about and understand quantum technologies at a high level (De Wolf, 2017). De Wolf examines the possible effects of particular quantum computing capabilities, such as its potential impact on cryptography and secure communications. Going beyond discussion of post-quantum cryptography, he points to the delicate balance between privacy and security and how that might be disrupted should a few powerful agents have access to quantum computing that is denied to others. The related question of ‘*who has access*’ to this technology is also discussed, and he argues that inequity of access to quantum computing may vastly amplify existing inequalities of wealth and power (this concern is also seen elsewhere, eg Ten Holter, Inglesant, et al., 2022). One possible solution he suggests – of companies providing wide access to quantum computing capability via cloud services – is currently being offered by AWS’s Braket system, among others, though it remains to be seen whether this will be either a long-term or effective counter to the centralisation of quantum computing in the hands of a few

powerful actors. Finally, De Wolf points out that it is not just lack of access to a quantum computer that may increase inequity, but access to the knowledge that has gone into creating it – the quantum research carried out within companies (as opposed to within academia) is indeed regarded as highly valuable intellectual property, and is guarded accordingly. He calls for governments to do more to democratise both the knowledge and the practical access to quantum computing.

Moving further along the ‘practical outcomes’ spectrum, DiVincenzo takes an entirely ‘applied’ view of the impacts on society of quantum, making estimates about what actual technologies society might expect to see, and refraining from judgement about whether these technologies might be beneficial or harmful, or what other externalities might emerge (DiVincenzo, 2017). Möller & Vuik, (2017) also restrict their considerations to the technical impacts of quantum computing, focusing on its impact on high-performance scientific computing. Arguably, these limited investigations of ‘societal impact’ do not significantly contribute to the wider discourse that Vermaas was hoping to stimulate with the special issue.

The remaining contribution in the issue, from Coenen and Grunwald, compares a small group of policy-oriented reports that discuss the use of responsible research and innovation in quantum technologies within a European, a British, and a German setting (Coenen & Grunwald, 2017). Comparing and contrasting the three approaches is a useful exercise and it is valuable for RI technologists – particularly in a rapidly developing field – to learn from and with each other. Coenen & Grunwald discuss similarities and differences between the application of RI in quantum and its earlier application in nanotechnology fields, seeking to draw lessons from RI work in these areas. The authors also touch on the importance of the ‘narratives’ of quantum – such as those discussed by Grinbaum, (2017) and Dihal, (2017) – and the ascription of meaning (a hermeneutic approach) to framings of quantum technologies for publics. However, their main analysis focuses on the wide interpretation of RI, frequently expressed as “*science with and for society*”. This becomes operationalised as questions of public engagement and how it was built in to the approaches in the policy documents they analysed. Coenen & Grunwald’s position is that in projects embedding RI there must be a significant public engagement or even co-creation aspect – which they term ‘strong’ RI. They describe the UK’s approach to RI in quantum as ‘weak’, contrasting it unfavourably with the equivalent approaches in Germany – a country that is also making significant research investments into quantum computing (Burke, 2020). However, the establishment of the UK’s National Quantum Computing Centre (UK NQTP, 2019) may provide a fresh opportunity for the UK to draw prospective governance approaches into its quantum computing work.

Overall, the discussions in the Special Issue demonstrate the breadth of potential ethical and responsible challenges, with many overlapping issues and concerns to take forward to the practical development of quantum computing.

## 2.5.2: Operationalising RI in quantum computing

### Anticipation

Possibly because of the relatively public nature of advances in quantum computing, which are frequently discussed in the business press and other mainstream media as well as in academic journals (for example Murgia & Waters, 2019), there is already some anticipatory work underway. For example, as has been seen, it is suggested that a quantum computer may be able to break current cryptographic methods in far shorter timeframes than a classical computer. This has implications not only for secure communications in the future, but for the future data-security of present-day communications (Sharma et al., 2020). The field of post-quantum cryptography, which is working to ensure online security even after the advent of a cryptographically-significant quantum computer, is already advanced.

As discussed in section 2.5.1, ‘narratives’ of technologies can play an important role in anticipatory work, particularly when it comes to shaping public attitudes. The analysis by Dihal, (2017) of the narratives of quantum – both in fiction and non-fiction – discusses many of the prevailing themes of quantum computing. Dihal shows that public impressions of quantum do not (currently) suffer from the same negative narratives as robots or artificial intelligence. Coenen & Grunwald, (2017) would agree that the public understanding of quantum has not suffered from as many misleading or unhelpful narratives in the same way as some other technologies have (notably nanotech, which became identified with a narrative of ‘transhumanism’<sup>7</sup>).

However, there will certainly be many unforeseen and currently unforeseeable impacts of quantum computing, and anticipatory work that can help prepare for these will be critical in retaining public confidence, per the House of Commons Science and Technology Committee, (2019). Anticipatory work in quantum computing, then, may benefit from current generally positive connotations in terms of public trust and understanding, but it will therefore be desirable for ongoing trust to avoid undermining those views.

---

<sup>7</sup> <https://en.wikipedia.org/wiki/Transhumanism>

## Reflection

As previously seen in Section 2.2.2, reflection and reflexivity are included in the RI in their pragmatic sense, and at various levels of granularity from the individual level up to institutional reflexivity. All these forms are relevant for quantum computing, but given the complex and highly collaborative nature of the discipline, the more collective and institutional forms may be the most relevant. In particular the ‘collaborative’ reflexivity discussed by Grimpe et al., (2020) may be especially significant given the high levels of funding that are required within the quantum computing field and therefore the importance of relationships with policymaking and funding bodies, as well as international collaborators.

The individual level of reflexivity remains relevant, however, in the context of *professionalism* – also discussed in Section 2.2.2. Within ICT and innovation sectors, there is frequently significant interdisciplinarity, and therefore considerable variation in the degrees of ‘professionalisation’ exhibited by different disciplines. The desirability of professionalisation continues to be debated in reference to computer science disciplines, not least because of the flexibility and breadth of fields covered, and despite the professional ‘code of conduct’ set out by the ACM decades ago (ACM, 1992), which explicitly calls for reflection and anticipation. Suddaby & Laasch, (2020) hypothesise that increasing professionalisation is associated with increasing institutionalisation – as a discipline becomes more mainstream and more embedded within societal contexts, the likelier it is that professional structure will be required. The link between societal embedding or institutionalisation (see also section 2.4 above) and the associated need for professional bodies, standards and codes of practice, is one that fields such as quantum computing may need to examine in more detail.

## Engagement

When it comes to highly complex technologies such as quantum computing, Davies & Selin, (2012) offer an especially relevant case study from the NanoFutures project. They discovered in their workshops a gulf between what needed to be understood in order for stakeholders to engage meaningfully with the material, and the level of scientific understanding exhibited by the publics they were engaging with. The challenges they describe include the question of how to frame a problem for discussion when there may be an unmanageable number of variables and a high level of uncertainty. As discussed above, the Public Dialogue exercise on quantum technologies commissioned by EPSRC, (2017) found that although citizens were generally familiar with the term, they did not have a good understanding of quantum phenomena, or the possibilities of quantum computing. However, the exercise noted that once the affordances of quantum devices were explained and participants gained greater understanding of the relevant phenomena, they became increasingly involved and enthusiastic, which may suggest an appetite for further engagement. Guston, (2014) also discusses

how publics are frequently keen to engage with scientific questions and Wynne, (1993) was an early voice pointing out that publics' understandings are frequently context-specific and highly nuanced.

However, as well as the question of how to facilitate meaningful stakeholder engagement in very technical fields, there is also the challenge mentioned in Section 2.3.2 of exactly how early in the innovation process to engage with stakeholders – if the technology is at such an early stage that only disciplinary experts could realistically have a meaningful exchange, then is it useful to involve publics? The timing of any consultation may be key, but very difficult to determine.

A different perspective comes from Stelzer, (2020) who argues that the 'participatory turn' seen in many disciplines (discussed in more detail in Section 2.2.3) is actively harmful to research. Stelzer suggests that complex technological matters should *not* be placed into the hands of wider publics, but rather that their assessment should be limited to experts who understand the science. However it could be argued that the findings of the Quantum Public Dialogue exercise discussed above demonstrate an eagerness on the part of publics to engage with even highly complex science.

### Responsiveness

In some senses, the quantum computing community and related disciplinary fields have already demonstrated the value of an RI approach. As discussed above, one of the earliest possibilities for a quantum computer is suggested to be large-number factorisation that could breach current internet communications security protocols, based as they are on mathematical operations that are extremely time-consuming for a classical computer (Sharma et al., 2020). This anticipatory work has enabled a field of post-quantum cryptography (Bernstein & Lange, 2017), which is already creating post-quantum cryptographic keys to ensure that communications can remain secure (Maeda et al., 2019). There is also an engagement requirement here, in order to ensure that the widest audience possible is aware of the need to become quantum-computing-ready.

### Questions of power

Questions around the relationships of power within, for example, the engagement activities of an RI approach have been evident throughout this discussion of RI and its utility in quantum computing – these questions are also seen in relation to policymaking in the area. Nielsen, (2016) suggests that RI debates have an essential tension at their centre – that of a risk-mitigation approach versus an innovation-driven emphasis in policymaking. This balancing act may be particularly evident in relation to quantum computing because of the extremely high levels of investment involved and the potential that is seen by governments for national leadership in the field, which may be in tension with evidence that the public wishes for robust governance in this area. Additionally, a perspective drawn from Participatory Design, which takes a strongly political stance, may recognise an additional tension

between essentially democratic and inclusive methodologies such as RI being subject to a type of institutional capture as they become institutionalised by funders and other policymakers (Ten Holter, 2022).

## 2.6: Conclusions

This review has discussed the importance of governance for novel technologies, and described why prospective, or anticipatory governance methods may be the most useful form of governance in quantum computing. It is contended that these methods are valuable both in terms of supporting societal acceptance and confidence, and in providing pre-emptive, agile, responsive governance. It has taken a brief look at the basis and current use of some prospective governance methods, focusing on Responsible Innovation (RI) as the broadest and most flexible but also incorporating useful notes from other methods such as Participatory Design. It has also discussed some of the theoretical and methodological challenges for such methods, and some approaches for potentially resolving or addressing some of these challenges, followed by a short examination of the use of prospective governance techniques in quantum computing.

Quantum computing, as discussed in section 1, may offer significant benefits for society, but if society is to receive those benefits then technology must perforce be societally acceptable. Technologies that experience societal pushback, as in the example of genetically modified organisms, are rarely accepted thereafter.

Taking into account the various strengths and challenges of RI, as discussed above, it will be valuable to investigate as the basis for a prospective governance methodology for quantum computing. RI is far from a perfect method, but, with its combination of anticipatory work, engagement, reflection and responsiveness, may provide the most useful framework currently available for supporting public trust in quantum computing and thereby enabling its future value for society.

---

# CHAPTER 3: CASE STUDY - THE QUANTUM COMPUTING ECOSYSTEM

---

*“It's important for all concerned working in this field to have those thought processes around the potential impacts”*

*Study participant*

This chapter takes a broad perspective on the entirety of the quantum computing landscape, evaluating some of the key concerns and themes, and discussing how these will impact further studies and the development of an anticipatory governance framework.

## 3.1: Introduction

As detailed in the previous chapters, this thesis has sought to develop a framework for prospective governance in novel technologies such as quantum computing. Before examining possible techniques or applications of prospective governance it was necessary to begin by mapping the landscape of this novel technology. Accordingly this first field study served as a scoping exercise that looked across the current state of the art in the deployment of quantum computing – its objective was to generate contextual information and specific research foci for the remaining studies, which each then examined one of the key areas of the triple-helix in more depth.

## 3.2: Background

As outlined in section 1.1, over the last decade and backed by significant public and private investment, there has been significant and accelerating development in quantum computing fields, outstripping the most optimistic predictions. Discussed in the previous chapters were some of the possible affordances of quantum computing, such as large-number factoring; optimisation of multiple variables; true randomisation; and simulation of quantum processes (Aaronson, 2014; Feynman, 1982; Lu, 2019), with suggested use-cases including quantum chemistry, machine learning, and drug discovery among others (Reiher et al., 2017, Biamonte et al., 2017, Cao et al., 2018) – all of which present the possibility of significant societal impacts. As has been seen, the possible commercial

potential of quantum computing has generated excitement both among companies and governments, with very large sums of money pushing the field forward, even without further confirmation of the potential use-cases. The rapid development of quantum technologies as a sector is illustrated by the ‘adjunct’ companies that are also springing up – for instance recruitment companies in the UK that specialise in recruiting, training, and placing people in quantum companies.

### 3.2.1: Societal effects

The use-cases potentially available from quantum computing may have significant effects on society – as outlined above, the potential for high-speed optimisation of large numbers of variables, or simulation of quantum processes, is likely to make them valuable for new drug and materials discovery, and many other fields – indeed it is the opinion of some in the sector that pharmaceuticals may be the first profitable use-case (Quantum Tech, 2021). On the negative side of these possible affordances, however, it is widely agreed that a quantum computer of sufficient capacity is likely to have the ability, via Shor’s algorithm (Sharma et al., 2020), to break existing cybersecurity protocols using rapid large-number factoring (Vermeer & Peet, 2020), and this is a topic of intense research at present (Yan et al., 2022). Clearly, the importance of secure communications in an era of embedded and growing Internet dependence is vital, so this serves as a useful example of how in order for the positive aspects of quantum computing to be realised, potential negative effects such as the impact on cybersecurity must be minimised to give society confidence in the technology. While it must be acknowledged that predictions of quantum computing affordances are varied – and in the case of the threat to secure communications much is already being addressed by post-quantum cryptography (Bernstein, 2009) – as with Machine Learning and internet technologies, it is not possible to predict with any certainty what the societal impacts will be some years into the future.

There have been some efforts to engage with publics on quantum computing technologies – the Public Dialogue exercise carried out in 2017 as part of the NQIT-RRRI project (Ten Holter et al., 2021b) represented a significant outreach operation that aimed to assess public understandings of quantum and levels of confidence in the technology (EPSRC, 2017). It found that *“support for the development of quantum technologies grew as people’s understanding increased, provided that research was subject to proportionate governance mechanisms.”* The House of Commons Science & Technology Committee too has addressed this, supporting the call for governance and the need for public confidence in the technology (House of Commons Science and Technology Committee, 2019).

However, at the time of writing, the sector was moving extremely rapidly, with some arguing that Noisy Intermediate-Stage Quantum Computing (NISQ) would be useful (Preskill, 2018), while others

proposed networked or modular quantum chips as a next step forward<sup>8</sup>, and tech giants such as IBM and Google announcing significant strides.<sup>9</sup> The entire field, both research and commercial, has grown enormously through the lifetime of this research, and does not currently show signs of slowing.

This gives rise to the research question for this case study below:

### 3.2.2: Research question

*RQ: Are there particular features exhibited by the quantum computing landscape that differentiate it from other novel technologies?*

#### Objectives

The study detailed below sought to map elements of the quantum computing landscape in academia, in policy, and in industry, in order to address the research question above. Three objectives were identified:

1. What are the relationships between academia, industry and policy, and how do these manifest?
2. What is the state of readiness in terms of public awareness of/engagement with quantum computing technologies?
3. What are the unique features of quantum computing (if any), and what are the commonalities with other emerging technologies?

## 3.3: Methodology

### 3.3.1: Data collection

Two datasets were utilised for this case study. Dataset 1 was collected by the NQIT-RRR research project on responsible innovation in quantum computing from 2015-2016, while Dataset 2 was collected by the researcher during the winter of 2020. Working with both datasets provided the opportunity to take a longitudinal perspective on the opinions and thoughts of scientists, industry personnel, and policymakers within the quantum computing ecosystem. Both datasets consist of a set of semi-structured interviews. Although the interview structure and precise questions differed, the same central topics were used as a starting point for discussion in both datasets. These included (but were not limited to):

---

<sup>8</sup> <https://www.technologyreview.com/2023/01/06/1066317/whats-next-for-quantum-computing>

<sup>9</sup> IBM Quantum roadmap: <https://research.ibm.com/blog/ibm-quantum-roadmap>, Google Quantum Computing Journey: <https://quantumai.google/learn/map>

- Relationships between industry, academia, and policy
- The role of funding (both state and private) in the quantum computing sector
- Quantum computing's status as a nation-state level technology
- Awareness and application of anticipatory governance approaches such as responsible innovation
- Challenges and enabling factors for the technology

### Dataset 1

This dataset comprised a set of qualitative interviews carried out during the life of the Networked Quantum Information Technologies Hub (NQIT) by the responsible innovation team embedded within the Hub<sup>10</sup>. These interviews were undertaken from 2015-2017, with a broad cross-section of researchers and innovators working with and within the Hub. The dataset was made available to the researcher during collaboration with the NQIT-RRI team on a subsequent paper. A breakdown of the interviewees is provided in Table 1.

### Ethical considerations

The consent forms signed by participants at the time of the original study were checked to ensure that the interviews could be reused for subsequent studies.

All findings and analysis discussed below were carried out independently by the researcher.

**Table 1**

Number of interviewees	Role	Code
4	Senior researcher	LD
1	Combined research/industry	LI
9	Researcher	LS

### Dataset 2

This data was collected over a four-month period during the winter of 2020-21, through the process detailed below.

### 3.3.2: Selection and recruitment of participants

After ethical approval had been applied for and granted, potential participants were identified through online research into the relevant quantum computing projects in the UK. The quantum computing community is relatively small, and also relatively open, and thus it was not difficult to identify

---

<sup>10</sup> NQIT-RRI project <https://www.cs.ox.ac.uk/projects/NQITRRI/>

potential participants. 20 individuals working in the quantum computing sector were selected to be approached for interview – the selected individuals worked across academia, in industry start-ups, and in policymaking. Potential interviewees were invited through an email outlining the topic of the research. Once they had indicated their acceptance, they were sent an information sheet and a consent form to sign and return. Of the 20 originally identified, 14 participated in the study. Table 2 shows the distribution of roles among the participants: some participants held more than one position (eg a university position and also a business role) and this is detailed within the table. Where relevant, participants were asked to elaborate on their views from both perspectives.

**Table 2**

Number of interviewees	Role	Code
2	Academia and industry	IA
4	Academia	A
3	Industry	I
4	Policy	P
1	Policy and academia	AP

### 3.3.3: Limitations

1. The research study was relatively small and relatively brief. Further work may be needed across a broader spectrum of quantum sector participants to gain a wider perspective on these findings – this could be done through the use of a survey instrument.
2. Given the aim of the study, it is very possible that those who agreed to participate in the study were already biased in favour of anticipatory governance or responsible innovation approaches – a broader survey as suggested above could attract a wider variety of views. In particular it would be valuable to obtain critical viewpoints in order to understand whether and where there is opposition to responsible innovation approaches.

### 3.3.4: Interview technique

Within a semi-structured interview format, it is necessary to form a rapport quickly with the participant in order to ensure that they feel comfortable discussing the topics at hand and to elicit the fullest possible response. Commencing the interview with questions around their particular area of expertise (rather than focusing immediately upon the interview questions) and then branching off from this to reach the topic of the interview can both support the creation of this rapport and also provide context for the remainder of the interview. This tends to create a portion of the interview that will not be productive from a coding standpoint, but is important nonetheless for ensuring the participant feels confident handling the interview topics. This can additionally be useful in generating

supplementary questions that may be specific to the individual's area of expertise. The average length of a discussion was 35-45 minutes, giving scope to enter into topics in depth, should this be useful.

### 3.3.5: Data management

Interviews were audio-recorded, then transcribed in their entirety and repeatedly checked against the voice-recording in order to obtain and retain as much context and accuracy as possible. Transcriptions were redacted at the point of transfer from voice to text and were given a randomly allocated numerical code. The connection between codes and participants was listed in an offline spreadsheet. Once redacted, the transcripts were then loaded into NVivo qualitative analysis software<sup>11</sup>.

### 3.3.6: Data analysis

Analysis was carried out inductively, without any pre-existing guiding framework (Braun & Clarke, 2006; Charmaz, 2014). This is a method *"consisting of flexible, successive analytic strategies to construct inductive theories from the data"* (Charmaz & Henwood, 2017 p238)

Each interview was open-coded at the sentence-level without looking for themes at this point. Every transcript was analysed twice in this manner, with at least a month between the first and second coding, with this double process allowing for code-checking but also to potentially break codes into more detailed levels (for example on the second run-through, the code 'quantum economy' was split into finer categories such as 'commercialisation' and 'supply-chain'). This particular process of analysis was selected for this study because of the importance of clarifying directions of travel, possible interview questions, and objectives for the next stages of the research.

During the second coding sequence, annotation and memos were also used to capture reflections on the data, note cross-references, and provide aides-memoire for followup. Memos in particular were valuable for capturing insights in an organic way while handling the data, without waiting for the next stage of code-grouping. This process is useful for retaining richness of analysis.

After the interviews had been reviewed twice in this way, the coding framework contained 32 codes. At this point it was useful to group codes into common themes – no codes were combined, in order to preserve the nuance of the data – and these themes are detailed in the Findings section below before being discussed further in Section 3.5. In accordance with usual qualitative practice, the Findings detailed in section 3.4 incorporate direct quotations from participants.

---

<sup>11</sup> <https://www.qsrinternational.com/nvivo-qualitative-data-analysis-software/home>

## 3.4: Findings

The findings below seek to address the concerns of the research question, and also to discuss additional areas of interest.

*RQ What are the particular features exhibited by the quantum computing landscape that may differentiate it from other novel technologies?*

Some of the findings are relevant to only one strand of the triple-helix (policy-industry-academia) model, some to all three. The cross-cutting issues are listed below, but for clarity will be examined in relation to their position in the triple-helix model (eg all the policy-relevant findings will be discussed together in the policy section).

### Cross-cutting issues

The themes of capacity, engagement and financial investment were visible in all three strands of the triple-helix, with different emphasis or significance depending on the strand, but highly interdependent.

	Academia	Industry	Policy
<b>Capacity</b>	<ul style="list-style-type: none"> <li>• Difficulty of recruiting staff when competing with general shortage of trained personnel</li> <li>• Frequently losing staff to more profitable or better-paying startups</li> </ul>	<ul style="list-style-type: none"> <li>• Shortage of trained staff proving to be a limiting factor on expansion</li> <li>• Challenge of finding more rapid routes to qualification</li> </ul>	<ul style="list-style-type: none"> <li>• Providing funding and support to increase places despite no certainty of success in the technology</li> <li>• Challenge of anticipating demand over long timescales</li> </ul>
<b>Engagement</b>	<ul style="list-style-type: none"> <li>• Researchers need to listen to publics in order to understand concerns</li> <li>• Attitudes that quantum is 'too difficult' to understand miss the point.</li> <li>• Little systematic engagement with industry or policy</li> </ul>	<ul style="list-style-type: none"> <li>• Little acknowledgement that industry has a role to play in either innovating responsibly or engaging with publics</li> <li>• Engagement with academia on ad hoc, unsystematic basis</li> </ul>	<ul style="list-style-type: none"> <li>• Understanding that public engagement improves likelihood of acceptance, but little systematic work on this</li> <li>• Hedging risks by engaging with academia and industry for strategic planning</li> </ul>
<b>Funding/ investment</b>	<ul style="list-style-type: none"> <li>• Need support for knowledge transfer/ spinouts &amp; start-ups that can feed development of quantum sector.</li> </ul>	<ul style="list-style-type: none"> <li>• Difficult to find long term, 'patient' funding that admits of possibly high number of failures</li> </ul>	<ul style="list-style-type: none"> <li>• Support for research can be incorporated into Spending Reviews, but more challenging to justify government procurement.</li> </ul>

### 3.4.1: Policy: shaping the landscape

As discussed in Chapter 2, quantum computing is an extremely specialised set of fields – this means that it requires not just highly-trained staff, but also significant infrastructure. These requirements alone would make it imperative to interview policymakers, as enabling these capacities is a policy-level concern at varying levels of granularity (for instance: whether a field of research is prioritised; how much funding is made available and how it is allocated; what criteria are used to assess proposals, and so on). Additionally, however, there is a scale of ambition for quantum computing and broader quantum technologies that makes the policy perspective essential to understand. These ambitions mean that at the policy level there are active and wide-ranging efforts directed towards the creation of a quantum sector – as with Machine-Learning it is seen as an enabling technology and, although use-cases may be currently unclear, participants understand that *“it’s a significant opportunity for the UK – we’re good at it”* (P19) and *“we don’t want to be left behind”* (P05). Playing into the context behind many of these decisions are national prestige, geopolitical, and economic concerns.

#### Building a quantum economy

Participants from policy fields believed that when it comes to quantum there needs to be more of a holistic and overarching view than in some other fields – *“we look across all the technologies for quantum”* (P19) – and that there were conscious efforts to look ahead and understand the ways in which the field could be supported.

*“we identified some barriers to growth ...and then tried to think about what are the things that we need to do from a government perspective to support growth” I19*

This incorporated many different kinds of action:- *“it’s kind of a whole package of things that we feel we need to get going”* (P19). This might include: providing infrastructure (for example the National Quantum Computing Centre, supported by the Science and Technology Facilities Council); ensuring funding (such as creating and supporting the Hubs: EPSRC, 2014); and encouraging ‘knowledge transfer’ by providing support and incentives for the translation of foundational research into spin-out companies. There is a strongly apparent sense of needing to support the whole, integrated, development trajectory, and an awareness that the creation of a sector from scratch is a very granular activity.

*“a big part of the strategy is to make sure that the [research teams] take things to the next level and push towards commercialisation” P05*

There is, however, also a recognition that the UK’s highly fragmented and silo-ed policymaking system – *“we have lots of different government departments and lots of different regulators”* (P20) – can find

it a significant challenge to pull together the type of unified approach needed for these holistic approaches.

*“how you bring that policy together in a cohesive way that everybody can benefit from, including the country and including consumers, becomes quite challenging to be honest”*

*P20*

These efforts to build and support a quantum sector are also linked to underlying theories that draw together ‘prosperity’ and ‘security’ concerns. Such theories suggest that some of the policy concerns around the development of the technology – such as economic advantage, national interests, and security – are intrinsically intertwined. They may also translate into an understanding that government has many levers available in the policy sphere. In this instance, the combination of national security and national prosperity is known as the Fusion Doctrine.

### Fusion Doctrine

Policy and policy-adjacent participants were clear that quantum technologies, and particularly quantum computing, are seen as an example of the ‘Fusion Doctrine’ (Cabinet Office, 2018) – an area where significant ‘security’ and ‘prosperity’ concerns are melded and therefore require a holistic perspective. The Fusion Doctrine, deployed for the first time under Theresa May’s government, brings this understanding to the fore and renders it concrete in policy terms. The impact of the doctrine will be discussed in greater detail below, but fundamentally this is a recognition at policy level of some of the concerns that arise around dual-use technologies, of which quantum computing (given its cybersecurity aspects) is one.

*“that’s a political decision that’s being made, which is quite clear now, and everything has to pass both ... of those tests of whether it’s economically viable but also is it secure.” P27*

*“there are many reasons to justify investing in technologies, one of which is prosperity and the other is security, and ... efforts are being made to join these two together, because maybe in the past, they’ve been thought of in different buckets.” P19*

This has implications both for investment and for governance, as the Fusion Doctrine recognises that – both in terms of national security and national economic interests – government has many possible routes to its desired ends. The Fusion Doctrine and its associated policy implications could therefore be characterised as an exercise of ‘soft power’<sup>12</sup> in the national interest.

---

<sup>12</sup> *“a persuasive approach to international relations, typically involving the use of economic or cultural influence.”*Oxford Languages

## 'Sovereign capability' and the national interest

Another aspect to the efforts to build a 'quantum sector' described above is the national interest on the global stage, in terms of comparison with international allies and competitors. There was an ambitious theme within interviews with policymakers, of the UK as a 'powerhouse' for scientific discovery and development, and the strategic advantage this might confer. There is a vision of the UK as a 'science superpower' which is seen as a matter of significant national prestige but also economic advantage. This may be particularly crucial in the post-Brexit period as the UK seeks to distinguish itself internationally. The corollary to this was an awareness that if efforts were not sustained to retain such a position in the vanguard of world-leading research, the UK could very quickly lose out to rivals. Quantum computing, with its dual-use characteristics, is seen as a key area in which to try to maintain the UK's position at the forefront of research and development.

*"it's seen as one of the kind of key strengths of the UK to be a leader in technology development." P19*

*"government has gotten involved as well, because they can't sit at the sidelines anymore, because ... other countries are absolutely treating this as a very sovereign capability issue"*  
P27

The framing of quantum computing as a 'sovereign capability' question is another aspect of the Fusion Doctrine – it generates a protectionist perspective, which is reflected in other policy priorities. For example the National Security & Investment Act (National Security and Investment Act, 2021), which passed into law in the spring of 2021, allows the Secretary of State for the Department for Business, Energy and Industrial Strategy (BEIS) to 'call in' for inspection any investment in a selected list of technologies if they believe there may be national security considerations. This is not least because for a large corporation, investing in (or buying out) a company is a quick method of obtaining its research and development base.

*"for Alibaba, you know, a million or two million is nothing ... but if Alibaba gets the IP then that's it, they don't really need that company anymore" P27*

This particular piece of legislation has been criticised as being business-unfriendly and "a powerful deterrent to foreign direct investment" (Fingleton, 2021), with the retroactive facility in particular drawing fire for giving government the ability to repudiate an investment in a company up to five years after the fact. The power was used by Kwasi Kwarteng in 2021 to investigate an investment in a graphene startup<sup>13</sup>. However, it is not only the UK that regards such protectionist initiatives as important – other countries and blocs have similarly-focused legislation (for example the multi-state

---

<sup>13</sup> <https://www.theguardian.com/business/2021/sep/07/uk-government-orders-national-security-review-of-graphene-firms-takeover>

Wassenaar Arrangement,<sup>14</sup> and the ITAR regulations in the US<sup>15</sup>) in technology areas considered sensitive or particularly crucial for national security or to maintain a country's position on the international stage.

*“technology is increasingly being used globally as a kind of a strategic advantage” P19*

### Investment and capacity

One serious hindrance to this question of the UK's international positioning and the development of a quantum sector or 'science-superpower' status would be a lack of investment. Despite the very large sums that have been invested in research, a shortage of sufficiently granular development support shows itself as an inability for small companies to obtain startup or expansion capital, and an inability to source staff either in industrial settings or academia because the funding has not been in place to support PhD programmes. This leaves the sector playing catchup as expansion is limited by funding and staffing issues.

*“the bigger issue for me is how do you keep up with that demand for people as it goes forward” P27*

*“enabling companies to get access to finance ... is a key issue ... because there are concerns that the UK market just doesn't have the right sorts of funds to continue to invest.” P19*

This challenge demonstrates clearly some of the trade-offs that have to be made at the policy level, as policymakers try to retain the balance between creating conditions that make the UK attractive as an investment proposition and can produce a thriving sector, while also protecting what is perceived as the UK's intellectual property. These are not solely issues of financial provision – policymakers need to work with academic and industrial researchers to try and understand how the landscape may develop in several years time, and plan for those capacity needs ahead of time, including where investment should be focused. For example the funding for the 2019 Centres for Doctoral Training, some of which are focused on quantum, was agreed in 2017 and would have been extensively consulted on prior to this, but the first cohort of these new postdoctoral researchers will not qualify until 2023.

*“how do you map what **these** people are theoretically designing, with what **these** people are practically trying to build ... that's the puzzle for us as government is 'how do we make that happen?’” P27*

This long-range planning is necessarily something of a gamble, as either technology or markets may pivot rapidly in an unanticipated direction, so policymakers therefore need to try to build in sufficient adaptivity and resilience to be able to manage these risks that may not transpire for a significant time.

---

<sup>14</sup> [https://en.wikipedia.org/wiki/Wassenaar\\_Arrangement](https://en.wikipedia.org/wiki/Wassenaar_Arrangement)

<sup>15</sup> [https://en.wikipedia.org/wiki/International\\_Traffic\\_in\\_Arms\\_Regulations](https://en.wikipedia.org/wiki/International_Traffic_in_Arms_Regulations)

## Engagement

Managing these strategic planning risks is therefore a key part of the policymaking role, and engagement with industry and academia is part of how these risks are hedged. Policymakers carry out horizon-scanning activities to try and foresee where there will be demand for capacity and attempt to consult widely to identify tensions, challenges, and address concerns. Within a relatively small and focused field such as quantum computing, it is relatively easy to consult with a large percentage of those working in this area.

*“everybody who's connected with quantum computing ... was part of that, so that's industry, academia, government – these very small players so startups, scaleups, spinouts, also the big players so the Googles, Microsofts etc”*

However, it is easy to conceive that with a more widespread, and less defined technology such as Machine Learning, it would not be possible to consult with such a high proportion of the whole field, rendering consultation exercises that much more difficult to manage effectively.

Engagement and consultation are also valuable tools for managing other risks, such as public pushback against technologies. Opening up dialogue with citizens is a way to try to shape and guide public awareness and response but also to encourage researchers to consider their responsibilities to the public, who are both the ultimate funders and end-users of research.

*“I saw the importance of engaging with the public. And I saw the importance of being responsible for your research and being able to explain what you're doing to people” P05*

This can be particularly crucial in fields such as quantum technologies generally, where public awareness of the technology is relatively low.

### 3.4.2: Academia: progressing the science

The academic ecosystem can be thought of as the central strand of quantum computing development, around which are wrapped the policy concerns, alongside the funding, and the industrial or commercialised elements. There is understandably more of a focus in this strand on foundational science – however, many of the concerns seen at policy level are also seen within the academic research environment, albeit with a different emphasis.

#### (Non-)neutral technology

Researchers in many STEM fields often see their work as ‘neutral’ – ie research that is carried out from motives of curiosity – *“It's just interesting.”* (LS1). Quantum computing researchers are little different in this regard, with a frequent trend of viewing quantum use-cases as entirely separate in terms of responsibility from the foundational research that led to those use-cases – responsibility concerns are

in effect ‘thrown over the wall’ to later stages of development. Interviewees here demonstrated a range of views on whether or not it was the responsibility of a researcher to consider possible impacts.

*“You don't want to say, well, we've done this beautiful theory and it could be used to destroy the world.” LS3*

*“I'm building the device. So how it's being used is a secondary question.” LS6*

The view of technology as necessarily ‘value-neutral’ – in that it cannot be either ‘good’ or ‘bad’ – is often positively reinforced by years of training. This training may combine with a lack of input from, and understanding of, other disciplines and fields to create a type of disciplinary ‘bubble’. Greater familiarity with social science, for example, might tend to suggest wider or different perspectives be incorporated – eg Winner (1980) on whether artefacts can be inherently political. Approaches such as RI, with their emphasis on stakeholder consultation, seek to provide alternative perspectives that can support these broader viewpoints. RI work can also increase understandings of ‘who are stakeholders’ – for example the concept of ‘stakeholders’ is not necessarily limited to ultimate users, and during early-stage development ‘stakeholders’ might be peers within a research team, or those in similar fields.

### National concerns

Participants in academic research, particularly those that are also involved in start-ups or other industrial and commercialised enterprises, perceived the same trend in quantum computing fields towards ‘national’ concerns as discussed by policymaking participants: *“it's becoming a bit geopolitical”* (A11). This direction of travel was concerning for many participants, as a ‘race’ to quantum computing is not necessarily beneficial for the field as a whole.

*“there's a ...big US-China rivalry that's ...developed. And we're also seeing it quantum computing” IA10*

There was also a degree of relative scepticism about whether the UK would be able to compete in the long term with organisations, countries, and international blocs that have substantially larger budgets, in terms of building a scalable quantum computer:

*“America wants one, and China wants one, and ... you can only have a certain number” IA14*

However, there was agreement that in many areas of the foundational science, the UK currently occupies a genuinely world-leading position and is in a good position to both retain and leverage this position with the right level of investment. In terms of a model, other small countries with focused and advanced technology sectors include Japan, Singapore, and Taiwan, and this type of approach was seen as a more viable long-term position than being the first country with a scalable quantum

computer. It will not be possible to sustain a position at the top of the world's quantum computing research ladder, however, without significant increases in capacity.

### Investment and capacity

Although those with experience running quantum startups reported issues with recruitment, it was in the academic research environment that the capacity element seemed most strained.

*"We have a very substantial challenge in recruiting people, because there are some very attractive opportunities in companies" A29*

*"at the moment there's a scarcity, not of money but ... of talent" A07*

This is not just due to a lack of direct investment (for example in doctoral training), but also because academic environments can be extremely demanding and by contrast the recompense of commercial work can be significantly greater. Academia has traditionally not been well-rewarded – reputation is often regarded as more important than financial recompense (Ten Holter, 2020) – but the quantum sector is unusual in the extremely high number of researchers who are either participating in startups or have left academia entirely to move into commercial organisations. There may be many reasons for this, including the better salaries available in industry and the fact that quantum computing requires a high degree of training, but the net effect is that academic projects frequently find it very difficult to recruit.

*"we're already seeing this at senior level, where... many of the people... have been snapped up by large companies and ... start-ups... We have several projects ... that we're really struggling to hire people for" IA10*

Conversely, there were concerns expressed that the availability of large amounts of research funding was skewing the marketplace – not only creating opportunity costs (in terms of what else does not receive funding as a consequence) but that the sums available were encouraging people to give their projects a "quantum-y" (A11) slant in order to be able to respond to funding calls. Some noted that large awards had generated jealousies and tensions in a field that until fairly recently had been a relatively small and closely collaborative community. These outcomes may be viewed as undesirable but need to be balanced by alternate concerns that insufficient investment over a number of years has led to the capacity challenges being seen in academia.

*"lots of places want to fund this research but they're not finding the people to hire ... because ... governments were not funding the PhD students and postdocs to work in this area." IA10*

There were perceptions that although the UK is an acknowledged leader in fields such as quantum computing, the country does not have a strong track record with leveraging its position.

*"the UK is very good at creating stuff and then selling it off or giving it away" A15*

In particular, several interviewees wanted more engagement at earlier stages of education. They expressed frustration that the basic tenets of quantum physics are not part of the National Curriculum and, as a consequence, students commence their University courses without any grounding in quantum physics. They considered it obvious that this could make a real difference to the level of work that students were able to achieve at an earlier stage of their careers.

*“high-school education hasn’t actually caught up with what are very well-known and well-understood concepts now in quantum mechanics” LS8*

Capacity is also an issue in terms of the diversity of the field. This has been highlighted recently by debates around whether girls ‘enjoy’ physics (amid accusations that they find the associated mathematics ‘too hard’ <sup>16</sup>) and the widely acknowledged problem of the ‘leaky pipeline’ (Pagel, 2022). Ethnic minorities, too, are hugely under-represented in quantum fields. As seen in Chapter 2, diversity is crucial for responsible innovation approaches. This element will be revisited in more depth in the later studies.

The challenges around recruiting and retaining staff to work on projects in physics and engineering, and in particular the shortage of women and minority academics, are likely to have long-term effects on the field. However, it is possible that wider public engagement could spur greater interest in the field and a deeper awareness of the opportunities that a growing sector represents.

## Engagement

In general, engagement is not a well-understood concept among researchers, particularly as it relates to responsible innovation concerns. In section 2.3.2 above, it was seen that there are many challenges for engagement, but a fundamental one is a lack of understanding about who might be termed a stakeholder. There was often a preconception amongst participants that ‘engagement’ necessarily meant public engagement activities, which in general were regarded with mixed views. There was awareness among some academic researchers that citizens may not only have an interest in the work that they are carrying out, but that the type of work they are doing carries with it a responsibility to engage, to explain, to address concerns and to listen to publics.

*“we must answer questions of society about any concerns or worries or hopes that they may have for this technology” A07*

Interviewees had different levels of experience with public engagement, and this tended to affect their view of its value. Reflecting the work in synbio from Pansera et al., (2020), which found that there could be a learning-curve associated with becoming acculturated to engagement, some

---

<sup>16</sup> <https://www.bbc.co.uk/news/uk-politics-61247374>

researchers (especially those involved in the more foundational end of the research spectrum) were sceptical about engagement, believing that non-specialists could not conceivably have an informed view.

*“it would be a mistake to imagine that ... even a very alert and engaged individual is in a position to contribute ... thoughts about the ... direction a researcher can go in, because you are just ultra-simplifying things ... so that they can make decisions” LD3*

This conviction among some researchers that it is not feasible to open a dialogue with stakeholders because they are not trained in the relevant technology also fails to recognise that it is not the object of engagement to be told how to develop the research. Instead, as discussed in Chapter 2, the purpose is to draw on the varied expertise of a diverse group of stakeholders to anticipate and prepare for a range of outcomes. However, again echoing Pansera et al (2020), more often there was agreement among researchers that they had frequently been surprised by the level of self-education and interest among the general public.

*“they’re often coming up with detailed high-tech questions because they already know about some aspect of programming or computers or quantum” LS1*

This understanding was still focused on the ‘public engagement’ aspect, however, whereas a true responsible innovation approach adopts an inclusive approach to stakeholders, recognising that they may be drawn from many areas and will bring a range of perspectives to the table.

Among some participants there was also understanding of the role of engagement in creating a deeper connection between science and society, as discussed in section 2.1. The purpose of such a connection would be to improve society’s ability to engage with broad questions on how science and innovation could affect them, their communities, and interests, as well as to enable the work of anticipating and preparing for possible impacts. Some participants showed awareness of the reciprocity of such a relationship – the importance of a two-way conversation – and also how such a mindset would encourage researchers to reflect on the end-uses of the technology or innovation they were working on.

*“if you’re working on that thing which might ... change society, ... maybe the duty is... for a sort of communication from researchers to the wider public, so that everyone can then discuss that thing.” LS9*

This engagement can be an important element of raising public awareness and – ideally – acceptance, as well as providing researchers with direct experience of public concerns and tensions. One key public concern, discussed in Chapter 2, is that of governance – the public requirement that quantum technologies, and quantum computing in particular, should be subject to good governance that could support public trust.

## Anticipatory governance

Linked to the engagement theme is the wider concept of anticipatory governance – exemplified in this context through responsible innovation practices. Academic participants were often very aware of their own inexperience with responsible innovation and its processes – although many were familiar with the term it was not clear to them how they should carry it out in practice, or its relevance to the very theoretical or foundational work in which they were often engaged.

*“I’d love to have an RRI red phone, like a Batphone, and say ‘drop what you’re doing, we’ve got this issue, how do we approach it?’” A15*

However, they often connected responsible innovation practice with questions of impact, demonstrating that even without a deep understanding, there was a consciousness of end-goals and societal need. The theme of *anticipatory governance*, although relatively unfamiliar to most participants as a term, was familiar as a concept, as researchers understood that they could look forward and imagine possible futures, at the same time as accepting that future applications would emerge which could not be predicted. There was a conflation of anticipatory work with ‘prediction’ – as discussed in section 2.3.1 this is a frequent misconception in responsible innovation, which uses participatory approaches to prepare for numerous possibilities rather than ‘picking a winner’.

*“that’s the continual lesson ... there’s just example after example after example of how impossible it is to predict what’s going to make the ... biggest changes to society” LS4*

*“we do talk, what is the best way to move forward and what would be the best mode of operation. More as a brain exercise at this stage, but these are important issues to discuss.” LS5*

These “brain exercises” are a useful practice nonetheless – partly because this implies that researchers are reflecting on their work, but also because a further point to emerge from the discussions around the uses of technology is the challenge of identifying the appropriate juncture at which anticipatory work can be most effective.

Although a responsible innovation approach argues for a consistently anticipatory and reflective mindset, there are doubtless key decision-making points at which such activities become more than usually crucial. Stilgoe et al., (2013) discussed the use of ‘stage-gates’ in the development of technology – built-in pause-points at which a technology under development should be reviewed to ensure it could pass through the ‘gate’ to the next stage – but these are not a common tool in research and indeed researchers may find it extremely challenging to identify such points in their work. However, comments such as the one above demonstrate that participants who may be unfamiliar with specific responsible innovation concepts are often articulating what may be termed *de facto* responsible innovation.

### *De facto* responsible innovation

As discussed earlier (section 2.4: *Institutionalisation? Embedding RI in research*) the idea of *de facto* responsible innovation is well-recognised (eg Randles, Laredo, Loconto, Walhout, & Lindner, 2016, Lindner et al., 2016; Porcari et al., 2021). It was clear from the discussions with some participants that responsible innovation-type processes were familiar, even if they had not previously considered or framed them as ‘responsible innovation’.

*“society and the scientific community is more and more aware that we need to communicate where science impacts ... society” LS5*

*“you can never predict quite where research is going to go, but the next bit, which is much harder, is ... we've identified something, what procedures should we put in place?” LS3*

Considerations around engagement, reflection and anticipation were also familiar. However, researchers frequently lacked a structured method of drawing these threads together in a conscious way into a cohesive process where engagement with stakeholders and consideration of outcomes could inform each other and, ultimately, a response. These grassroots understandings align with the Randles & Laasch (2016) description of the ‘ground-upwards’ type of RI institutionalisation described in Chapter 2.

### 3.4.3: Industry: creating a sector

#### Investment and capacity

Towards the innovation and commercialisation end of the technological development process, the same concerns were apparent in terms of there being insufficient funding available to truly enable and support a flourishing quantum sector. There were also anxieties that the capacity shortages being seen in academia could well be carried forward to commercial sectors in the future, meaning that start-ups may struggle to expand. This is also affected by concerns that lack of investment will make it challenging to retain talent in the UK – foreign or private investment may encourage teams to move to other regions – creating further capacity issues. Some participants had seen this happen in areas close to their own work.

*“he's raised \$400 million... and he's taken the workforce out of the UK.” I15*

Capacity is also not the only issue caused by a lack of investment. A further domino-effect concern was that this type of research – once it starts to be increasingly carried out within private companies – becomes subject to many more limitations that are imposed by the need of large corporates to protect their intellectual property, as they operate in a highly competitive environment.

*“there are probably increasing numbers of programmes ... in various regions, where it's not ... visible what researchers are doing anymore.” I07*

Also expressed as a potential issue was the likelihood that – should private companies become the primary investors in this potentially extremely high-value sector – the impact on societies was likely to be less beneficial than if control over quantum computing capacity is widely and equitably shared, supported by public investment.

*“if the result is that these companies end up being the ones that have the only quantum computer in the world ... then this is not necessarily a good thing for humanity” I10*

There were concerns that as well as a potential lack of investment and the various issues discussed here that would result from that, academics who are the primary source of start-ups in the sector may have a lack of ‘soft’ skills and business acumen, which would put them at a disadvantage. This type of support is also seen as something that could be supported by public investment. Many participants took the view that government needs to be a major early investor in this sector due to the long gestation period that quantum start-ups are likely to have before returns can be made. Additionally, like any nascent sector, there will inevitably be a high percentage of failures. Investment is needed for these relatively high-risk, slow-burn start-ups from investors with deep pockets who are willing to potentially wait a considerable time for a return on investment.

*“an industry has to grow and then be self-sustaining and expanding, but you’ve got to support it until then and you’ve got to give it confidence - so government, in my view, needs to be the first customer.”A15*

However, there may be apprehensions that this type of approach places government in the position of having to ‘pick winners’, and the issues that could arise if companies should fail after receiving taxpayer funding.

### 3.5: Discussion

This landscape study demonstrated the particular importance of understanding the policy perspective, which provides an important framing to the development of quantum computing within academic research fields and industrial development. Indeed, as has been shown, it is critical to understand the policymaking context in order to appreciate the ultimate source of many of the impacts that are perceived by academic and commercial interests on the development of quantum computing technologies – these impacts also ripple out into the public domain. Unless the drivers and motivations operating in the policy space are taken into account, the way in which quantum computing technologies may develop both in the lab and in the marketplace cannot be fully understood. These factors give quantum computing some distinguishing characteristics that set it apart from some other novel technologies. In particular the number of national-level variables affecting policy decisions must be explored, along with the context that they provide for the industrial

arena and the experience in academia. The policymaking context is also important for understanding the perspective on the importance of engaging with citizens and publics – for such a technology, which is deemed to be of national significance, it might be expected that policymakers are conscious of the need to establish high levels of trust and confidence, and take seriously public requests for good governance.

The effects on various ‘levels’ of the relationship between science and society must also be distinguished. There is very significant work on the development of quantum computing taking place within all three strands of the triple-helix, but in the absence of a comprehensive engagement programme, this work is by and large invisible to general publics. There is little opportunity for societal involvement in the development of this crucial technology, whether it is to express concerns or simply to learn about it – as Coenen et al., (2022) point out, terminology such as ‘entanglement’ and ‘superposition’ should be matters of everyday parlance, *“Instead of trite appeals to Schrödinger’s cat”* (Coenen et al., 2022, p3). It is here that participatory methodologies may prove useful in offering an overarching framework of engagement, anticipation and response to the possible impact of quantum computing on society.

### 3.5.1: Societal impacts

In some respects, the potential impacts on society are already being considered, particularly those that affect security. The dual-use nature of quantum computing has ensured that it has been within the purview of organisations such as the MoD, NCSC and GCHQ for some time, and an anticipatory governance approach has supported the development of post-quantum cryptography (although this is not yet especially widespread). However, interviewees were clear that these consultations and anticipatory considerations are not necessarily being applied in wider societal contexts: *“What are we ... missing here? When it comes to societal impact, whether it’s economic or quality of life?”* (I07)

Interviewees were almost uniformly aware of a sense of urgency in terms of considering the possible impacts of quantum computing technologies on society, and also displayed a strong awareness that no organisation or group is currently tasked with either undertaking or drawing together work on societal impact. There is research and development taking place on an increasingly broad front and in numerous new areas, but there appears to be little cohesion. Indeed, even though quantum computing does not yet exist at a useful scale, it is already having impacts due to the resources that have been committed to it (and the opportunity costs that those commitments represent); the increasing apprehension about quantum ‘hype’ in both academia and research; and a rising concern about lack of governance – anticipatory or otherwise – among quantum computing developers (eg Khan, 2021).

### 3.5.2: Connecting policy and academia

The study revealed many areas of challenge around capacity in the field, one of which is the necessity for extremely distant long-range planning by policymakers. To take only one example, if postgraduate capacity shortages in a given discipline or industry are identified, then a number of policy decisions – influenced by many variables – must be taken in order to respond to this and for funding to be made available. If funding is made available for a Centre for Doctoral Training, for example, (which is not by any means a certainty) then it must be applied for by a training body. Even moving at pace, it may be 18 months after that (successful) application before the first cohort begins training – that cohort will likely then take a further four years to graduate, and the last cohort will not have finished their training until eight years after the first cohort started. There could therefore be up to 10 years between identification of a need, and a sufficiently large supply of postgraduates. Strategic planning at such a distance is therefore likely to require significant adaptivity to be built in through the course of the policy deployment. There are also likely to be tensions throughout the process about sources and directions of funding, not to mention political influences and implications. These challenges are likely to be particularly acute in a field that requires the depth of training that quantum computing does. Deeper investigation of this and its possible effects on the development of the field would be valuable.

This work also highlights the importance of the pathways between research and policy and the need for access from both sides. Policymakers rely partially on academic expertise as sources of information in the formation of policy, but such consultation mechanisms frequently draw on the expertise of senior academics, who may be invited to give evidence to parliamentary groups, or sit on committees. However, not only will these views represent only a small fraction of the views available, they will inevitably be filtered through a number of competing priorities – as work on consultation-type initiatives like ‘Sciencewise’ shows, it can sometimes be very difficult to influence policymaking (Smallman, 2018). An example is the need identified in section 3.4.2 for quantum to be taught as part of the national curriculum – but there is no direct route for researchers to suggest this. Another challenge is that senior researchers will have their own pressures, such as the need to maintain a flow of funding into their research group, which may affect their role in consultations.

A further example of the impact of high-level policy on research is the requirement for some commonly accepted standards within quantum technologies that can ensure interoperability – “*no standards: no interoperability*” (L11). Without such interoperability, not only do developers tend to create silo-ed work, but the downstream choices of purchasers are necessarily limited because they may become locked into particular procurement choices. This has implications for the operation of the marketplace. Standards also create efficiencies and certainties, eg “*The buyer has a certain assurance that what he buys is what he thinks he buys*” (L11)

At present there are not obvious routes for researchers to pass such concerns on to policymakers, and so important considerations may be insufficiently weighted or lost entirely.

### 3.5.3: Prestige

Questions of national ‘prestige’ may seem to lack relevance at the research or commercial level, but financial markets, and indeed government finance, are significantly affected by the degree of ‘confidence’ that markets and banks have – whether that is in a company, a country or a technology. So although the notion of national prestige may seem peripheral to academic research, its correlation with national prosperity and security – the fundamentals of the Fusion Doctrine – may mean that it has a significant influence on the level of funding available to a technology such as quantum computing. This also links to tensions seen around questions of the size and direction of current funding, as discussed below.

### 3.5.4: Tensions

One area of tension is between the pervading free-market economy ethos in UK government, and the desire of policymakers to have some control over the ‘quantum sector’. This means that despite the sector having been identified as important to the UK economy and receiving ongoing investment, free-market *laissez-faire* means that policymakers do not intervene further along the pipeline. For example there may be little incentive for those trained by the public sector to remain in academia, and private investment can make a move into the commercial sector very appealing, causing continual capacity challenges within academia. There may also be a lack of joined-up thinking in the way tax-incentives are formulated. It is very difficult to ‘grow’ a sector under these circumstances, and it may prove necessary for policymakers to create more structured incentives for national investment – especially as the kind of protectionism that led to the National Security & Investment Act may act as a significant disincentive to inward investment.

Tensions are also visible between national initiatives such as the National Quantum Computing Centre, and distributed programmes such as the Quantum Computing Simulation Hub and university courses – there is a perception that, rather than acting as a focal point for national efforts, the NQCC may simply be an additional competitor in the same marketplaces for resources, including staff and funding, exacerbating the capacity challenges. Academic incentives, too, may lead to fracturing of the field as team-leaders create their own labs and research directions. There is little evidence, as yet, of the NQCC drawing together the numerous areas of the UK’s quantum ecosystem.

As noted in section 3.4, funding is a particular point of tension. Not only are participants in the academic and commercial areas concerned about levels of funding and the pressured hype-cycle they

may be encouraging, but tensions were visible about the wisdom of the current focus of funding programmes, with some viewing the effort to build a quantum computer in the UK as ultimately a waste of resources and a mistaken desire to be seen as a global leader. Additionally, the availability of large sums for research and/or development creates not only incentives but extreme pressure to deliver results.

## 3.6: Conclusions

The landscape study that has been the subject of this chapter was designed to address the research question: *What are the particular features exhibited by the quantum computing landscape that may differentiate it from other novel technologies?*

In order to examine this, the study sought to sketch the quantum computing ecosystem in academia, in policy, and in industry, investigating the relationships between these strands and how those relationships manifest, the state of readiness in terms of public engagement with quantum computing technologies, and any unique features of quantum computing that might set it apart from other novel technologies.

The study has demonstrated that while many factors tend towards the view that quantum is ‘set apart’ from other novel technologies – even some of the technologies with which it is sometimes grouped as being ‘of interest’ at the national level (such as AI) – it has enough characteristics in common to serve as an exemplar for the manner in which novel technologies can be embedded into society. One of the most important factors (and one which it shares with AI, robotics and other fields) is the perceived importance of the field as an element of national security.

### 3.6.1: Security concerns

The development of quantum computing is closely monitored by agencies such as the National CyberSecurity Centre, Government Communications Headquarters, the Ministry of Defence, the Centre for the Protection of National Infrastructure and others. Those working in the field cannot but be aware of this interest and can anticipate the effect that it may have not only on funding streams in the future, but on questions such as *what can be published*, and *where, or from where* they are allowed to hire staff, or receive investment monies. Quantum computing’s profile as a research area characterised as important for technological sovereignty (and the attendant concern to become less reliant on US- and China-based tech) makes it one of only a handful of ‘enabling’ technologies seen as crucial in the UK context.

### 3.6.2: Academic entrepreneurialism

This small study engaged with a relatively limited number of academic researchers working in the field, yet encountered several within that group who were involved in start-up companies, along with an air of urgency around the possibility of missing opportunities (eg *“Everyone’s got a company. I need to get one!”* A15). Although this could not be regarded as a representative sample, the general trend of researchers becoming involved in startup businesses of various sizes indicates that the potential rewards (particularly when compared to the rewards of academia) are seen as extremely attractive in a way that other novel technologies, such as synthetic biology, are not.

This is also indicative of the lengthy and technical training that is required in order to engage with quantum computing – it may simply not be possible for those without an understanding of quantum mechanical principles and the challenges of creating a quantum system, to start up a business enterprise. This has implications for the creation of the quantum sector, because it suggests that the academic pipeline may be the only mechanism whereby new companies can be created – and the capacity challenges seen here thus represent a significant ‘pinch point’ for the commercialisation of quantum computing. It is likely that these capacity issues will need addressing not solely through an increase in funding for doctoral students, but through creating new mechanisms that can attract and train personnel from other disciplines (for example the use of conversion courses).

### 3.6.3: A quantum ‘bubble’?

The scale of global investment in quantum computing, its framing as a matter of sovereign capability, and the potential instability of an economic sector that is seeing high-profile mergers and acquisitions<sup>17</sup> without accompanying governance, legislative frameworks or provable technology combine to create what has been described a ‘Wild West’ environment<sup>18</sup>. Levels of excitement such as this have been seen before in technology sectors, and they are frequently followed by periods of disappointment and retrenchment, as with the various AI ‘winters’. When a ‘bubble’ of hype is large enough, its bursting can cause serious economic consequences and this may also generate societal impact – a parallel could be drawn with the dotcom bubble of the 90s. In that case stock market crashes in response to overvaluations and market hype caused not only economic shrinkage (the NASDAQ lost \$5tn in value<sup>19</sup>) but loss of general confidence in tech companies and subsequent land-

---

<sup>17</sup> <https://www.forbes.com/sites/moorinsights/2021/06/08/honeywell-quantum-solutions-and-cambridge-quantum-computing-merge-and-plan-to-go-public-by-end-of-year/>

<sup>18</sup> Ilyas Khan, CEO of Quantinuum, speaking at the UK Industry Forum on 18 November 2021

<sup>19</sup> <https://www.latimes.com/archives/la-xpm-2006-jul-16-fi-overheat16-story.html>

grabs by the surviving companies. It is clear that – before it is even a proven technology – the development of quantum computing is creating societal impact.

### 3.6.4: Implications for further case studies

This study will help to shape the research approach to the remaining case studies, and also for the development of the Framework.

#### RI study

The study on responsible innovation in practice will examine in more depth the concept of *de facto* RI, and whether the degree of cohesion that a tool such as the AREA Framework can provide is indeed valuable. It will take the opportunity to assess different interpretations of the meaning of RI in different fields, allowing the operationalising of such anticipatory governance methods in widely varying disciplines to be compared and contrasted. It will also interrogate how a novel methodology such as RI can be institutionalised, and the enablers and constrainers on such a process.

#### Policy study

This case study revealed the degree of fragmentation and silo-ing within policymaking – different departments are not always communicating well with one another and there is a degree of tension over areas of jurisdiction. This could have significant impacts on the effectiveness of governance and oversight of a novel technology and bears further investigation – this fragmentation could be highly significant in terms of any attempt to create cohesive governance. Quantum computing's categorisation as susceptible to the Fusion Doctrine also mean that policymakers are likely to be faced with challenging tradeoffs between economic and security concerns. Quantum policy has both internal concerns, such as prosperity, and economic security, and also external concerns – international security and the 'prestige' factor discussed above.

In general, given the direct and immediate influence of policy on this area – for example through the NS&I Act – there is a need to investigate how decisions are made in relation to which technologies are supported; infrastructure requirements and capacity; governance and oversight; availability of investment (whether from government or other sources); and the development of a 'quantum economy'. In particular, given the earlier discussion of the difficulties inherent in stakeholder participation, it will be valuable to investigate how policymakers choose whom to consult on these questions, and the implications of those choices.

## Industry study

The pool of potential participants in a study of the industrial context for quantum computing is relatively small (though expanding fast), and the findings of this case study will influence discussions with industry participants. In particular it is necessary to examine the relationship between academia and industry; as noted above, it may be that the only realistic path into industry is via an academic route, making the necessity for training and expansion in the academy even more acute. It will be beneficial to investigate as much of the industrial ecosystem as is practicable, including elements of infrastructure such as funding, knowledge transfer, and recruitment.

### 3.6.5: Implications for the Framework

The findings from this case study will help to inform the building of the Framework. In particular, it draws attention to the consistent nature of tensions – it will be important that the Framework provides a mechanism to retain nuance and knowledge about the sources of conflict, the tensions, and the possible trade-offs, that have been surfaced in engagement phases of work.

This study also highlights the need for different approaches at different levels of granularity, and the importance of understanding the various perspectives on the challenges and concerns faced in each thread of the triple-helix. These may be broadly the same, but will be experienced differently and will require different responses or a different application.

Finally, this landscape case study reveals significant gaps in current methods of anticipation, engagement, and responsiveness, and highlights the need for a more programmatic approach that can reassure publics, create agile governance, and support the academics, policymakers and commercial organisations working to develop useful, safe quantum computing.

---

# CHAPTER 4: CASE STUDY – RI IN PRACTICE

---

*“So persuasive is the power of the institutions we have created that they shape not only our preferences but actually our sense of possibilities”.*

*Ivan Illich*

This chapter examines how an anticipatory governance approach such as responsible innovation can be adapted for different types of research and in different disciplines, and discusses what will influence the embedding of such governance into processes and procedures.

## 4.1: Introduction

In early 2018 EPSRC issued a call for Centres for Doctoral Training (CDTs), to fund new centres that would accept their first cohort in autumn 2019. The call ran across all its research themes and eventually a total of 92 CDTs were funded across an enormously diverse range of topics, from games intelligence to AI in healthcare. The purpose of funding Centres in this way is generally to address perceived capacity issues in particular sectors. This is demonstrated by the fact that 16 of the CDTs were in artificial-intelligence-related domains – a topic that is not only extremely broad but is seen as key for the UK’s claim to be a “science superpower”. Two quantum technology CDTs were funded.

Setting this call apart from previous funding rounds, however, was the inclusion of Responsible Innovation as a mandatory element of training for students. The exact phrasing varied between research priorities, but *all* CDTs were required to provide training to their students in RI. Section 3.5 of the overall call document read:

### ***“Responsible Innovation***

*All students must receive training related to the Responsible Innovation Framework. Responsible Innovation (RI) is a process that seeks to promote creativity and opportunities for science and innovation that are socially desirable and undertaken in the public interest. EPSRC introduced its framework and AREA approach for RI in 2013. Often described as Responsible Research and Innovation (RRI) which highlights the important role of research in the framework, here it is referred to as RI. This*

*is to reduce confusion with the use of the RRI term developed by the European Commission which has an emphasis on broader thematic elements.*

*We would expect students to receive training in the general topic of RI as well as in issues more specific to the scientific areas relevant to the Centre. The amount of training and consideration taken of RI should be a proportionate response to the Centre's vision and topic, the requirements outlined in a priority area description (if relevant), and individual student's projects. UKRI expects that all CDTs are able to demonstrate that resources have been committed to activities relevant to RI."*

The significance of this is clear – this was the first attempt to embed responsible innovation approaches across the training of a whole generation of students. In order to teach responsible innovation, training would likely be necessary for supervisory staff as well. Such an undertaking represented a significant shift in policy, and may lead to the inference that responsible innovation as an ethos and a practice is becoming institutionalised – at least within this particular funder. The purpose of this study, therefore, is to examine in detail this question of institutionalisation of an anticipatory governance technique and its implications, as well as to look at the different ways in which responsible innovation might be understood across this very wide range of domains.

#### 4.1.1: Institutionalising RI

As described in Chapter 2, responsible innovation has been promulgated throughout the last decade by EPSRC (2.2: *Responsible innovation and the AREA Framework*) – it has become incorporated to an increasing degree in research calls but prior to the 2018 call had not previously been mandated in CDTs. Also discussed in Chapter 2 (2.4: *Institutionalisation? Embedding RI in research*) was the point that this mandate does not necessarily imply that responsible innovation has been 'institutionalised' either at EPSRC or within the universities hosting the CDTs. Both universities and funders are a type of institution. Although 'institutions' can be conceptually framed in various ways, an institution is a system founded on rules (Genus & Iskandarova, 2018) – it is by definition a 'stable' entity possessing a high degree of inertia and is therefore generally resistant to change (Pansera et al., 2020). For Randles & Laasch, (2016) institutionalisation is an exercise in embedding new norms. This, importantly, requires old norms to be discarded, a process they term (*de*)*institutionalisation*, in order for new norms to take their place. Given the inbuilt resistance to change, this can be a long process.

Owen et al, (2021) have examined the progress of institutionalising responsible innovation within EPSRC, and conclude that it has "*focussed ... largely on advocacy, mobilising support and developing the skills and knowledge necessary for RI implementation*" – which while necessary, is deemed insufficient. Owen et al (2021) call for "*more assertive*" governance – this reflects many factors, not least that, despite the inclusion of RI in funding calls, RI is rarely (if ever) assessed as a critical part of a

funding proposal's acceptance or rejection. It is also not habitually examined as part of mid-term or end-term reviews of projects. This supports the previous work from Genus & Iskandarova, (2018), who describe the institutionalisation of RI in the EU and US as "partial", and stemming more from a wish to "secure desirable outcomes from scientific research and promising innovations" than from a drive to improve society's trust in scientific endeavours. As seen in section 2.5.1, Coenen & Grunwald, (2017) also looked at the integration of responsible innovation in the UK's quantum technology landscape, and concluded that the UK's approach was inadequate to the task.

However, this may not be a complete picture. Institutionalisation can occur in at least two ways: through the top-down imposition and promulgation of new methods, rules, and standards – seen in mandates, policies, strategic documents etc – but also in a 'bottom-up' mode, evidenced through alterations in grassroots behaviour, where people may shift their own patterns of activity and create new modes of operation according to their own understandings or requirements (Randles et al., 2016). EPSRC's adoption of RI may be seen as the 'top-down' variety, but institutionalisation of responsible innovation that follows the latter, bottom-up, pattern is often termed '*de facto*' responsible innovation (Lindner et al., 2016, following Rip, (2010) who discusses *de facto* governance). Broadly defined, such *de facto* RI covers activities that include engagement, reflection, anticipatory work, and responsiveness, without necessarily being formally framed as responsible innovation. This also aligns with Thelen's work on how institutions adapt to changing times without necessarily altering their foundational structures (Thelen, 2004) and can experience significant shifts in the way they do things without ever formalising these processes. Institutions, although *constructed* as rule-based entities, are *comprised* of people with a certain amount of agency and discretion who create structures and routines, and may choose to alter these through daily choices, communications and collective decisionmaking (Grimpe et al., 2020). Projects that have investigated these differing forms of institutionalisation include the RESPONSIBILITY project, which aimed to create a network of stakeholders that would adopt and diffuse a common understanding in Responsible Research and Innovation between different actors in Europe and around the globe, and the ORBIT project, which set up an Observatory for Responsible Innovation in ICT.

It is here, in the area of individual and communal action, but aside from formal rule-based structures, that theories of the *learning* organisation may be instructive.

#### 4.1.2: The learning organisation and responsible innovation

The topic of 'change' in organisations and institutions has been widely studied in recent decades, whether the subject is greater efficiency in managing projects (Tinoco, Sato, & Hasan, 2016), culture (Handy, 1976; Harrison, 1972; Koberg & Chusmir, 1987) knowledge management (Cross, 1998; Wilson,

2002) or other foci of interest. The challenge of creating change and entrenching new norms, such as we might expect to see occurring through institutionalisation, is therefore a vast field, particularly in the management literature. These processes of change have also been a focus of interest within responsible innovation research, as responsible innovation approaches need to be adopted in a widespread manner in order to embed new mindsets and ways of carrying out research. It is therefore of value to investigate these topics within the management literature in order to draw lessons for the ways in which responsible innovation approaches might be adopted – whether through top-down or bottom-up mechanisms.

Theories of the learning organisation are one of the most enduring management tools of recent decades, in which the seminal work is *The Fifth Discipline* (Senge, 1990), published more than 30 years ago. Senge tackled the management-heavy theories of organisational operation that prevailed at the time and argued that the task-based modes of working that were part of management theory at the time, may cause employees to be unable to see the ‘big picture’. Individuals tend to “lose [their] *intrinsic sense of connection to a larger whole*” (Senge 1990, p17). His solution to this was to empower employees and give them agency in their own fields, to encourage them to comprehend the ‘whole’. Senge lists five “disciplines”<sup>20</sup> or activities that, if practised well and thoroughly, can lead to the creation of a *learning organisation* – a forceful, collaborative enterprise – which could not only survive but thrive in commercial ecosystems.

Since the publication of Senge’s work over thirty years ago, his theory of the learning organisation has been regularly revisited and critiqued, although the height of its popularity in management theory was around the turn of the millennium. The ‘learning’ element of the theory in particular has been the focus of much scholarly discussion (see, for example, Easterby-Smith et al., 2000) and this may be linked to the vagueness of the definition and exactly what the ‘learning’ should encompass. The loose definition of a ‘learning organisation’ may also be a contributor to its ongoing usefulness, however, as new generations adapt it to novel circumstances and changed commercial landscapes – Hoe, (2019) suggests that although theories in management tend to be transient, the elements of the learning organisation concept have become accepted as simply part of good management practice. This offers insight into the processes of institutionalisation discussed above, where new norms become embedded and accepted as part of the organisational ‘culture’ (Handy, 1976).

Hsu & Lamb, (2020) address this point when reviewing some of the critiques of learning organisation theory, including the question of whether or not Senge’s theory can improve a company’s bottom line.

---

<sup>20</sup> Senge used the term ‘discipline’, but this can be confused with academic disciplines, therefore in this chapter the term ‘activity’ will be used.

Rather, they propose that – instead of providing a blueprint for economic gain – learning organisation theory can provide different benefits. These might include the removal of layers of middle-management through creating greater agency at levels of the organisation: *“learning organisations realize that empowered workers can make better decisions than managers”* (Marquardt & Reynolds, (1993) p56, in Hsu & Lamb, 2020). This insight is useful because it has parallels with the concept of *de facto* RI discussed previously, where individuals and small groups are empowered to make decisions, and create processes and structures within an unchanging overall framework that, over time, result in an institution with a potentially significantly altered position. The lack of evidence on whether learning organizations perform better economically, versus acceptance of these ideas as simply good practice, demonstrates that organisations can evolve without an economic incentive but simply because the new norm represents an improvement. It is possible that this may also apply to the rollout of responsible innovation within CDTs.

A theoretical link between learning organisation principles, and processes of embedding responsible innovation, is also discussed by Hansen, Jensen, & Nguyen, (2020), who have interrogated whether theories of the learning organization are able potentially to facilitate responsible innovation approaches. As with Hsu & Lamb (2020), discussed above, Hansen et al (2020) recognise the point that learning organisation theory does not necessarily provide companies with economic advantage, but suggest that the additional benefits of adopting these practices are sufficient to ensure that they remain useful. (It is also salient to note that non-economic advantages frequently translate to economic ones in the longer term – eg increased consumer trust that can result in greater brand loyalty). Such a point may also apply to the adoption of RI methods. Hansen et al (2020) additionally discuss the ethical basis of learning organisation theory, and suggest that it bears strong similarities to the underlying philosophies of responsible innovation. They note, for example, that there is no equivalent in learning organisation theory for the inclusion of stakeholders in consultation, but argue that there is sufficient common ground that it is only *“a small step from being a learning organisation to becoming a responsibly innovating learning organisation”* (Hansen et al 2020, p65). It could also be argued that learning organisation theory’s empowerment of sub-management employees is a *de facto* inclusivity that is well-aligned with responsible innovation values. Although the authors do not discuss it, their argument parallels the points about *de facto* RI seen in Hsu & Lamb (2020).

It can therefore be seen from the discussion of institutionalisation in section 4.1.1 above that there are significant areas of overlap between, in particular, *de facto* responsible innovation and learning organisation theory. These insights are examined in this case study on the institutionalisation of responsible innovation - whether *de facto* or explicit – being carried out in UK CDTs.

### 4.1.3: Research questions

The purpose of this study was therefore to evaluate and investigate understanding and depth of knowledge of responsible innovation theory and methods in the CDTs funded by EPSRC in the 2018 funding round (outlined in section 4.1 above). Opinions and perceptions were sought about the various ways in which RI is professed, understood, and actioned by the CDTs. This included an effort to appraise both the extent and the penetration of understandings of responsible innovation, and gain a wider comprehension of how RI was becoming embedded. The study investigates the research question:

*How are different fields and disciplines addressing the challenge of incorporating responsible innovation into their work, and what can this tell us about how new norms such as responsible innovation become institutionalised?*

As a supplementary research question, this study also aimed to answer the question:

*What can the institutionalisation (or not) of responsible innovation tell us about the use of anticipatory governance methodologies in novel technologies?*

## 4.2: Research method

As the research questions required significant focus on the individual views and research domains of the various Centres, the research method utilised here was to carry out a set of semi-structured qualitative interviews with CDT directors.

Once ethical approval had been received, all the Directors of CDTs funded in the 2018 call were invited to participate. This gave a total of 91 CDTs being invited, including 16 on artificial intelligence topics – 28 were eventually carried out. (35 Directors had responded initially, but in seven cases interviews could not be arranged due to the exigencies of coronavirus lockdowns.) The CDTs that did participate represented research fields across the entire EPSRC remit, giving a broad spread of types of novel research and technology. It was therefore considered that these interviews were a fair representation of a number of expanding research fields.

As with the previous study, a wide-ranging, semi-structured interview format was used, to draw out the Directors' experience of incorporating RI into their Centre. Questions focused on: their own prior and current understanding of RI; the depth of understanding within their Centre (eg awareness of RI amongst supervisory staff); RI training activities they had planned; how and whether they perceived RI-type activity within their associated industry partners; RI assessment, particularly in student projects; and their experience of integrating a new theme into their teaching and training. These

questions enabled the surfacing of themes around levels of understanding and interpretation of RI, institutionalisation and institutional change, amongst others. Using the same method as the first case study also ensured that where similar topics were discussed or addressed, the responses of different respondents could potentially be compared and contrasted.

As with the previous study, each audio-recorded interview was transcribed in full, redacted at the point of transcription and assigned a code. Interviewees are referred to by this code – eg D23.

### 4.2.3: Analytical processes

The review of literature outlined above, incorporating work on themes around change in institutions and embedding new norms, had provided some indication as to the findings that might emerge from the data, but nonetheless an inductive approach was taken to the coding. This approach was adopted in order to ensure that preconceptions were not brought to the interviews, per Braun et al., (2019).

Interviews were open-coded at the sentence level, then the codes reviewed once all the interviews had been assessed once. The transcripts were then coded a second time to iterate on any additional themes, and to confirm that earlier-coded interviews were checked against themes that had emerged later on in the process.

As an additional layer of analysis, a data session was conducted within the wider research group in order to draw out any further topics for discussion.

## 4.3: Findings

The study demonstrated overall that the assessment from Genus & Iskandarova (2018) of the institutionalisation of responsible innovation as ‘partial’ is likely justified. Even after the processes of applying for CDT funding and therefore having to include discussion of responsible innovation training programmes within their proposal, participants reported a low level of familiarity – for example D02 who was only confident that they possessed “*a basic understanding*”. For many, the process of writing the proposal for the EPSRC funding was the first time they had encountered responsible innovation. However, differing levels of familiarity were reported dependent upon discipline and experience.

*“I’m very close to ethics myself so I was very keen to have this kind of component in the new CDT” D05*

There was also a view that some participants would be naturally more familiar with RI, as their research had more direct societal impact.

*“I think some people probably do have a much deeper handle on it ... depending on the implications of their research” D72*

The question of comprehension is key, because the requirement to train students may not be fulfilled without such an understanding.

### 4.3.1: Understanding responsible innovation

#### Disciplinary/industry influence

The way responsible innovation was understood and deployed as part of the training in the Centres was frequently closely tied not only to the academic discipline of the CDT, but more broadly, to the industrial context for the CDT's area. This is well-illustrated by the Centres that have a healthcare focus, and may be partnered with local teaching hospitals or medical-device manufacturers. In such areas, the context for the CDT's work is not only a highly regulated one, but also brings with it a deep understanding of the need for ethical approaches that can supplement these regulations. The Centre itself will therefore conform to both the regulatory context and the associated ethical standpoint of its industry partners.

*“because we're in a very applied field ... our researchers in this field would think they're quite strong in RRI because they're thinking about addressing healthcare needs” D73*

*“I think that there are aspects of it that are ... that are obviously the things ... that seems to me that you must obviously consider, in terms of how you do things” D80*

There was a clear division here from areas of industry and commerce that have neither such a long tradition of ethical self-examination, nor a well-established governance framework. Such areas – for example in applications of artificial intelligence – offer a far less rich environment to which the CDTs are able to relate their RI work. This lack of contextual information, or abundance of disciplinary discourse, can mean that supervisors and students have far fewer resources to draw on to assist them in understanding how to tailor responsible innovation to their own area's particular requirements.

*“[our] scientists, perhaps, are not exposed to that ... type of thinking” D27*

Participants were very aware of these distinctions, understanding that, in looking laterally across the EPSRC remit, there were inevitably going to be significant differences in the sophistication and depth of RI understanding.

*“there's a [disciplinary] perspective on it that I will have and ... obviously in different subject areas that's going to be different” D80*

There will inevitably be knock-on effects here in terms of the way RI is or can be institutionalised, as the varying industrial contexts influence how RI is conceptualised within the CDT.

## Ethical underpinnings

This may also extend to the design of RI training – many CDTs incorporate RI as part of their ethics training. It was relatively rare, particularly within the CDTs addressing these questions from very foundational positions, to find RI framed as a research methodology, a practice, or a set of design guidelines. Responsible innovation approaches are indeed grounded in normative ethics (eg Stahl et al., 2016), but RI practice does not overlap significantly with philosophical discussions of ethics, or training in applied ethics such as that seen in medical practice. Several participants, despite having familiarised themselves with responsible innovation during the processes of preparing proposals for the funding call, were uncertain of where boundaries lay, and how responsible innovation might be differentiated from other frameworks.

*“there are rules, and there are things you have to do. But just abiding by the rules is not sufficient ... and this is more I suppose ethics” D71*

*“I’ll be honest sometimes I think I get RRI and ethics a little bit confused or I’m not sure ... if I’m talking about ethics or RRI” D72*

Responsible innovation work, for example, incorporates the involvement of stakeholders “very upstream” (RRI Tools, 2013) – springing from this requirement will be practical questions such as ‘who is a stakeholder?’, ‘when to invite them?’, ‘what to discuss with them?’. Many of these questions will in turn be influenced by the political questions discussed in 2.3.2: *The politics of inclusion*. Although ethical considerations may guide interaction with stakeholders, these processes of participation and inclusion cannot really be elided with ethics training. The responses from participants in this study underline how differently responsible innovation can be interpreted, depending upon the disciplinary context. Other approaches were demonstrated by those who understood it as a practice, and those who interpreted it as a set of criteria.

*“it needs to be part of how you exist in the future” D71*

*“the **how** you do it is much, much easier to teach and to just tell people to adhere to certain rules” D71*

## Wider ethics-based and other approaches

Other interpretations were also seen, including those who connected these requirements for inclusivity with broader Equality, Diversity and Inclusion (EDI) practice. It is certainly the case that gender equality is regarded as one of the EU’s RI ‘pillars’, but EDI work should be differentiated from the AREA Framework. EPSRC separate these concepts on their website and in calls, making it clear that they should not be conflated as far as the funder is concerned.

Within some CDTs, especially those with connections to manufacturers as industry partners, it was common to find responsible innovation strongly identified with concerns around environmental

impacts. Sustainability is also one of the ‘pillars’ of RI as seen within EU frameworks. There have indeed been calls in the responsible innovation discourse to adapt the remit of RI to more explicitly incorporate sustainability concerns, or examining its utility in sustainability work (eg Stahl et al., 2019), and to CDT Directors it seemed to be a natural extension.

*“they would look at ... the whole process, and the implications in terms of the environment and the societal impact” D80*

Given that, by comparison with ED&I, sustainability is not yet included in funding calls as a general rule, this seems an important and useful addition to the RI toolbox – however, Directors sometimes struggled with the boundaries of these issues; *“is this ethics? is this sustainability? is this RRI?”* (D72).

### ‘Hard’ v ‘soft’ skills

The question of boundaries between areas of learning also arose elsewhere in training programmes. In particular some Directors were unsure of whether they should categorise responsible innovation as a ‘soft skill’ akin to eg leadership, or a ‘hard’ skill such as technical knowledge, or training. In reality there is not a clear delineation between skills in this way (writing skills, for example, would fall between these categories, as it obeys rules but also has stylistic elements) and in RI practice, considerations should spring naturally out of the understanding of the work. For example coding may be classified as a ‘hard’ skill, but there is no particular rationale for classifying the consideration of issues arising from that coding as a ‘soft’ skill. It is important to note this categorisation of skills, however, as classifications such as these may have a bearing on how seriously (or otherwise) training in these skills is taken – especially in disciplines that traditionally focus on ‘hard’ skills.

There was also recognition from participants that in many respects, responsible innovation requires students to adopt new perspectives on their work in the type of ‘systems thinking’ approach described above (Section 4.1.2). Being encouraged to think holistically from the very beginning, for example when considering stakeholders, is likely to be an entirely novel concept to many students. These processes, and the externalities that such a process can draw attention to, were often unfamiliar territory and required new approaches and novel skills.

*“it’s just getting the students to think ... right at the start in terms of their research and all the different impacts it can have, and what they need to consider. Because ... they’re not trained to really step back and think about the long term and also the indirect impacts that our research is having.” D73*

Finally, there were some participants who understood that a formal descriptor of an activity as ‘responsible innovation’ was not necessary for responsible innovation practices to be adopted, in an understanding of the *de facto* RI discussed in previous chapters.

*“we do try and practice the principles of RRI but but not formally recognising it as RRI” D03*

This recognition may be particularly significant in terms of wider processes of institutionalisation.

### 4.3.2: Institutionalisation processes, and embedding responsible innovation

As discussed above, widely varying levels of familiarity with responsible innovation approaches were seen across the CDTs, and often correlated with pre-existing disciplinary association around ethics, or around industrial or commercial regulatory context. Personal interest also played a part, with some Directors acknowledging the impact on their own practice of having to incorporate it into their plans for their Centre.

*“I’ve thought about it much more carefully as a result of putting in the CDT bid” D71*

Some Directors and supervisors were not only familiar with the general precepts but had published their own work in RI or RI-adjacent topics such as Technology Assessment. Others were coming to RI for the first time and had not yet had the opportunity become familiar with it (*“it’s not reached the top of my to-do list” D68*). Others were finding it challenging to relate to their own discipline. Any examination of institutionalisation, or organisational change, therefore has to take into account this variety of familiarity. These differences are also reflected in the number of different approaches to training structure.

#### Depth of training provision

The mandate from EPSRC did not provide specifics on how training in RI was to be offered – in some CDTs, training in responsible innovation was essentially being outsourced as a discrete piece of work from an external provider, or central training resources at their institutions. This was especially the case at universities with multiple CDTs, where training several sets of students at once might be seen as most efficient.

*“It seems like a ... great way to do this type of training is not ... for everyone to reinvent it in every single CDT, but to share some of this effort across the different CDTs” D21*

This type of approach, taking advantage of network effects to train large groups of students together, has obvious attractions, but may limit the degree to which students are able to work on tailoring RI approaches to their own particular field or project, which should form a key element of the training. Additionally, bringing external training providers into the institution would likely mean limiting the extent to which responsible innovation approaches were brought into the institution more widely.

In other CDTs, almost the opposite approach was taken, with Directors booking supervisory staff onto training courses in order to ensure that responsible innovation approaches were understood at all levels of the Centre.

*“the new thing that we’re doing isn’t just about training up the students, it’s about training up the team” D73*

*“the supervisors have to ... buy in.” D07*

In many CDTs, however, there had not been systematic attempts to establish the extent or depth of existing supervisory understanding of RI, with some Directors justifiably pointing out that there are significant training requirements for supervisors and academic staff already. Others were clearly aware that supervisory familiarity with RI might be considered to be essential in terms of assessing how well students are adopting an RI approach within their own projects.

*“We ... thought about from the very beginning ... there will be supervisors who aren’t aware of responsible research and innovation and how we’re going to bring them up to speed” D02*

Wider supervisory training may also be key in terms of the dissemination of RI approaches – supervisors within CDTs will be conducting other teaching and research as well as working within the CDT, and accordingly the level and type of training they receive may have a significant impact on the institutionalisation of RI.

### Wider adoption of RI

These differences in the CDTs were also reflected in differences in the ability to ‘embed’ RI into wider research groups – which many Directors stated as their aim – as this was often dependent on the attitudes of colleagues, who might or might not be receptive to responsible innovation considerations.

*“[the Director] is very committed to RRI ... other faculty are also very committed ...and then there’s a few people that are holdouts. They go along with it. ...They’re not a large number, but ... they do the absolute minimum.” D07*

The pattern seen across the CDTs, where there were widely varying understandings of the concepts of anticipatory governances and responsible innovation, was also reflected within the broader institutional context. Some Centres were working within or across teams where there was significant understanding of RI already – in some CDTs this even included supervisors and other academics acting as RI ‘champions’. Other institutions took the view that RI could be dealt with in a relatively brief and informal way, as essentially a piece of administrative work to be managed in a similar manner to applications for ethical approval for a research project. These differences could be exacerbated in Centres that were collaborations between several institutions.

These ongoing processes of institutionalisation also demonstrate the importance, discussed in Watkins & Marsick, (2019), of individuals and their decisions. One Director reported that their CDT had a “*very rich archive of case studies on RRI*” (D73) because at one point they had had an RI specialist as part of the team who had been extremely committed to developing such an archive, and this was still seen as

a valuable resource. Other Directors were confident that, as heads of department or teams, they were in a strong position to ensure that responsible innovation was translated widely across teams and into projects. In such cases, where an individual Director is enthusiastic about the value and importance of responsible innovation, a mandate such as that from EPSRC provides critical support for their initiatives. Some Directors were already looking across their institution and considering ways in which they could broaden the remit and impact of their RI work.

*“I’m confident that everybody in our CDT will ... get some exposure to RRI and ... some training and skill in that area and then it’s ... ‘how can we maximise the impact of that?’”*  
D20

Others were more doubtful about how far RI approaches were meeting with a positive reception, and becoming part of organisational culture.

*“I would hope it’s much more seeded in people’s minds, but I don’t know the answer”* D23

There was frequently an understanding that the incorporation of RI was not necessarily limited to the CDT and there could and possibly should be wider institutional adoption. Many Directors understood that for RI to truly take root in an institution, it needed top-down, strategic-level support, with buy-in at senior levels, as well as the relatively low-impact (from an institutional perspective) bottom-up training work being carried out in the CDTs.

*“that overarching level, I think is - it’s almost a driver for quite a lot of what we’re doing anyway, so that that level of thinking comes into it from from the start.”* D80

There was a recognition that without the RI mandate in the CDT call, RI would have been very unlikely to feature in so many training programmes, demonstrating the importance of top-down funder decisions – including mandates where necessary – in processes of institutionalisation.

*“if funding agencies don’t kind of ... prod us into this direction I think people would generally not care too much about it”* D21

### ‘De’ institutionalisation

The processes of scoping, designing, and embedding a new CDT within a University carry a heavy bureaucratic and administrative burden, and in many cases Directors had found themselves battling the “incumbent logics” of their own institution (as discussed in Section 2.2). Numerous challenges were described by participants, including the difficulty of getting funding agreed, agreements signed, management bureaucracy for new courses handled – and in many cases a lack of familiarity with new criteria and difficult choices to be made about incorporating new training either alongside or instead of existing programmes. These are even more the case when there is a new element, such as the requirement for responsible innovation training.

*“there’s an entire programme bureaucracy that makes it very hard for you to ... suddenly make changes” D22*

These processes could arguably be characterised as the ‘de’institutionalisation discussed by Randles & Laasch, (2016) – setting up a new CDT requires adaptation on the part of the institution. Adaptability is generally not a feature of large, rule-based entities, and new norms are not easily adopted. D65 demonstrates an understanding that incorporating a new approach in this way requires a change to established modes of thinking.

*“it has been something of a kind of culture change at the university in terms of embedding some of the thought processes that sit around this” D65*

This Director also showed a clear understanding of the *impact* of changing an embedded approach in this way – once responsible innovation approaches have been institutionalised, mindsets change so that:

*“rather than thinking of it as something distinctive and different it's just part of what you normally do. And ... I think that comes across in your research applications thereafter because it's not something you ... have to think about, it's just something that automatically comes out in the applications that you're putting forward.” D65*

The recognition here is of RI as a way of thinking, as well as a methodology, and a mindset that – as it becomes more accepted and taken-for-granted – will also play a useful role in future funding applications.

*“Those who are not as familiar with it are sort of becoming more so, as it's becoming more and more sort of centrally important” D49*

As D49 shows here, and as discussed in Section 2.2, RI is becoming more common in research calls, where Owen et al., (2021) have called for more “assertive” governance and assessment.

### 4.3.3: The importance of assessment

The issue of ‘assessment’ of RI activity, as noted above in Section 2.2, will be critical to any question of whether a process is becoming institutionalised. This understanding is of course not limited to this area of research - there is a decades-long volume of management literature, for example, from Kaplan & Norton, (1992), to van der Kolk, (2022) devoted to the importance of measurement in improving performance. As institutions subject to the Research Excellence Framework (REF) well know, there is an understanding that measurement or assessment can support activity regarded as important but that without the evaluation might potentially be practised less assiduously. However, there is also voluminous research about the way in which measurement metrics themselves can alter performance in sometimes-undesirable ways, for example ‘teaching to the test’ (Volante, 2004). Assessment retains its value in terms of demonstrating the importance of a given factor, however, or how seriously it is

taken, and is therefore something of a useful proxy for institutionalisation. In terms of institutionalising responsible innovation, such assessment would be important at student level, at supervisory or leadership level, at the level of the institution, and within the funder – without evaluation at all these different levels of granularity, RI approaches could hardly be said to be truly ‘embedded’.

## Students

Within the CDTs there were indeed some that planned to assess student inclusion of RI.

*“students are expected to ... clearly articulate how and where they have considered RRI within their research and how that has influenced their project progression” D03*

The way in which this can affect the progress of their research varies widely between CDTs, however, with some supervisors needing to understand and ‘pass’ a student’s RI work.

*“each student’s main supervisor will have to look at that RRI information, and then give it sign off” D49*

At the same time in some other Centres, students were not expected or required to reflect on their RI training at all. The importance of individual decision-making is seen here, as Directors decide for themselves what is relevant for their own CDT, without necessarily referring to the funder’s original requirements. The degree of embedding of RI concerns – in some cases very little, and in some Centres throughout the course of the work (for example, *“we expect them to be able to talk about responsible innovation as part of the viva” D68*), is therefore highly inconsistent.

## Supervisors

A few CDTs were also making efforts to ensure that supervisory staff understanding of RI was not just supported by training but was also underpinned by assessment. Others were aware that supervisory training might be a supporting factor they could usefully consider.

*“if we’re preaching to the students that this is important and they should do it, it really should be more of a “we’ve done this as well’ and there is value in the top team knowing about it,” D12*

*“we don’t actually get the academics ... to reflect on that and that actually might be something that would be good for us to explore.” D41*

It would seem likely that in order to be able to assess student understandings of RI, supervisors will need an understanding themselves. In order for responsible innovation approaches to be taken seriously, individual supervisory attitudes will also be a key factor.

*“At the end of the day, it’s the supervisors that say either ‘that’s important’ or ‘don’t bother about that’. We want the supervisors to acknowledge it’s important, and take it seriously” D07*

*“All the supervisors’ proposals to get funding for PhD students underwent a review, which included the review of RRI. [Supervisors] all had to write something [on RI]. And they all got feedback. And not all of them got passing grades.” D07*

The above comment from D07 demonstrates the variety of approach – this was the only participating CDT in which supervisors might ‘fail’ their RI assessment.

## Funders

At the highest level of analysis, in the question of the funder’s mid-term assessment, all the interviewees were anticipating that the review from the funding council would contain an element of reporting on the CDT’s RI activities and training. This understanding demonstrates a recognition of the importance of measuring and reporting for institutionalising new behaviours and norms.

### 4.3.4: Student approaches to responsible innovation

Many Directors whose students had already undertaken some responsible innovation training, reported that it had received positive reactions. These Directors often discussed how students had seized the opportunity to consider wider issues – these might include, in particular, environmental concerns. They also often welcomed the chance to place their work in context vis-à-vis societal impact, and to discuss with peers the necessity of given trade-offs or compromises.

*“when we introduced our students to [RI] they were very interested in that and ... engaged with it.” D20*

There was a sense in which RI was understood as a way to ‘do the right thing’, and an almost instinctive linking of it to ethical approaches, as well as a long-term perspective that recognises the ‘engaging with the future’ element of responsible innovation (section 2.2).

*“they very quickly get why RRI would be important to them” D27*

*“they show ... awareness of how important it is. They have a very strong sense of responsibility” D62*

A strong theme was the provision of space in the training curriculum to discuss and debate these questions, and some Directors had created peer-to-peer spaces for cohorts to talk over these sometimes-challenging issues.

*“one of the things ... with the cohort is that they learn things from each other. And if you have people coming in who can offer experiences that are different, that’s something I think very much that some students can learn from others” D80*

However these positive attitudes towards RI were not entirely universal – some Directors reported seeing different approaches in different fields, with small minorities of students exhibiting impatience with questions that they saw as being irrelevant or constraining.

*“students that I meet in other programmes – responsible innovation seems more like a barrier or an inconvenience that stops them from doing really interesting ... things. ... Whereas it’s seen obviously in places like social sciences as much, much more central to the research agenda.” D42*

Further investigation might reveal a variety of motivations for this, including a recognition that students may not fully appreciate the broader context into which their work would fit. Potentially, therefore, greater understanding would generate greater enthusiasm.

*“I think students in particular, I think it takes a while to appreciate that” D80*

Supervisors, with their relatively greater general awareness of the industrial context into which their students might be moving, tended to think that RI training could be advantageous. This was clearly most relevant where that was likely to include a move into industry – being sensitised to questions of impact, environmental concerns, and societal acceptance, would provide students with a more mature awareness of the industrial context.

*“what a lot of [our industry partners] are particularly enthusiastic about is ... our aim to develop students, researchers that are ... well-rounded and ... have this range of skills” D15*

#### 4.3.5: Broader familiarity – RI practice in industry

The question of the level and type of industry involvement varied between CDTs, with some Centres involving their industry partners on a relatively regular basis, whether to deliver training or occasionally as supervisors, and other Centres inviting their partners in on an infrequent basis for industry days etc. In some Centres, Directors were considering how to open or expand discussions about responsible innovation with industry partners – some regarded the mandate in the call as a good jumping-off point for this. As discussed above, prior awareness of RI-type concerns within industry partners was often dependent upon the industry sector associated with the disciplinary area. Healthcare-focused CDTs, for example, tended to have industry partners with high levels of awareness of ethical challenges and the need for new practitioners to be well-versed in responsible approaches. For some Directors, the provision of RI training was thus an absolute necessity.

*“it’s kind of our selling point that we do Responsible Research and Innovation” D05*

Depending on the industrial context and the particular arrangements of the CDT, there might also be a need to discuss RI in more depth with the partner – this was particularly prevalent in Centres where industry partners were more directly involved in student training. Industry partners needed to understand what the students were being trained in and how it would affect their outlook.

*“[industry partners] are listed as co-supervisors so they’ll have to have the training” D73*

*“as research scientists, certainly that's something you talk about with companies ... there's some aspects of it, they're very keen to get out there and to establish that that's something that they're thinking about.” D80*

Industry partners were frequently reported as having very pragmatic views, understanding that new generations of employees might potentially expect ethical approaches to have been adopted by their employers. These include their approaches to questions such as inclusivity and diversity, sustainability and the environment, and so on. Indeed, in some sectors a consciousness of an overtly ethical position was a feature of the overall industrial context, with Directors being aware that their particular industrial environment would require students to be well grounded in ethics-based approaches.

*“sustainability, and evidence of sustainability, for employee satisfaction, retention and recruitment, is important” D07*

Several participants reported that their industry partners were aware of the need to take RI-type concerns into account, although they were unlikely to term it ‘responsible research and innovation’, and Directors who had previously worked in industry agreed that the terminology was not something they had encountered in those environments. The concepts of social responsibility and ethical alignment, however, were familiar from a terminological point of view and also as influencing practices that might be undertaken within industry. As well as these pragmatic approaches, some Directors noted that although industry partners may not use the term ‘RI’, they nonetheless engage in activities that reflect RI-type concerns and activities. This reflects a *de facto* RI approach as discussed in Section 2.2.

*“ they may not recognise what they're doing as RRI but actually some of the things will be addressing some of the issues of responsible research” D03*

In terms of actual activities undertaken, as opposed to use of a particular terminology, some Directors were aware that their industry partners were in fact more cognisant of some of these issues – as being further downstream – than they themselves had been. These activities might be driven by a number of factors, including the employee satisfaction mentioned above; the desire to have a good ‘story’ to tell customers (and/or a recognition that such issues are important to customers); regulatory compliance, and other motivators – the net effect was that in many cases the commercial environment was potentially very receptive to RI concerns.

*“just from dealing with industry partners, they're probably more tuned into this than universities” D23*

*“I could be cynical and say, it's part of their PR, which, yeah, I'm sure it is, but I'm also sure there are aspects behind it where because of changing legislation they've actually been forced to think about things in a little bit more of a socially responsible way.” D80*

There was also a recognition that, from a practical point of view, larger companies may simply have more resource with which to tackle questions of responsible innovation. A small startup, even with the best of intentions, might not have the capacity to spend a lot of time reflecting on RI.

*“I would imagine [large partner] has fairly advanced policies that may relate to RRI... on the other hand ... we work with a lot of SMEs ... who may be worrying about whether they will be here next year” D20*

However, as noted previously, the range of disciplines in the CDTs was reflected in the range of industry partners working with them, and it was not surprising that some Directors encountered resistance to RI approaches, particularly in areas where regulatory compliance or pre-existing ethical frameworks are absent.

*“actually a couple of them were quite antagonistic when we first raised it ... and we had to convince them that it wasn't about stifling innovation and it wasn't about political correctness” D68*

## 4.4: Discussion

The EPSRC requirement for responsible innovation training in the 2018 call for new CDTs was, as far as can be established, the first and the largest attempt to incorporate and embed RI practices within a higher education training context. The EPSRC remit is wide, and accordingly the spread of themes in the new CDTs was also very broad. In terms of addressing the research question for this study, framed at the outset; *“How are different fields and disciplines addressing the challenge of incorporating responsible innovation into their work, and what can this tell us about how new norms such as responsible innovation become institutionalised?”* it was clear that this breadth of remit is resulting in a correspondingly wide variety of responses, and numerous interpretations of concepts such as engagement, and responsibility. These responses are governed not only by degree of expertise, but by industry context. Some Centres had instigated highly detailed programmes for students, which were deeply embedded into the syllabus and training they were engaging in, whereas in others, where the experience of social-science based methods of critique was relatively novel, RI training programmes were frequently less in-depth. This illustrates the discussion above (4.1.1: *Institutionalising RI*), firstly around the challenge of displacing *“incumbent logics”* (Owen, von Schomberg, et al., 2021) in order to embed new ways of thinking into institutions; and secondly how limited the impact of an individual may be within the institution (Watkins & Marsick, 2019). Individual CDT Directors may have a good understanding of RI and its concerns, but if that is not reflected at an organisational level – if the organisation has not intentionally embedded RI into its culture and practice – then the institution remains unchanged and new understandings can remain superficial at best.

The problem of superficiality in the centre or the university can, of course, be partly addressed by widening the scope of the training, including supervisory staff to create greater strength in depth – and some Directors were indeed planning for this. However, new training will necessarily either increase staff workload or entail the removal of other elements – these choices and trade-offs reflect in a conscious way the process of embedding new norms (or not) referred to by Randles & Laasch (2016). These choices are likely to have a substantial effect on institutionalisation – as pointed out in Ten Holter et al., (2022), the question of whether supervisors are trained is more significant from an institutionalisation perspective than student training. Even though all students will receive RI training, there may be a relatively low impact on the organisation as the majority of students are not a long-term part of the institution and will largely move on without affecting it. They may, of course, take their insights and training into industry.

#### 4.4.1: RI in industry and commerce

The concerns, noted by D20 (section 4.3.5), around larger companies having time and resources to direct towards RI efforts in a way that smaller SMEs may not, underlines the real need for an institutionalised RI approach – RI training and practice cannot be the preserve of specialists, but needs to become part of the mindset for all, such that, as D65 noted, *“it becomes part of what you ... do”*. It is in this point around ‘mindsets’ that the second research question, *“What can the institutionalisation (or not) of responsible innovation tell us about the use of anticipatory governance methodologies in novel technologies?”* can be addressed.

There are several implications from the finding that the prevailing regulatory and ethical practice within an industry sector can feed back into the approach of associated disciplinary CDTs. It suggests the possibility that regulated industrial ecosystems have an influence on ‘partner’ sectors even if those sectors are not themselves regulated in the same way. Student projects in a healthcare CDT, for example, would not generally be subject to the same regulatory frameworks or ethical compliance requirements as a teaching hospital, but nevertheless, many of the same rules may be adopted both because of the training and experience of supervisory staff but also the need to align with the industrial context and prepare students for their future employment.

Another implication of this ‘linking’ effect between academic discipline and industrial use-cases might lie further downstream. As has been seen, in some sectors (such as healthcare) ethical and highly-regulated approaches are prioritised and this influences the work in the CDT. If this effect were also to flow in the opposite direction, this could mean that an ethical framework or anticipatory governance framework that was strongly embedded at the research phase may be taken on into further study or commercialisation, and have positive knock-on effects in these contexts. This could mean that EPSRC’s

decision to move towards institutionalisation of RI could feed outwards into industry from research organisations, influencing industry approaches in the long term.

#### 4.4.2: Institutionalisation of RI at funder level

Even though the funding body was not directly studied here, some inferences can be drawn about institutionalisation of RI at the funder level, providing an interesting coda to Owen, Pansera et al., (2021). It could be argued that mandating RI in the CDT call not only demonstrates a commitment to it in the longer term, but suggests a degree of *de facto* RI on the part of the funder itself. EPSRC has consulted with stakeholders, reflected on possible outcomes, and shaped its response accordingly, by embedding RI into the call.

In another footnote to Owen, Pansera et al (2021), with regard to the appeal in that paper for RI assessment, it was also noteworthy that Directors were anticipating that an assessment of their RI activity would be included in the mid-term review, even though at the time of the interviews EPSRC had not made any stipulation about the mid-term reviews. As seen in the discussion of the importance of assessment above (section 4.3.3), it is unlikely that, without evaluative processes, RI will become as embedded as arguably it needs to be in order to be effective.

Returning to the ‘learning organisation’ lens can also be instructive here. Marquardt & Reynolds, (1994) describe a process of ‘distribution of empowerment’. This is a method whereby empowered individuals are enabled to empower others, and can ‘cascade’ a process through their networks of influence. It seems possible to perceive a chain of such empowerment, incorporating reporting and assessment, on responsible innovation that runs down from the funding body through the CDT’s management structure, and supervisory staff, and from there to student projects. The key decisions of individuals, the importance of which was discussed in section 4.4.3, can also therefore be perceived here.

#### 4.4.3: Importance of individuals

Following the analysis of Watkins & Marsick, (2019), who stress the importance of identifying and examining change and learning in organisations at different levels of granularity, the findings of this study demonstrate the link between the particular individual in post and the likely spread of RI within their institution. For example, it might be an entirely fortuitous circumstance that a CDT Director has a prior familiarity with RI, but this familiarity may make a substantial difference to the way RI is conceptualised and taught. It also makes clear, however, that an individual’s knowledge cannot be regarded as part of the overall knowledge-base of that particular organisation – if that person leaves, their knowledge leaves with them. Such an individual may also be in charge of key decisions about

how far training is extended through the supervisory layer, or how far it is raised as a topic at management level. This was clearly demonstrated in the interviews, where some Directors were training supervisors, while others had not considered this to be a necessary step (or may be reluctant to add to existing supervisory workloads).

This lack of strength in depth when it comes to RI – the shallowness of the impact so far made by RI in terms of its institutionalisation within universities – supports the view of Owen, von Schomberg, et al, (2021) that AREA-based RI remains “*patchy*” and that there is still no “*self-sustaining culture*” of RI (Owen, von Schomberg, et al., 2021, p11).

## 4.5: Conclusions

This study has demonstrated the slow and variable progress of the institutionalisation of RI, suggesting that even an anticipatory governance technique that has been mandated, with the persuasive weight of a funding body behind it, may fall prey to both the “incumbent logics” of the parent institution and also of the related industrial ecosystem. The implication is that the institutionalisation of RI in universities may be necessary but will likely prove to be insufficient to make wider change. However, it may be the case that an adequate ‘tipping point’ has not yet been reached within certain fields, which might suggest the value of amassing a body of evidence for researchers in such fields, through being able to analyse another project’s RI approach. It could be the case that, as RI becomes more frequently mandated and – to whatever degree – institutionalised, that record-keeping and publication of RI-activity datasets would be a valuable resource for building knowledge and wider expertise in its application within particular fields or technologies. In particular, this study demonstrates not only the general challenge of the concept of responsible innovation, but the value of interpreting it in ways that reflect the individual challenges of a discipline or industry – an understanding that may also be seen in the later case studies.

The study also found that work from the discourse on ‘learning organisations’ can provide a helpful lens for examining how responsible innovation is progressing in terms of its institutionalisation, and potentially ways in which its impact could be extended. Additionally, using the ‘granularity’ lens from learning organisation literature to examine different units of analysis – the student, the Director, supervisors and the wider CDT – it is possible to evaluate more closely the different challenges and opportunities at each grade. For example, comparing the student response to their RI training, with the supervisory response to needing to become familiar with a new technique, is helpful in terms of providing a deeper understanding of relative frictions at each level of the institution, and the ways in which responsible innovation and anticipatory governance might become a ‘new norm’.

It is also clear from this study that the embedding of RI – here exemplified by its incorporation into CDTs – almost necessarily progresses at different rates, depending on factors such as prior experience and discipline of its members, supporting the ‘granularity’ theories of Watkins & Marsick, (2019). The findings from the study also highlight the role of assessment as a way to reinforce organisational learning, as assessment not only empowers champions of new norms, but drives a need to take new processes seriously. As student-focused organisations the CDTs place great value on student experience, which means the student reception of RI also plays an important role in the institutionalisation of RI. The set-up of the CDT furthermore offers an opportunity for wider integration of RI beyond organisational boundaries, through the industrial partners that CDTs have collaborative and funding ties with.

The study also highlights how reflecting on and coming to a new understanding of existing practices can illustrate where norms may not, in fact, require replacement. This could be regarded as purely a semantic exercise, ie simply changing the descriptions or language around an existing practice to re-frame it as responsible innovation. However, it can also be understood as a way to draw on deeper reflections around the underlying *purposes* of a norm, creating fresh comprehension of incumbent practices as contextually part of a more coherent, wide-ranging and longitudinal mindset. This is complementary to, rather than contradicting, the Randles & Laasch (2016) description of de-institutionalisation – the removal of ‘old’ norms to create space for the new. An example was seen in this study as some CDT Directors re-assessed some of their existing training provision, and drew together modules with underlying ‘responsibility’ themes (such as research ethics and research integrity) alongside new work on responsible innovation. This could be viewed as simply a pragmatic response to needing to incorporate new training requirements, but it can also be understood as a revisiting of the *purposes* of such training and placing it in broader context – as a way to connect science and innovation with the society in which they operate. This broader view relies on an understanding that much of RI is founded on internalised insights – whether this is reflection and anticipation in RI, or ‘mental models’ and systems thinking in learning organisation theory – rather than visible, action-oriented processes.

#### 4.5.1: Implications for the Framework

The findings from this case study will help to inform the design of the Framework, not least by drawing attention to the importance of examining responsible innovation at several levels of magnification – the individual, the research-group or small community, and the wider organisation. All will be relevant in terms of assessing how responsible innovation is understood, actioned, and assessed. This ability to ‘zoom’ in and out in order to understand the perspective at different layers of an organisation, or

project, as well as understanding the importance of clarifying the unit of analysis, will be key in using the Framework.

It also suggests that information-sharing across and between projects may be a key factor in institutionalisation, and that the development of a body of evidence could be of material assistance in fields where 'responsibility' concepts are not yet familiar.

#### 4.5.2: Implications for further studies

The findings of this study also open up interesting avenues for the case studies in the other strands of the triple-helix model – for example, the ways in which responsible innovation is understood and/or operationalised in industry. The question of how a technique such as responsible innovation becomes part of policy will also generate queries for the policy case study – not least because it is clear that institutionalisation, and organisational learning, take place over long timeframes that require ongoing strategic support. This will be particularly key in the case of novel technologies that do not yet operate within a coherent regulatory framework.

An additional point is that, in an example of an unintended consequence, this study may itself have operated as a piece of 'action research' – ie research that, through interrogating a topic, can bring about change. Several Directors were struck, during interviews, by the question as to whether they were working with industry partners on responsible innovation elements, or the prospect of discussing RI with them. It was in many cases a novel idea, but one that appeared to them to have many benefits. It may be that a followup study with the Directors (although outside the scope of this thesis) could therefore provide a valuable waypoint in any examination of the institutionalisation of RI both in universities and in industry.

---

# CHAPTER 5: CASE STUDY – POLITICS, POLICY, AND QCT

---

*“The art of economics consists in looking not merely at the immediate but at the longer effects of any act or policy; it consists in tracing the consequences of that policy not merely for one group but for all groups.”*

*Henry Hazlitt*

This chapter interrogates the ways in which policy is made for and around novel technologies, and investigates the attitudes and approaches of policymakers to what they perceive their role and responsibilities to be.

## 5.1: Introduction

The work carried out for the quantum landscape study (Chapter 3), reinforced the hypothesis formed in the literature review (Chapter 2) that both the politics and the policymaking<sup>21</sup> around quantum computing were highly instrumental in shaping foundations and assumptions in the academic and commercial areas of the sector. As discussed in those chapters, one of the factors that makes quantum computing a particularly distinctive technology is the tight wrapping-together of the policy, research, and industrial elements of the innovation pipeline. As well as this closely-knit relationship between the threads of the triple helix, one of the factors that sets quantum computing apart from other technologies – even from other quantum technologies – is its importance at UK national level. Its status as a ‘sovereign capability’ technology, with regulation in place governing inward investment (the National Security & Investment Act 2021), and the degree of national and international interest that are driving high levels of investment and a push towards commercialisation, combine to make it distinctive in the current novel-technology landscape. This drive towards the commercialisation of quantum computing and the growth of a new sector is generating an encouraging environment for large numbers of startups. Objectives such as the creation of a commercial quantum sector are the product of policy drivers that include the Fusion Doctrine (discussed in section 3.4.1 *Fusion Doctrine*) – the combining of ‘prosperity’ and ‘security’ concerns into a single analysis.

---

<sup>21</sup> These are usually differentiated in terms of *policy* being essentially a guide or a collection of regulations, while *politics* is the activity of governing

This Chapter therefore reviews and interrogates some of the prevailing systems of policymaking in the UK, identifying ways in which the multiple factors at play come together, and examining both the breadth and the granularity of policymaking processes. This examination will provide a setting for the questions around governance, and will inform the design of the Framework.

## 5.2: Background

The priorities of policymakers, driven as they are by global, national, and local economic and geopolitical requirements, provide some context for the development of quantum computing at the national level. The term ‘policymaker’ is used here in the broad sense of “a member of a government department, legislature, or other organisation who is responsible for making new rules, laws, etc”<sup>22</sup> and does not distinguish between, for example, policy analysts and special advisers. Nor does it differentiate politician (elected representative) from bureaucrat (unelected official) – although policymakers themselves are very aware of this difference – as for the purposes of this study all are involved in the development, creation and promulgation of ‘rules’.

With its publication of the Integrated Review (HM Government, 2021) and the Innovation Strategy (Department for Business, Energy, and Industrial Strategy, 2021) published in March and July of 2021 respectively, the UK government has repeatedly made clear that it regards ‘enabling’ technologies such as AI and quantum computing as part of its strategic capability (the ability to harness skills and resources to gain or retain competitive advantage), and a fundamental pillar of its high-level commitment to becoming a ‘science superpower’. As seen in Chapter 3, in recent years the Fusion Doctrine – the understanding that security and prosperity are closely tied together – has led to the development of policy themes around innovations that are seen as having the potential either to help or harm the UK both economically and in terms of defence. This type of commitment to the ‘science superpower’ ambition at the highest levels therefore turns a spotlight onto technologies such as AI and quantum computing, with policy frameworks around investment and national security focusing more closely on these innovations that are seen as key for the UK’s future security and economic prosperity.

Both the Integrated Review and the Innovation Strategy make clear that quantum computing is considered to be a crucial enabling technology. Despite use-cases still being speculative (Inglesant et al., 2021), there is agreement that quantum computing is highly likely to be powerful *in some way* (Khan, 2021). The government’s own research shows that publics require good governance in the case

---

<sup>22</sup> Cambridge English Dictionary

of novel technologies such as quantum computing (House of Commons Science and Technology Committee, 2019), but there is little theoretical or practical agreement about the degree to which public concerns should drive policy (see, for example Savulescu et al., 2021). These various pressures on policymaking give rise to numerous research questions.

### 5.2.1: Research questions

The overarching question for this Chapter is related to understanding the sources on which policymaking processes draw, the aims and objectives they serve, and the outputs they create in terms of governance, enablers, and constrainers.

*RQ1: How is policy developed for novel technologies such as quantum computing?*

This incorporates sub-questions

*RQ1a How can policymakers retain the ability to respond to changing circumstances? – for example the unforeseen rapidity of developments in quantum computing*

and

*RQ1b How do policymakers address challenges around consultation and governance? – especially in technologies where governance is such a priority for citizens.*

This Chapter discusses the difficulties faced by policymakers in long-term strategic planning for novel technologies like quantum computing. For example, quantum computing has been a relatively distant goal for most of its development, but in the last few years has accelerated beyond all predictions. Adapting to these quickly-changing circumstances is extremely demanding for bureaucratic processes that are not designed for rapid changes of gear

Overall, an understanding of these policymaking processes, their advantages and disadvantages, is required. Quantum computing may also serve as a useful case-study for the way policymaking may operate in general in the development of novel technologies.

## 5.3: Methodology

The qualitative research carried out for this chapter was completed in the summer and autumn of 2021. It comprised interviews with 12 policymakers from various walks of policy and policymaking-adjacent professional standpoints, including the Government Office for Science, the National CyberSecurity Centre, and Research Council funders.

### 5.3.1: Identification of participants

Searches were conducted on LinkedIn for policymakers, particularly those involved in the development of policy around novel or quantum technologies, and messages were sent to those whose roles aligned with the focus for this study. Snowballing within interviews was also used to identify additional contacts (eg ‘can you recommend another good person for me to speak to?’) and via invitation emails (‘if you are not able to participate, perhaps you could suggest an alternative?’)

An anonymised list of the various participants is provided below:

Role	Code
Former policy adviser	P02
Head of cross-party thinktank	P04
Senior policy adviser	P05
Science adviser	P06
Head of industry, technology and innovation at think-tank	P07
Former senior policy adviser	P08
Policy consultant	P09
Policy manager	P10
Former policy adviser, now industry consultant	I03
Head of theme at funder	P11
Technical director at government agency	P14
Policy adviser, government department	P15

Interviewees were asked questions including, but not limited to:

1. Could you tell me about your role and the work of your team/department/organisation?
2. What type of policy does your organisation formulate?
3. What roles are involved?
  - a. And how do they work together?
4. How is policy developed and brought together?
5. What about public consultations – how/when are those done, and what happens to the results?
  - a. How about ‘Blackett’ reviews and other types of consultation?
6. What other teams and departments do you work with?
7. Are you concerned with governance?

- a. Do you consider governance part of your role?
- b. Are policies designed with governance in mind?
8. Is policy development different for different areas of research?
  - a. What factors affect that?
  - b. Are there policy ‘templates’?
9. What is your process for choosing what research to support/fund?

These questions formed the basic framework of the interviews, but other questions were included dependent upon the particular expertise of the participant – for example those involved in strategic or horizon-scanning activities were asked about the difficulties inherent in long-range planning, while those with responsibility for engaging with stakeholders were asked about their process for this. Background was also significant – for example some participants had good understandings of local government processes.

### 5.3.2: Analysis

The analysis for this study was carried out in the same manner described previously, via transcription of audio files and then deductive coding in NVivo to categorise responses before analysis.

The policy interviews from the landscape Case Study (Chapter 3) were also reviewed in the light of the codes generated from this study, in order to potentially draw out new issues. There was no overlap in interviewees between the two studies.

## 5.4: Findings

This investigation into the theory, practice, and processes of policy as applied to novel technologies such as quantum computing examines from another perspective the concerns raised in the landscape study (Chapter 3) about the policy background to the quantum sector. The policymakers interviewed in the landscape study were focused on the creation and nurture of a new sector of the economy based around quantum technologies and specifically quantum computing. They drew attention to not only the dual-use nature of quantum computing, but also the possibilities seen for it on both security and prosperity grounds – they perceived its importance to both national defence and to the economy. Their concerns were centred on creating greater capacity, supporting the growing sector, and ensuring that inward investment did not stem from sources that might be regarded as problematic at some future point. There was little attention given to engaging with citizens and civil society to ensure that quantum computing was regarded as trustworthy, or concern that governance was a necessary precondition for that trust.

This study therefore investigates in more detail this impetus for the creation of a new sector, the concerns and pressures that might be shaping that, and – most importantly – the processes that policymakers engage with in the modelling of policy for a new area. The context for this is the set of linked concerns previously discussed around national security, and national prosperity.

### 5.4.1: National interests

As seen in Chapter 3 (3.4.1: *Policy: shaping the landscape*), questions of sovereign capability are central to the assumptions that go into developing policy for novel technologies. These concerns are high on the list of priorities for policymakers as they advance strategies for getting the best out of novel technology.

*“you don't want Alibaba or Tencent or somebody just paying into a startup company ... and getting all the IP for ... pennies to the pound.” P08*

*“we've got a very good view about how things can affect the UK's national security, and I use that [term] in the broadest possible sense” P14*

The point from P14 in particular underlines that national security is thought of as a multi-factoral issue. It is not simply a case of addressing P08's concern about inward investment; instead, issues such as cybersecurity, supply-lines, and economic security are all important for a holistic view of 'national' security. This is a key point with regard to quantum computing and other disciplines that are drawn under the banner of 'sovereign capability' technologies. The big-picture perspective takes a very cautious view of even investment from 'friendly' nations.

*“the market's not 'nice'. And the market is funded at the moment by Chinese money and US money, and neither of those are good for the long term health of the UK” P14*

*“government faces choices... Where does it want to be leading? Where does it want domestic capability? Where is it okay ... for British companies to access technology on the open market? That is a series of trade-offs and choices” P06*

The approach described above by P06 is known as the “own-collaborate-access” framework and is elaborated upon in the Integrated Review (HM Government, 2021). Capabilities can be arranged along a spectrum, with policymakers needing to make decisions about the relative importance of technologies. Some will be prioritised in terms of financial and policy support as critical for the country. Less critical ones can potentially be a joint effort with other countries, and some capabilities can be sourced from 'friendly' nations. There is a high degree of nuance in the discussions about individual technologies with individual countries, and many pressures and considerations – both historic and current – will influence these discussions. An example is seen in the recent moves in both Europe and America to ramp up domestic chip production in order to not have to rely on imported chips for industry (eg European Commission, 2022)

*“some people we talk to on a strategic level ... some people we talk to on a very tactical level” P14*

International relationships are key in other ways too – being able to persuade allies of the value of a particular approach, or even just leading by example, may produce the desired result. There may also be elements of national prestige at stake, with unspoken international competitiveness playing a role.

*“I know that if I stand up and bang my fist on the table and say ‘do it this way because it’s right!’ nobody gives a chuff. If the Indians do it because there’s a market there, suddenly the world takes notice. That’s fine by me, because I get the outcome I want.” P14*

*“there’s always going to be an element of ‘if France and Germany are doing it, we should be doing it too’.” P05*

These types of relationship, including the soft power of providing a governance template to more numerically significant nations, the complex geopolitics of defence and global technological leadership, and overarching concerns around encouraging markets in the ‘right’ directions are important to bear in mind as background to the efforts to build a quantum sector.

#### 5.4.2: Building a quantum sector

As previously seen (Chapter 3), there is a significant policy focus on the commercialisation of quantum technologies and a concomitant pressure to feed out potentially profitable work from universities into marketplaces.

*“In the end, the point is to get it into the world” P08*

Policymakers are well aware of the number of elements that need to be in place in order to facilitate this, however, and the strain that this places on investment and capacity pipelines (as seen in Chapter 3’s study). There is a need to create an environment where there is not only funding from government for research, but a supply of skilled researchers to carry out the work, and numerous kinds of support (for example infrastructure, tax breaks, professional advice) to bring innovations into the commercial world. This objective may generate its own problems, however.

##### ***Encouraging and supporting innovation***

There was scepticism in some quarters around the question of whether it is possible for government to be encouraging innovation, in terms of whether it is the right entity to be trying to achieve this at all. This scepticism is rooted in an observation that governmental and bureaucratic roles tend not to attract or recruit employees who are risk-takers, for obvious reasons. The encouragement of innovation, however, necessitates the acceptance of risk and an acknowledgement that ventures will fail.

*“there is extraordinary density – as in stupidity – when it comes to government engagement with innovation. An unwillingness to take risks ... and to a greater or lesser degree, a disinclination to quite get their arms around what innovation is about.” I03*

This unwillingness or inability to understand the need to accept a certain amount of risk within the innovation space is understandable, given that these endeavours are undertaken with public money.

*“there’s a raft of issues including culture, quality of the personnel involved in it, risk aversion ... government is not doing a very good job. The closest we’ve got to doing a good job was 1948.” I03*

It may also not be simple risk-aversion. Investing in innovative technologies requires choosing a well-balanced portfolio in the knowledge that a proportion of the ventures will fail, and therefore carrying a heavy responsibility to pick winning or successful teams or technologies.

There may also be other reasons for participants taking the view that direct investment from government into risky start-ups is not the most effective way to stimulate the sector – the space is also occupied by large tech companies with deep pockets.

*“the government can never compete with a Google or an Amazon or Microsoft or IBM ... actually, the job of government is - it’s not to pick winners” P08*

There may therefore be a fundamental disconnect between the setting up of funds such as the British Business Bank, and the essential needs of the innovation system.

*“the British Business Bank ... is spending British government money to try and support and stimulate innovation ... so there are efforts, but they tend to be underfunded and over-bureaucratized.” I03*

These are not the only challenges facing policymakers – skills shortages, as discussed in Chapter 3, are a serious issue.

### **Capacity**

The problem of capacity-shortages in the system is recognised at policy level, with high awareness of the challenges many academics and commercial organisations are facing in terms of recruitment. The issue is keenly felt across government departments and within funders as well, frequently connected to wider issues about recruiting international staff. Naturally, there is also competition from large tech companies.

*“this is looking at how to make the UK a more attractive place for innovative talent” P05*

*“the problem is the big companies ... are all poaching the startup people, the researchers or whoever it is – offering them £300,000 a month, and saying, you can set up your own lab in Palo Alto ... who is going to refuse that?” P08*

There is an understanding that “maintaining that pipeline of skilled people and fellowships...” (P11) is a policy-level responsibility, and also recognition of the fact that the current schools and training

curriculum makes sustaining this pipeline a challenge – reiterating the concern seen in Chapter 3 that students receive no quantum training before undergraduate level. This lack of a ‘quantum curriculum can mean that students “*come from engineering, physics, different backgrounds.*” (P05). While interdisciplinarity is undoubtedly a strength in the long run, in the immediate term there are clear challenges here. For example it is possible that the lack of a defined quantum curriculum creates space for some to ‘quantum-ify’ their work in order to gain access to funding. This was noted by researcher participants in the earlier study and is recognised on the other side of the funding equation by those working in the policy space.

*“It’s quite difficult to say ‘is it driven by the real scientific enthusiasm – or where the money’s going?’” P11*

This might justifiably raise concerns around not only the legitimacy of the research, but also the possibility of creating or inflating a quantum computing bubble – a risk identified in the landscape study.

These concerns and challenges around national interests, and the development of a new sector centred on quantum technologies, as well as the priorities set out at strategic level by government commitments like the Integrated Review, create severe pressures on policymaking processes, which are already complex and multi-layered.

### 5.4.3: Process in policymaking

In a study of this size and necessarily limited scope it is of course not possible to completely map out policymaking processes, which vary widely within and between teams, across departments, and between branches of government. Processes are highly contextualised and can be heavily influenced by the personalities, exigencies, and desired outcomes of a given situation. The focus of this study was on policymaking for novel technologies, which is frequently shaped by considerations of economic growth. This is particularly the case for quantum computing (as noted above in 5.1), which is seen as a potent driver of future growth and prosperity. Given the deliberate shaping and development of the quantum computing commercial ecosystem by policymakers, it is important to acknowledge that “*bad policymaking can have really bad effects*” (P07) – and so the processes that are generating policy are significant as an object of study in this area. The particular focus on economic benefit may also be a shaping factor.

In broad terms, policymaking has given aims and objectives in view, particularly in relation to novel technologies, and especially when it comes to the development of a quantum sector.

*“we’re supporting the pipeline, the constrainers and the enablers of developing and commercialising new technologies” P05*

*“the bit that I look after is the bit that thinks about the future” P06*

Policymaking was often characterised in general terms by participants as a largely linear, one-way apparatus that is designed to take a range of inputs at one end and produce policy outputs at the other.

*“you come across a policy area ... you'll have those forums and those discussions, and then you'll have a strategy created and that's delivered, and there might be ... some more feedback loops within that process. But often you have a strategy, 'that's delivered', and then we move on to the next one.” P07*

From a distant viewpoint this ‘sausage-machine’ approach may largely appear to be the case, but the reality on the ground is experienced somewhat differently from such high-level overviews. Processes as they are carried out day-to-day are frequently described as anywhere on the spectrum from “organic” to “messy”.

*“we're all led to believe there is a 'process' of policy, and I know formally there is, but I don't think you could say it's particularly consistent.” P10*

*“there wasn't really a very set way ... there wasn't really a template ... of how to do it” P05*

This degree of flexibility does mean that there is an element of agility built in to the system – in theory there is the potential to respond to a particular need in short order without having to adhere to some kind of checklist or set of guidelines. However, flexibility such as this can also lead to uncertainty, which may counter-intuitively slow down or impede responsiveness. For instance:

*“having things like, 'Oh yeah, all these documents were in different formats and then ... we didn't know who to clear that [with] - that was really stressful’.”*

The responsive, ad hoc nature of much policymaking is well-understood by those working in the area, and there were undertones of concern from participants that the lack of a ‘template’, a formula, or a checklist for processes could lead to important areas being missed.

*“... the thing that should save you in all of this is having a good [team] that will see things coming and that ... is well connected up to the [external stakeholders] ... that works pretty well.” P02*

There are clear tensions between the linear mechanisms that together comprise the way in which policymaking is ‘deemed’ to take place, and the more impromptu, unstructured responses to the numerous drivers of the system. Some participants were less concerned about this than others.

*“we have a proactive and a responsive mode. And we'll balance the two” P06*

Whilst some displayed a strong element of pragmatism in the reasoning for adopting a formal process or otherwise.

*“you do the really formal stuff if you think you're going to get a judicial review.” P14*

Although in this case, ‘formal’ refers less to a particular design of process (such as the the steps to be taken and the order in which to take them) and more to the need to document what was done and the decisions taken (and especially the reasons for a decision) in case of challenge.

Questions of what to prioritise – where a need for policy is seen to be most urgent – are similarly multi-factoral, and influenced by a range of pressures.

### Identifying priorities and anticipating needs

Policymaking frequently needs to respond not only to *external* changes and challenges in society, and in science and innovation, but to *internal* government drivers. There were concerns from participants around the degree to which these various pressures and drivers are frequently not well-aligned, and that the need to constantly balance competing and rapidly-shifting factors creates both redundancies and gaps.

*“continually having that long term strategic vision in mind, whilst also having to be responsive to the changing landscape” P11*

This ‘balancing act’ and competing pressures outlined above have led policymakers to create multiple mechanisms for identifying areas of interest or concern. Sometimes priorities are identified proactively as a result of anticipatory work: this may take the form of purposeful horizon-scanning activity by ‘futures’ teams, or there may be data-mining exercises (such as analysing an increase in academic publication rates around a given area, or finding that the private sector is taking an interest in a particular sector) that will raise flags around areas to investigate.

*“there is quite a big government community of forecasters and future thinkers” P05*

*“a growth in publications, ... a spike in investment – this looks like something we should care about” P06*

As anticipated, there are significant tensions in planning for future activity and taking decisions that will shape a sector or research area. This may be driven by the capacity challenges discussed above, or a perception that a field, or a discipline is going to become significant. There will also be pressures from ministers, who may have their own priorities.

*“you’re looking at trends, and anticipating the best that you can ... And that might depend on the political agenda of the time” P05*

*“the things that really changed it was whether this was a ministerial priority” P09*

Policymakers were frank about the difficulties they experience in balancing the need for long-term plans with the need for agility and responsiveness. They find similar challenges when trying to support existing research areas and commitments, whilst also exploring new challenges and encouraging diverse fields of research.

*“How do you ... maintain that core base of excellent research and also expand and diversify these opportunities?” P11*

However, very frequently this priority-setting is driven by exogenous events – for example a report may be produced by another department that is then passed over to them for action, or global events that overtake any internal drivers.

*“we were handed [a report on a novel technology] – what you then do is engage with your industry partners, and then create a process from that” P02*

*“the landscape changes. COVID for example has thrown a massive spanner in the works of all of our research programmes” P11*

There are pressures from numerous directions, but noticeably, these pressures are generally coming from above, in the form of government or ministerial concerns (in one department, priorities might be driven by something their political leaders had seen in the news), or laterally in the form of recommendations or requirements from other departments. There is an obvious gap in communication and priority-setting here – priorities seem rarely if ever to be driven from the ‘bottom up’ by users, researchers, or other stakeholders (including citizens). There is simply no official mechanism for stakeholders to approach a team and request support, governance, funding or other policy-driven support.

*“there's no line of sight ... for end-user back up to the policymaking decisions” P07*

This aligns with the findings of the NQIT-RRI study that could discover no path whereby individual researchers not already directly in contact with funders or policymakers would be able to report any concerns they had, or offer suggestions (Ten Holter et al., 2021b). This does not apply only to end-users – even those with devolved powers struggle with the challenge of passing their concerns up to the top of the decision-making chains.

*“do local authorities have the autonomy to make those decisions? Yes, to an extent, but they also don't have the autonomy to feed back into central government ... and say, ‘this is where we think investment needs to be made’” P07*

This is significant for many reasons, but not least because much policymaking is supposed to be consultative, drawing on public and democratically accountable fora as a source for both legitimacy and creativity.

### Consultation and stakeholder engagement

Consultations with stakeholders and the trust-relationships this can build are an important element in the creation of wider trust amongst citizens and societies.

*“we're only going to be able to develop these technologies and implement them further into public services or healthcare if we gain the trust of the people. And I think there are ways that we could be doing that better” P10*

An understanding of the importance of trust within consultation processes is very evident. For example, it would be easy to assume that participants are generally focused on bringing their own agendas to the table: in the case of academics this might mean a requirement for funding, while industry participants might be thought to be seeking to enhance their prestige, or be working to bring the government in as a client. However, this was largely repudiated by interviewees, who considered that when consultations take place – under Chatham House Rule<sup>23</sup> or otherwise – stakeholders’ opinions can generally be relied upon.

*“many organisations will upfront say ‘we tried this project and it didn't work’. And ... that builds trust between the private sector and the policymakers” P07*

*“we want to be able to have that open conversation, and that also comes from private sector being honest about the limitations of their technology, of what it can and can't deliver... policymakers - they want to be able to feel they can trust the evidence and the information that they're being sent.” P07*

Policymakers are keenly aware that they need to be able to rely on the information they receive from stakeholders, particularly where these are academic or industrial participants who might be considered to be ‘at the coalface’ in terms of their direct experience of the science and the commercial world. Should that information turn out not to be reliable, or if participants do bring their own agenda to the table, then there may be far-reaching implications for how a policy could be designed or deployed, and what its impacts may be.

*“there's been this highly politicised space ... where if you say the wrong thing or you give the wrong impression, that can completely knock out a whole conversation because there's this lack of trust, going back and forth, of ‘are you telling me this for your own financial interests or are you telling me this because it's actually true and that's what I need to know at this point in time’.” P07*

It is possible that policymakers therefore have little option but to trust their stakeholders, making the conversations around who those stakeholders might be, and the consultations with them, even more significant. In particular some institutions act as ‘trusted advisers’ and take seriously their responsibility to provide fair and impartial advice.

*“[our] role as being in between those academic, industrial expertise and be able to relay those messages to government” P11*

---

<sup>23</sup> “When a meeting, or part thereof, is held under the **Chatham House Rule**, participants are free to use the information received, but neither the identity nor the affiliation of the speaker(s), nor that of any other participant, may be revealed.” <https://www.chathamhouse.org/about-us/chatham-house-rule>

The question of whether these consultation processes are presently fit for purpose will be discussed below.

#### 5.4.4: Process gaps

Policymakers are aware that processes for designing, building and implementing policy are frequently sub-optimal, and various degrees of frustration were expressed at the degree of inconsistency, improvisation, and informality that can attend these processes.

*“I don't even understand what the pipeline is in the current context” P04*

*“we all want better public policy ... we have very different ways of getting there” P10*

Some of the clearest gaps in processes relates to the concerns seen around the handling of consultation (5.4.3 above) – that of failing to seek out a wide enough range of opinions, and the implications of this lack of diversity. This risks losing the well-established benefits of consultation and collaboration.

#### Engagement gaps

*“what we really needed was some kind of ... evidence from the policymakers, on actually ... what is the process that they're doing ... to make a public consultation happen” P09*

*“there are a handful of people in government that give an opinion. That's not healthy” P14*

There are long-understood challenges in designing and carrying out consultative processes, discussed previously in Chapter 2 (see 2.3.2: *Challenges for engagement*). One major concern in the policymaking ecosystem (as with much participatory work, eg Schubotz, 2019) is that simply by asking questions and engaging with publics, there may be an implied – or inferred – promise of action. However the consultative process in question may not have been designed for this, and may not produce any substantive change – it might not even lead to the concerns of the group being fed back to decision-makers. The possibility of creating assumptions among participants can act as one factor disincentivising those who might otherwise engage.

*“I think there's a risk-averse mindset generally ... the bit that worries most ... policymakers when doing participatory work is when they feel like they're overpromising” P09*

Some of the other challenges with participation discussed in Chapter 2 are also well-understood here, such as with *whom* to engage.

*“you need to create conversations, town hall meetings, diverse input from many different angles, from many different lives in order to truly think through the full implications of stuff that you're doing. And you'll still get it wrong” I03*

Decisions have to be taken around what groups, individuals, communities, bodies, organisations, institutions are included, and this selection process may be driven by both conscious and unconscious

assumptions. Any gaps can also be amplified by the difficulty of, firstly, finding out who the people and communities are that it would be desirable to engage with, and secondly, being able to engage with them.

*“we probably don't have quite the same ability to connect to the industrial R&D base that we do with the academics. Frankly, they're less Google-able.” P06*

However, even the most diligent and conscientious policymaker leading a consultation cannot possibly engage with every element of civil society.

*“we tried to ask the policy teams to be holding the pen on who they're inviting because they should know who their stakeholders are.” P09*

*“when it's more focus-groups or when there's more smaller stakeholder engagements or even co-creation things is ... a bit of fear about ‘is this representative’” P09*

Fears around representation may be well-founded – so-called ‘hard-to-reach’ groups may not be included, and this is certainly a legitimate concern in the construction of policies that may impact them. Another theme around the challenge of representation is the question of how ‘representative’ a representative democracy is considered to be. Hypothetically, an MP represents all their constituents regardless of age, gender, disability, employment status and so on – clearly this cannot, in fact, be the case.

*“something I think government must take accountability for - to start talking to their citizens. But often, I guess, the rationale is ‘we don't need to talk to the citizens because the MPs represent the citizens’” P07*

Even if representativeness was assumed to be ‘perfect’, there is immediately another set of decisions to be taken around the information shared and the questions asked – similar to the concerns expressed by academics undertaking engagement work (see 3.5.2: *Engagement*). A key theme in consultation processes acknowledges the importance of the framing and the specific questions that are addressed to consultation groups:

*“they don't quite know how to put these kind of policy challenges or policy questions that they have for the public in a language which people can really engage with” P09*

There is also an understanding that if questions are asked in certain ways, then the answers are likely to be very predictable – citizens will accept developments they perceive to be well-regulated and aimed at the public good.

*“if you've got those three things of public benefit, serious use and independently regulated, it's really quite difficult to find a public dialogue exercise that didn't conclude that [citizens will accept it].” P06*

It is possible that the degree of consensus displayed in such exercises indicates that there may have been a failure to explain the issues properly – P09 noted how important it was to use “*language which*

*people can really engage with*” in order for them to appreciate where challenges might lie – or potentially that the discussions or questions were framed at too high a level, so that deployment issues and/or specific impacts were not discussed. This loss of nuance, and the difficulty of articulating the possible impacts presented by a given technology, may be one factor that leads to relatively ‘blunt instrument’ policymaking.

Policymakers do not only face methodological challenges in consultative work, however; before even reaching this point there is the challenge of resourcing such exercises. This can be particularly acute at devolved level.

*“resources are a huge issue in local authorities ... The engagement that needs to be done needs to go through so many processes ... to engage with those citizens – it’ll almost ... destroys the benefit that you would get from it” P10*

Despite all of these difficulties, however, many policymakers are very clear that there are positive net benefits of carrying out consultation work.

*“when decision making bodies ... take [public] views and concerns into account and are transparently doing so ... that public debate goes better because people feel listened to. Even people that weren’t involved ... can see that ‘we’ have been listened to.” P06*

Participants also had examples to offer of occasions when they had seen processes or consultations go awry because they had not been sufficiently inclusive. There is an understanding that consultation is seen as a desirable process by citizens and civil society organisations, and should form part of the necessary apparatus of a democratically accountable system.

*“I don’t think it’s a winning situation for anybody if you can push through a policy without involving the end user or the individual.” P10*

As a process, it may also have beneficial effects in the longer term (beyond the creation of immediate policy) because consultation can improve the long-term relationship between government, either local or central, and its citizens, and thereby increase public confidence.

*“it’s not always efficient, but ... it’s essentially political suicide [not to do it] so then you think ‘well actually it would have made sense to go through all those processes because we wouldn’t have had to backtrack on this policy’.” P10*

The elements of confidence and trust are key factors in many of the relationships and networks formed throughout policymaking bodies.

### Diversity gaps

Allied to the challenge of ‘representativeness’ discussed above, there was concern among several participants that, particularly in some areas of research, a homogeneous set of voices is problematic in terms of ensuring that a breadth of opinions and new ideas are brought to the table, and that hearing

from broader groups would be valuable. This is allied with an awareness from some that these tendencies to work within tried-and-tested limits need to be challenged. Given the critical importance of having conversations with stakeholders, it can be seen how useful it is to hear from numerous perspectives.

*“there’s a set of misaligned incentives about how advice comes to government... I think we need more diversity of thought in there.” P14*

Some efforts are made to counter this trend to speak to just ‘the usual suspects’ – this is seen as useful not simply to bring more ideas into the mix, but also to ensure that there can be a type of review process where policy suggestions or decisions may be ‘tested’ amongst a variety of communities.

*“central government ... should be wanting to have more views than the ones that are just exist in government ... I don't think we currently have enough external challenging to government to make sure we’re making the right decisions ... we need to fix that.” P14*

The same awareness around the risks of homogeneity may be seen in taxpayer-funded research projects, and public-sector funders are aware that one of their roles is to try and ensure that – particularly in terms of governance frameworks such as steering groups and advisory boards – a range of opinion may be brought to bear.

*“we see that as quite a common issue around some of our advisory boards – they’re not necessarily the type of diversity that we’d expect” P11*

The absence of diverse voices in the policymaking spaces where decisions are made can mean that important perspectives – including evaluation of impact – are lost. The net effect of this is likely to be the creation of gaps of more or less significance – those factors, groups, and individuals that are absent from the process will have an effect on the way a policy is shaped at least as much as those that are involved. Choices about *whether* to consult, *whom* to consult, and *what to ask them*, are often not driven by systems of precedent or guidelines so much as they are by the decisions of individuals and their personal networks.

### Reliance on individuals

Another significant theme in the interviews was that – in the absence of well-understood or widely-adopted templates and guidance – mechanisms in policymaking are frequently person-led, rather than process-led. The approach and the decisions taken will often depend on the knowledge, expertise, or particular requirements of a given individual. Their views will shape their team, and will determine, for example, how often members of their group should be engaging with stakeholders.

*“there’s an unspoken rule in my team that you should be chatting to businesses and researchers once or twice a week minimum.” P05*

In another team, with different leadership, this engagement might be more or less frequent. The personal element may also factor into choosing workstreams, where people select their particular projects for taking forward. In many cases it may be the choice or the responsibility of an individual to kickstart the policy-creation process:

*“ [my head of department] may say, ‘... I think that looks like something that's important’ ... Or a senior customer in the central government could say, ‘well we've heard about this, we're increasingly interested in it, can you tell us what you know about it’” P06*

The converse of this is that other priorities may fall by the wayside, regardless of whether they might be viewed by other individuals as necessary to pursue.

*“If a minister's interested in something then things get driven through ... And if they're not interested they just fall down.” P02*

This can also mean that – should a given individual move on – a policy they were leading may fall into abeyance:

*“it started when [the minister] decided that it started and once he'd gone as a minister, it was on the slide” P02*

The reliance on individuals to lead on and shape process almost inevitably leads to gaps, especially when there is not a good ‘fit’ between those individuals and the topic they may be developing policy on.

*“if you're dealing with the sort of senior officials who don't know a damn thing about [the subject at hand], you've already got a problem” P02*

As with all large, complex organisations, there are challenges around organisational learning and knowledge-management, with a heavy reliance on the personal knowledge and experience of others in a team. Such an emphasis on institutional culture and the importance of individuals can also cause significant loss when individuals depart, with participants worrying *“how do we retain that corporate knowledge, institutional knowledge?”* (P09)

Such reliance on individuals is not necessarily problematic *per se* – institutional knowledge is built up within individual skillsets over the course of often lengthy experience, and such individuals may know the most efficient and effective means of accomplishing particular objectives. However, this is not always the case – as seen above, officials may not have expertise that is directly applicable to their brief, and in an ecosystem that is excessively dependent upon such individual expertise, this is likely to damage the operation of a team.

*“I think a lot of civil servants were allowed to isolate themselves from [knowing about] that. Which is quite bad, but then when you actually appoint one of those people to being director of [the] department, then it's inexcusable” P02*

This can also lead to uncertainty about where to go for information or how to obtain permission – this can cause unnecessary delays and stress.

The ‘personal’ element of this reliance on individuals also means that it is extremely important for officials at all levels to develop networks of relationships, as they are quickly aware that these networks will be a significant factor in the success or otherwise of their projects. Good networks can both extend a policymaker’s ‘reach’ in terms of the people they are able to consult with, and can amplify their own power to affect outcomes.

*“initially it is just about developing those personal relationships, to allow that to progress ... So it's just about finding those initial faces, and then going from there.” P10*

*“what you really need to be doing is ... have the relationships to be able to have more influence.” P05*

These networks can also operate as back-channels for communication and collaboration; *“you know someone who worked with us in a project then will recommend us to maybe someone else”* (P10). They may also be important on occasion for creating a multi-voice consensus that can push policy along.

*“having a consortium of people all saying the same thing, is a good thing” P07*

One of the key uses for these networks of contacts is in the need to cross the boundaries of roles, teams and departments.

### Boundary gaps

Within policymaking environments, boundaries are a key shaping concept (eg *“it's important to bound our work”* P06). Clear boundaries are drawn between particular job roles, between team foci, and between departments. This is not a matter of accident but of design – to try and ensure that all areas are covered, but with a minimum of overlap – *“you get very complicated constellations of decision making in a large organisation”* (I03). This may be designed to be as efficient as possible, but problems can arise when agility or responsiveness are required to adapt to new situations – for example when deciding in whose remit a new technology should be accommodated. This is a particularly acute problem in the case of cross-cutting or ‘enabling’ technologies. As one policymaker pointed out, this has been a vexed question ever since the deployment of electricity (and likely prior to that), and is still no easier to resolve. For example it has been a relatively recent challenge to decide which department should be responsible for the government’s ‘digital’ strategy, or even whether a single or unified strategy is possible or desirable. It is likely that quantum computing and other enabling technologies such as AI will also fall into this category.

*“it's institutionally quite hard for government to manage things that are spread across sectors. That question of ‘Oh, is it them? Or is it us’ ... keeps coming up.” P02*

These boundaries and questions of jurisdiction are designed to ensure that it is possible to be clear about who is *responsible* for a given task, or piece of the policymaking puzzle.

*“governments handle the world by breaking it all into bits” P02*

Unfortunately it is rarely possible to classify the real world with such precision. This sets up a constant tension between the ‘messy’, cross-cutting, multi-factoral challenges of day-to-day demands on policymakers, and the defined boundaries that they work within.

The lack of overlap between teams can lead to a silo-ed focus, unless countered by both strategic and operational efforts at communication to cross the boundaries. Ironically a lack of overlap may itself lead to duplication because departments as teams can remain unaware of what other departments or teams are focusing on. The exigencies of the pandemic may have improved this situation – *“I’m 100% sure that it’s the pandemic that has made us talk better”* (P14) – but it is highly probable that there are still seriously fragmented approaches and duplicated work.

*“policy teams are focused on their specific policy objectives ... which might be complementary but which might be pulling in slightly different directions as well.” P05*

This lack of overlap between teams or departments can even extend to being the only person in the room who is knowledgeable about a given subject – a situation that can create a significant level of discomfort and concern.

In particular, the boundaries between policy and politics are unmistakably defined. Civil servants advise, analyse, consult – but they are clear that the ultimate decision-making power rests with elected officials.

*“we can give advice but in the end it’s ministers who have to make the decisions.” P14*

However, boundaries are potentially able to be consciously set aside when it comes to a novel area of policymaking or technology, where it may become a question of establishing control over something new, interesting, or prestigious – a reflection of the Abbottian principles of professions (Abbott, 1988).

*“once everybody starts clawing for it ... it means more jurisdiction, more power, more – it’s a land grab.” P08*

*“turf wars, you know ... it’s already both big ‘P’ and small ‘p’ political.” P08*

The converse is also true – if there are challenges that officials in a team may not want to be involved with, adherence to boundaries can be a way of sidestepping these. This is particularly seen with cross-cutting issues, or ‘enabling’ technologies – it is frequently possible to make a case that ‘this should be handled elsewhere’.

*“a generic question is ... should this be in [X] department ... or [Y] department? and that question doesn't necessarily **have** to be answered, it can be shared ... But once senior officials start asking, you get into this nightmare of ‘Oh, can't we make them do it?’” P02*

There is understanding that the multi-layered, disparate shape of policymaking can create gaps of this type, and there are some efforts to address that:

*“there are functions that are trying to pool together what's happening across different departments, and ... I'm in a team which is trying to do just that ... pull together insights on technology from lots of different areas.” P05*

However these efforts are not seen across all areas, and this type of communication does not appear to be prioritised. This challenge may be particularly acute in the case of governance, which requires cohesive and well-informed responses between and across teams.

### Governance gaps

The question of the governance and oversight of novel technologies is absolutely central to their development, because of the need for public trust and societal benefit. As was seen in Chapter 2, public acceptance of novel technologies is predicated on factors that include a basis of sound governance, particularly in technologies with a high degree of uncertainty about use-cases. It is clear from the way AI has been developed – and the reputational damage it has suffered due to perceptions around poor quality datasets, unreliable machine-learning, and algorithms that amplify bias – that there may be performance gaps here between what is necessary for public trust, and what is in place. This gap is recognised within some quarters

*“we were very, very interventionist in the first few years [here] in order to reshape that market and make the incentive models much better. You do exactly the same in technology. That's what we should be doing in quantum computing ... - you shape a market to want what you want to deliver.” P14*

*“the decision to back a technology area is not just about how much money government puts into it, it's also about making sure that the regulation's appropriate and adaptive, making sure that you've got the right skills in place.” P06*

These comments are particularly relevant to the development of incentives within the quantum computing sector that can assist in shaping a governance strategy.

*“you often get policymakers or civil servants or parliamentarians who will have seen something in the news ... and that suddenly creates an opportunity to talk about the regulation and governance of technology” P07*

However, it is possible that existing mechanisms for alerting policymakers to the need for more robust or stringent governance in an area are left more to chance than to analysis.

*“if we think that there's signs that this technology doesn't fit within existing regulatory frameworks ... if experts were coming to us and saying, ‘... This technology doesn't fit*

*within what we've used to govern similar things anymore.' That's something that we would include in our advice to government." P06*

There may be a sense that – due to the mechanism of funding for quantum computing and other novel technologies being managed via the UK Research Councils – the governance side is ‘handled’. Research Councils and universities have strict compliance regimes, and retain oversight of large projects in order to ensure good governance as well as tracking impact, onward funding and other metrics via its reporting lines (eg through the ResearchFish tool).

*“them getting the next round of funding is contingent on them providing documentation, evidence, impact, evaluations, all that stuff.” P08*

At this stage in the process, therefore, governance is to some degree ‘built in’. It will be necessary to follow this thread through to the industry study in order to ascertain whether such compliance structures are retained through the spinout stages. It would also be an important (and very different) question to examine in the development of privately funded research projects, or research within large companies.

### **Procurement as governance**

*“the government can influence procurement ... that's quite a big thing” P04*

Procurement was discussed briefly in the literature review as a frequently-overlooked tool in the governance armoury, possibly because of its position at the tail-end of the innovation chain. In very simplified terms, by this point foundational research has been funded, spun out, and commercialised into goods and services. Governments, whether local or central, then become significant purchasers of these goods and services, and have within their power the ability to specify criteria that suppliers must meet before they can win a public-sector contract. Although processes are extremely carefully designed and followed (or in the words of one interviewee, *“procurement processes are incredibly stultifying, and calcified, and risk averse.”* I03), there is some evidence of a change in the way in which these processes are being designed. From being entirely focused on ‘value for money’ and the end result to the taxpayer’s pocket, government purchasing since the Public Services (Social Value) Act of 2012 is now expected to centre “social value”<sup>24</sup>.

*“Government is now obliged to apply a social value metric [in] every piece of procurement that it does. And what we're thinking about in terms of ethics and social good should be featuring therefore, in what government does. And that should ring bells in the minds of those who are involved in governance of companies,” I03*

---

<sup>24</sup> <https://www.legislation.gov.uk/ukpga/2012/3/enacted>

This requirement to consider “*economic, social and environmental well-being*” (Public Sector (Social Value) Act, 2012, Section 1, 3) should provide procuring bodies with powerful levers to consider questions of societal impact. The Act, which was renewed in 2018, is seen by some commentators as a response to failures of innovation policy that have left unaddressed such challenges as population ageing, poverty, or climate change (Uyarra et al., 2019) – the aim is to spur innovation in directions seen as societally desirable. However, participants saw little sign that the Act has made any real change to procurement.

*“technology that's used in the public services and public sector is acquired from the private sector ... And the people making those procurement decisions – are they having those open conversations? ... Are they really considering ‘does this fit into the social value framework’”*  
P07

The expertise of individuals – or lack of it, as discussed previously – is also a significant factor in this area.

*“[in]the procurement of technology... people who are having to assess technology but haven't necessarily had the skills in order to do that assessment”* P07

This shortage of in-depth knowledge in key areas – particularly when it comes to purchasing – may also have significant knock-on effects. As seen above, individual knowledge is key in policymaking processes. Where such knowledge is lacking, it is not just the knowledge itself that is absent, but also the bare awareness that such a gap exists.

*“There was literally nobody in government who could challenge anything I said on it. I could be talking complete nonsense, and nobody would have known ... that's not a comfortable place to be”* P14

This also raises the question of how these procurement processes are evaluated. Participants were unsure how or whether this (mandatory) incorporation of ‘social value’ was in fact being measured or assessed – it appears to be an open question as to where responsibility lies for ensuring that the social value framework is adhered to, or whether such a responsibility is even perceived.

### **Assessment gaps**

These unresolved questions of assessment and evaluation are also seen more broadly throughout policymaking. In the opinion of several interviewees, a noteworthy gap in the overall process of policy was the appraisal of work that had been completed.

*“that's something I'd like to rectify ... - people really thinking about previous policy and government interventions ... that has had an effect ... but there's no assessment for us to say whether that's true or not”* P07

*“the [Integrated Review] ... says ‘we're going to be this technology superpower, this science and tech superpower’. Okay – we'd better make some bloody good decisions then.*

*Right? And the current structures we have for making those decisions I don't think are quite right" P14*

As argued previously in Chapters 2 and 4 (eg Owen et al 2020), assessment and evaluation are key elements in the institutionalisation of processes – in order to become ‘taken-for-granted’ they need to be accountable through metrics that can provide transparency in the process, as well as providing understanding of where improvement is needed. This assessment gap applies not just to the processes of procurement and their alignment with social value concerns, but also to the concerns seen earlier around stakeholder inclusion, and boundary gaps. Assessment is a critical factor in governance and accountability – it is a management truism that *‘what doesn’t get measured, doesn’t get done’* – therefore if something needs to be done, it also needs to be measured. This understanding is not apparent in much of the policymaking sphere. An example of this is seen with the ambitions around the “science superpower” aspiration mentioned above. This target has been a stated aim of the last several UK governments, but a highly critical House of Lords report dismissed much of the rhetoric as hyperbole, pointing out the lack of joined-up policymaking, the belief that the provision of funding is all that is needed, and the absence of follow-through. The assessment gap is clearly seen here – without agreed metrics, indicators of success, and evaluation, the ambition to be a science superpower will inevitably fail (House of Lords Science & Technology Select Committee, 2022). Unless these challenges are addressed, this lack of cohesion – particularly for a cross-sector ‘enabling’ technology such as quantum computing – is likely to have serious knock-on effects for oversight of its responsible deployment.

## 5.5: Discussion

One of the largest challenges for policymakers is that they are at a minimum of one remove from the objects of many of their policies. They are not researchers running projects, nor are they business-owners managing companies – yet they must design and plan for the environments where these endeavours can flourish, while possibly having no direct experience of their own to bring to bear, and while also taking account of larger societal pictures and shifting priorities from government. The techniques that have been devised for managing these challenges bring tensions and difficulties of their own.

### 5.5.1: ‘Breaking it all into bits’ vs real life

As seen in section 5.4.4 (*Process gaps*), policymakers face a constant tension around the problem of new technology or information that is not easily categorised. Policymaking systems are predicated on a model of the world that is essentially neatly divisible – this is not a framing that is especially

representative of real-world issues and challenges. Their responsibilities are partitioned in ways that can be seriously tested when an issue or opportunity presents itself that straddles boundaries.

As demonstrated in 5.4.4, it is clear that within advisory processes and policymaking teams these boundaries are well established between, for example, those who perform horizon scanning activities and those who create strategy, or those who carry out analysis and those who generate policy. This may in part be simply a function of operating in a very large, fractured, and spread-out organisation (or indeed set of organisations). Such boundaries are at least partly designed to ensure that there are not overlaps in function that might make for wasteful duplication, and also of clearly outlining areas of jurisdiction and accountability. However, rigid boundaries are also liable to create gaps in responsibility when novel challenges or opportunities appear that require a policy response but do not fit neatly into the pre-existing framework. This may result in one of two outcomes – either a ‘land-grab’, when the opportunity is seen as a desirable one, or a hardening of the boundary to try and ensure that the challenge is allocated to another team’s remit. More critically, the gaps may go unrecognised, and issues that fall between jurisdictions may be missed. In any eventuality, the cross-boundary working that may be required in order to create a cohesive response will be difficult to achieve.

The seminal work on professions from Abbott provides a useful theoretical lens with which to examine these responses by policymakers to the challenges of their work (Abbott, 1988). His analysis focuses on the ways in which jurisdictional struggles such as those seen here can drive changes in internal organisational structure and *“where success is continuously reshaped by struggle with adjacent professions, and by wider social changes”* (per Cox & Corral, 2013, p1527 ). Although policymaking may not be recognised as a profession in the same way as teaching, or medicine, it bears many of the hallmarks of one (*“a knowledge base, autonomy, an association, a code of ethics and values, and high economic rewards and social status”* – Cox & Corral, 2013), and certainly there are many policymakers who regard themselves as professionals<sup>25</sup>. Abbott’s analysis of professions, and its grasp of the competitiveness inherent in many of the daily interactions within the space, may therefore provide a valuable model for understanding how policymakers interact.

This is directly linked to the second set of tensions seen in this space, which may indeed be a response to the challenges seen here.

---

<sup>25</sup> See, for example, <https://www.gov.uk/government/organisations/civil-service-policy-profession/about>

### 5.5.2: Organic vs systematic process

The tension here is between the ‘messy’ methods of policymaking described by several participants, and the linear, structured mechanisms that are at least nominally in place for policy-creation work. The ‘rational’ model of policymaking has been the accepted ideal for many years (eg Lasswell, 1956), and as Botterill & Fenna, (2019) point out, “*illusions of rationalism have enjoyed a charmed life*” (Botterill & Fenna, 2019 p77). Both structured and messy approaches have their benefits and disadvantages. Messy processes can be robust and flexible (Rawluk et al., 2019), with the ability to integrate interdisciplinarity and adapt quickly to new needs. This may be particularly seen in cases where individual participants have a degree of professional agency within their roles – they are potentially empowered to take the initiative in a new situation and act quickly. Messy processes can also incorporate a degree of redundancy, which can improve resilience. However, a messy process may also mean it can be difficult to identify who is the responsible party in a given situation – there can be challenges attendant on relatively basic tasks such as achieving ‘sign-off’ for a particular action. This may be complicated by the proliferation of leadership styles generated through the emphasis on individuals, where one manager perhaps has different requirements than another. A lack of clarity around who is allowed to permit (or forbid) a given action can thus exacerbate the risk-aversion for which policymakers are already well-known, as well as creating stressful situations for teams and individuals.

This tension between *messy* and *structured* represents a stability-plasticity dilemma – a constant and ongoing push-pull between the need for certainty and the need for adaptivity. Such a tension is comprehensible in the large and complex set of organisations that together make up the ‘constellation’ of decision-making bodies comprising policymaking structures. This tension, between certainty or stability, and adaptivity or plasticity has been explored by Thelen and other authors (eg Mahoney & Thelen, 2010; Thelen, 2002) in their work on the way institutions adapt and change. As seen in the earlier discussion around the institutionalisation of RI (Chapters 2 and 4), institutions have ‘incumbent logics’ that create inertia – but within these structures, individuals with agency are able to bring about substantive change.

The messiness seen in policymaking spaces is not necessarily, therefore, a problem *per se* – but problems are likely to arise if the tension between the ability to adapt and the need for certainty goes unrecognised. This may be what gives rise to some of the worries and uncertainties that surfaced in the interviews.

### 5.5.3: Individuals and networks

The need to form networks and personal relationships across teams may also be linked to the structural problem of “breaking it all into bits” – without defined boundary-straddling protocols, and systems for cross-team working, individuals fall back on the formation of unsystematic personal relationships. It is clear from the findings above that individual personalities are important in the policy space, and this may have knock-on effects when considering the interpersonal relationships that coalesce to form networks. Several participants pointed out that these networks are as often informal (eg “once you’ve gotten in with [Dept A] they can introduce you to [Dept B]” P10) as they are formal – for example the Heads of Horizon-Scanning group, or Chief Scientific Adviser Network. It is in the genesis of such informal networks that individual personalities can be key, and these may be critical when it comes to decisions around sources of expertise – for example a policymaker who has particularly good relationships with a given research institution may choose to consult their contacts there, rather than at another institution. This will affect the discussions they have, and then possibly the decisions they take as a result. Given the absence of formal stakeholder identification and consultation processes, this can amplify the impact of an individual as their particular network of contacts then feeds through into policymaking decisions.

## 5.6: Conclusions

The findings of this study underline the conclusions from Chapter 4 of the importance of using varying levels of magnification to examine challenges around governance. In that case study it was evident that individuals – their understanding, interest, and enthusiasm (or lack of any of these) – were key in the rollout and institutionalisation of the anticipatory governance processes of responsible innovation. In this policy study it is possible to see some of the same patterns with regard to policymaking priorities, directions, and processes. Individuals guide the shaping of policy, and their own experience, skills, and personal relationships influence the way in which policymaking processes are carried out – for example in decisions around what type of consultation to design, who to consult, and what to ask them. This relates directly to the ‘politics’ questions that arose as part of the literature review (2.3.2: *The politics of inclusion*) and is deserving of further research, as it may be hypothesised that individual biases in forming networks, conscious or unconscious, could have a marked effect here. The way these personal relationships affect the formation of teams and networks also supports the ‘granularity’ conclusions from Chapter 4.

The understanding of competition that comes from an Abbottian analysis of policymakers as professionals is also valuable in explaining some of the territoriality and jurisdictional struggles seen in these teams and departments.

Additionally, the findings of this study suggest that – at the macro level – policymaking processes could be characterised as a *Complex Adaptive System* (CAS), per Holland, (1992). CASs are systems with a large number of relatively independent, distributed parts that are able to act independently, but that as a whole “*adapt themselves to the problems posed by their surroundings*” (Holland, 1992, p18). They have certain distinguishing features: they show evolution (responding in new ways based on learned experience); aggregate behaviour (where the outcomes of the whole are more than simply the sum of the relevant parts); and anticipation (the ability to “*look ahead to the future consequences of current actions, without actually committing itself to those actions*” (Holland, 1992, p. 25). Although CASs were initially conceptualised as a means of translating large and complicated systems into computing terms, the theoretical framing of CASs has since been adopted to describe many different types of system, at micro and macro levels, from human immune systems to economies. This includes work within the organisational and managerial literatures that focuses on more grounded and less mechanistic models than have traditionally been promoted for business success. Work from Reeves & Levin, (2021) for example, seems particularly apt for the policymaking ecosystem, as they describe ‘messy’ management as a means of maintaining robustness, and recognising that the effective levers in a complex system are not always the obvious ones. A key feature of such systems is also ‘emergence’ – where events and interactions cascade to reshape the process, causing agents to respond again in an iterative, developing process. The final characteristic, also relevant to policymaking, is that such systems are always sub-optimal. They are always in the process of ‘becoming’ and so outcomes can potentially be the result of a series of compromises, rather than a one perfect solution. The need to trade-off multiple factors, balance competing interests, and manage personalities – as well as the numerous other processes that comprise the policymaking ecosystem – gives rise to exactly this type of outcome: sub-optimal and perpetually in the process of ‘becoming’.

The research questions outlined in 5.1 at the opening of this chapter asked how policy is developed for novel technologies such as quantum computing, and included a need to investigate how policymakers can adapt to rapidly changing circumstances. Viewing the policymaking ecosystem through a CAS lens, as above, it seems clear that the size and complexity of the numerous institutions that collectively make policy does not align well with the relatively simplistic, linear models that comprise the usual descriptions of policymaking. Rather, policy is developed for novel technologies through a constant need to balance multiple factors, work with incomplete information or contradictory requirements, and plan both for immediate and long-term outcomes. That this process is necessarily ‘messy’ should not be surprising, but the lack of frameworks or assessment is likely to mean that policy is persistently behind the curve of novel technologies. It is possible that a recognition of this necessary ‘messiness’ might permit of a more helpful modelling, and consequent improvements to the ecosystem. More

coherent consultative processes, better assessment, and an understanding of the need for iteration, could enhance responsiveness in policymaking.

The final research question centred on the ways in which challenges around consultation and governance can be addressed by policymakers – this is especially the case with regard to technologies where governance is a priority for citizens. It is here that the policymaking system may not be responding as well as could be desired – it is not clear that these questions of governance are, in fact, being addressed. The priorities of citizens fall a considerable way behind those of central government, whose drivers in relation to many new technologies – including both quantum computing and AI – are based around economic prosperity and national security (as defined by the Fusion Doctrine), and are likely to be relatively short term.

### 5.6.1: Responsible innovation in policy

Although responsible innovation is not a recognised methodology in policymaking, it is very possible to regard many of the processes and procedures policymakers use as inherently a form of *de facto* RI. Policymakers are familiar with horizon-scanning and planning activity (anticipation), consultation (stakeholder engagement), and response. It is highly likely that reflective processes – whether formal or informal - are also built in, whether or not they are recognised as such. This understanding aligns well with the previous discussions around the importance of understanding responsible innovation in both a macro and a micro sense – using different lenses at various levels of magnification affords a helpful degree of perspective on the processes at work in policymaking. It would be valuable to research this further with policymakers to assess whether this is a framing that chimes with their own thinking, as well as to provide additional insight into the institutionalisation of responsible innovation.

### 5.6.2: Implications for the Framework

This study, focused as it was on the processes in policymaking, gaps in those processes, and the consultative elements in particular, carries several important implications for the development of the Framework

The inherent ‘messiness’ of the policymaking process provides a valuable steer for the development of the Framework. It indicates the value of retaining elements of informality and redundancy that allow a system to adapt rapidly to new situations.

It also suggests that attempting to adhere to a linear, methodical process may be an unrewarding endeavour, as well as not reflecting the reality of developing technologies, and that consciously recognising and adopting more adaptive, robust approaches might reap rewards in terms of better

responsiveness in policymaking. Greater responsiveness may mean that – per the Collingridge dilemma (1980) – smaller ‘nudges’ may be possible rather than large-scale governance efforts.

The importance of assessment is also key, and is linked to the question of linearity – there may currently be no space in existing processes for the decisions and outcomes of policymaking process to be reviewed and assessed. Building in iteration, with its necessary reviewing and recombining of different elements, as well as the process of repeatedly ‘going through’, can provide an opportunity for these evaluative processes to take place in an informal way.

This study also further underlines the importance of the choice of issues that are selected for consultation topics. The experience of P06 in public consultation, where the answers were predictable because asked at a very high level, provides a useful shaping concept for the Framework. Taking complicated questions out of the equation acts as a type of filtration, and can further amplify the effect of a lack of diversity. The study therefore also supports the value of asking difficult questions and – even if those can’t be addressed immediately – retaining those questions for future iteration.

### 5.6.3: Implications for the industry study

Many of the questions and concerns raised by this policy study will be experienced ‘in real life’ for the day to day operations of businesses working in the nascent quantum sector. Different challenges will be seen at different levels; academics need support as they venture into the world of commercialising their research, while startups need to attract funding. Larger organisations face the challenge of generating turnover to justify their valuation and create a return on investment for markets. The ways in which policy succeeds (or not) in creating and supporting a quantum computing sector, the enablers and constrainers that it instigates, will have direct bearing on the experience and achievements of commercial organisations, and investigating their perspective will be valuable to assess how policy is affecting developments on the ground.

In particular, the power of direct or indirect government investment will be of interest, as well as the capacity challenges in the sector, the innovation space in general and the translation of governance from academia into the commercial world.

---

# CHAPTER 6: CASE STUDY – COMMERCIALISATION AND RESPONSIBILITY

---

*“It is not only what we do, for which we are held responsible, but also for what we do not do.”*

*Moliere*

This chapter investigates the commercial landscape of quantum computing, how it is developing, and what the implications are for governance of the ways in which it is growing and changing.

## 6.1: Introduction

The previous Chapters have examined the landscape for quantum computing within the triple helix of academia, policy and industry (Chapter 3), and then conducted detailed studies within the academic (Chapter 4) and policymaking (Chapter 5) contexts. The final field study examines how the effects of decisions taken in the policymaking and academic contexts are experienced within the commercial<sup>26</sup> environment for quantum computing. This includes not only the funding and governance decisions from policymakers that translate the foundational research into businesses and shape the sector – including the impacts on capacity and the availability of investment – but also the ways in which values, methods, or principles from academia such as responsible innovation may or may not be carried into the commercial environment. Externalities such as the prevalence and effects of quantum ‘hype’ will also be considered, as these may be part of the research and development story from the very beginning (for example in the concerns from landscape study participants that researchers were ‘quantum-ifying’ their projects in order to receive funding). This Chapter therefore re-examines from this new perspective the points raised in previous Chapters about the status of responsible innovation in industry, and how it becomes embedded in organisations, as well as investigating the specific

---

<sup>26</sup> Within economic practice, ‘industry’ is thought of as the procurement and processing of raw materials into finished products, while ‘commerce’ is the exchange of goods and services for value. The development of the quantum computing sector straddles both of these activities, and so the terms are used interchangeably.

factors that may be in operation around quantum computing and quantum-computing-adjacent companies.

## 6.2: Background

As seen in previous Chapters, the UK government is committed to the creation and support of a quantum sector, and Figures 2 and 3 from McKinsey (McKinsey & Company, 2021) provide some context for this. They demonstrate the global acceleration in investment in quantum computing over the last few years.

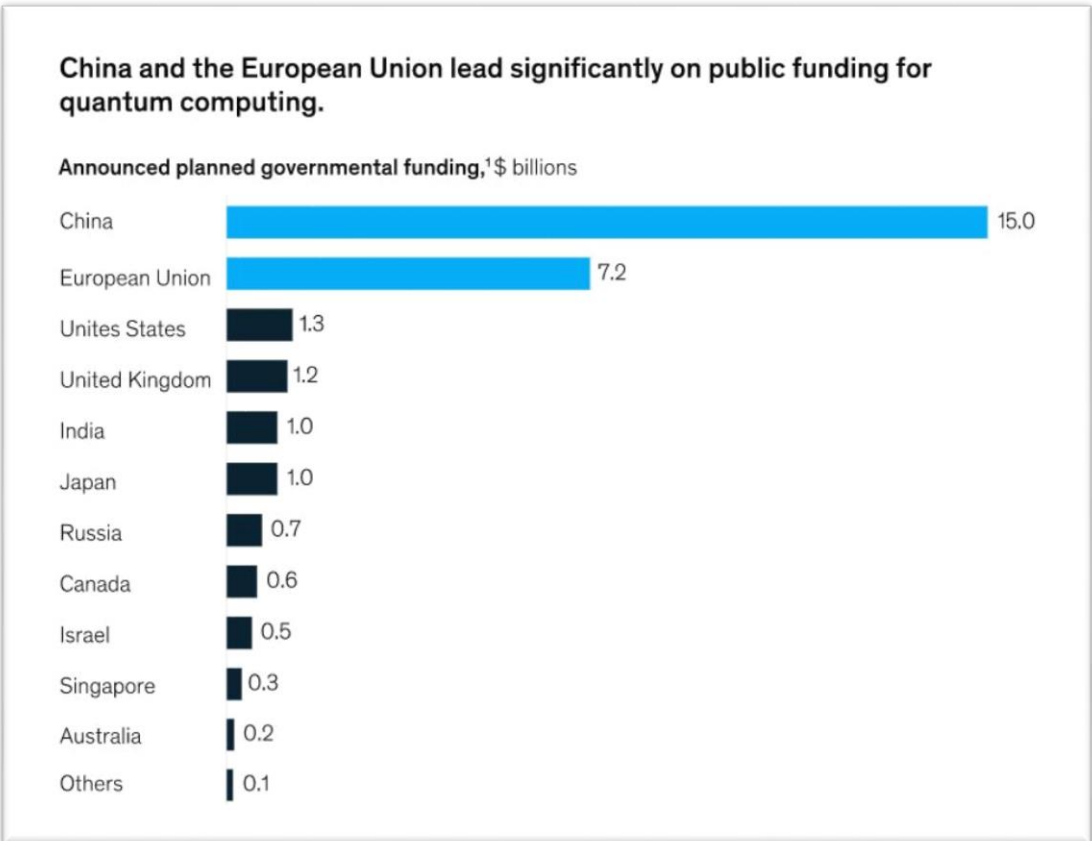


Figure 2: Global public funding for quantum computing (Source: McKinsey & Co 2021)

Figure 2 shows the sums committed by governments worldwide (ie public-sector funding) while Figure 3 shows the increase in private-sector investment and number of startups since 2015. These startups now populate every stage of the quantum computing 'stack', from hardware through operating

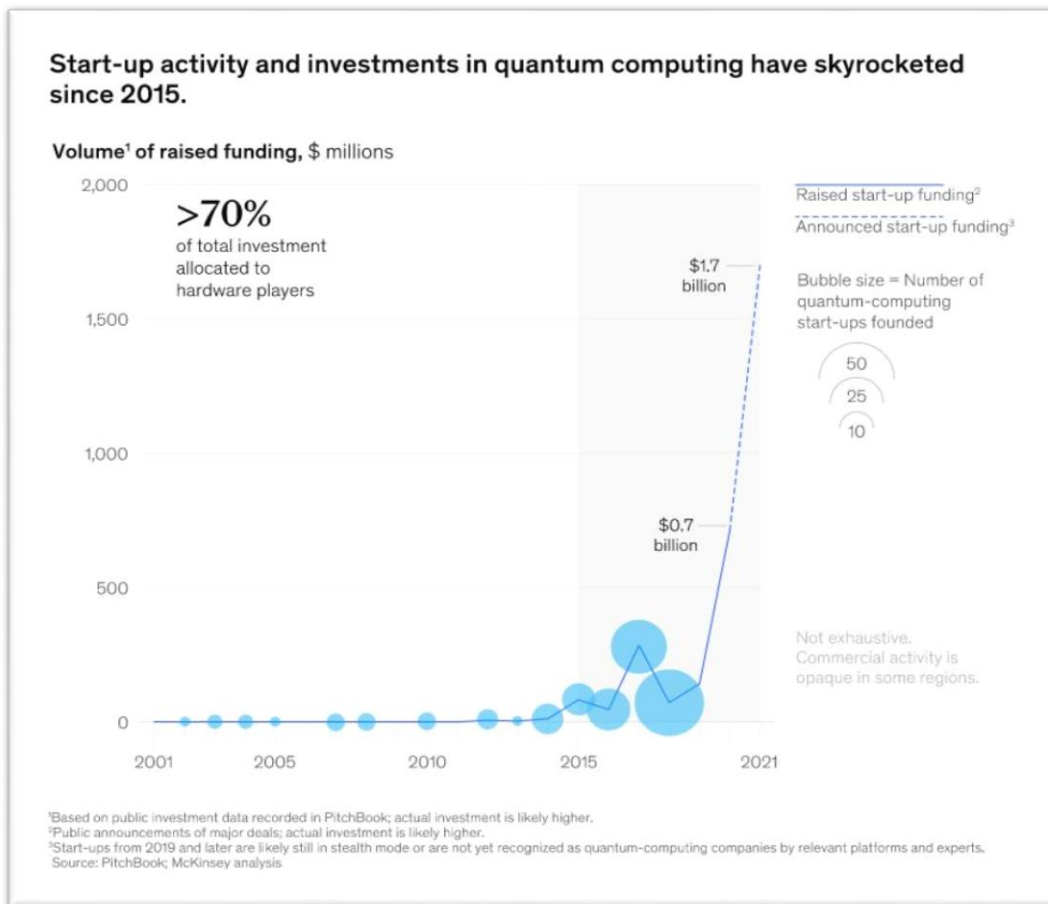


Figure 3: Private investment and quantum computing spinouts

systems to algorithms and programmes – despite the fact that no one form of hardware has yet emerged as the preferred option. This surge in investment can be seen across some of the largest national and international-group economies of the world. This sectoral development is also evidenced by a shift in mainstream press discussion from largely science-based innovation and technology stories towards business-focused news. This is seen in the focus of reporting in business and finance outlets such as the Financial Times – eg *“The 35-year-old field is finally transforming from a scientific endeavour into a burgeoning industry.”* (Murgia, 2021). There is also growth in sector-related areas such as specialist recruitment, professional services, and market analysis, while markets themselves demonstrate business-focused activity such as floatations, takeovers and consolidations. In October 2021 IonQ became the first quantum computing company on the New York Stock Exchange, and in the same month Rigetti announced that it would also float on the public market. Only three months previously Cambridge Quantum Computing had joined forces with Honeywell to form a new enterprise, Quantinuum – now the largest quantum computing company in the world. Japan has joined the race, seeking to have 10 million users of quantum computing (the number considered to have been a tipping point for the use of the internet in its early days) by 2030 (Baker, 2022). Markets

have responded rapidly to this perceived maturation of the quantum computing commercial sector, with May 2022 seeing the first attempt to deliberately ‘short’ a quantum stock (Konig, 2022).

The acceleration of movement and growth in the sector means that academic analysis of the commercialisation of quantum computing, and research into the translation of responsible innovation processes into industry, is essentially happening in real-time.

### 6.2.1: Impacts of funding availability

The landscape study (Chapter 3) demonstrated that the availability of finance for developing the quantum computing sector is potentially a limiting factor on sectoral development, as venture capital funds generally look for a defined exit route for their capital in order to be able to make returns to their investors (this may be anything from five to ten years). In looking at funding for hi-tech start-ups the rule of thumb is that of every ten ventures funded, seven will fail, two will perform solidly, and one will be a ‘star’ – performing well enough to make returns that cover the whole portfolio. However, this investment window is too constrained for a technology such as quantum computing, where there are such significant uncertainties around the viability of the hardware, and there is not even yet consensus around the ‘best’ way to create qubits (or whether there is one ‘best’ way). As discussed in section 3.4.2 (*Investment and capacity*), quantum computing requires non-traditional financing routes, and it was suggested in the earlier study that government finance could be an attractive option.

There is clearly also an impact on governance – public sector financing will have very different implications from a governance perspective, compared to private investment. Whether public or private, investment funding will impact the governance of a sector, depending on the priorities of those making the financial commitment. As has been previously argued, governance is linked to societal embedding for novel technologies, so funders’ priorities may be a key factor to consider.

The availability of funding is also directly linked to the capacity challenges being faced at all stages of the innovation chain. Without sufficient funding injections, companies are unable to hire the quantity or depth of expertise needed, and when highly-qualified employees are in such short supply, bidding-wars for talent both nationally and internationally are extremely probable. Questions of capacity are also linked to anticipatory governance concerns, as capacity-building can provide an opportunity to address the challenge of diversity and inclusion in the sector.

### 6.2.2: The problem of inclusivity in novel technologies

As has been argued from the outset, questions of diversity fall within most understandings of responsible innovation because of the requirement for inclusivity, which is crucial both for anticipatory work and for evaluation of societal impact: “*Responsible innovation should not just*

*welcome diversity; it should nurture it.*" (Stilgoe et al., 2013 p1573). These questions are important to consider within areas of novel technology, for many of the reasons discussed within the policy case study (Chapter 5) – representativeness and diversity can improve challenge, robustness, interdisciplinarity, the generation of new ideas, and help to ensure that multiple sectors of society are included within areas of growth and prosperity. However, the fields of physics and computing from which quantum computing entrepreneurs are in general drawn, are known for their homogeneity. Despite many campaigns, the percentage of girls taking physics at UK secondary schools has increased only fractionally – from 21% in 2016 to 23% in 2021 (Pagel, 2022), while the female to male ratio in physics in general is 23:77 despite women being in the majority in the overall population (Hillier, 2020). Representation of minorities is also extremely limited, with the number of Black professors of physics in the UK in single figures (for example the University of Oxford's Department of Physics did not, at the time of writing, have any Black faculty members). This is likely to have significant implications for questions of inclusivity in the commercial quantum computing sector, particularly when compounded by the challenges that women and minorities already face in raising venture capital<sup>27</sup>. These challenges are so pronounced that specifically Black VC funds have been set up both in the UK and US to support Black and minority entrepreneurs, and there are similarly focused funds supporting female entrepreneurs. Additionally, the language of physics is almost exclusively English, creating barriers to entry for Global South countries. Collectively these challenges mean that quantum physics and by extension the quantum computing ecosystem are in an ongoing position of failing to take advantage of the well-documented benefits of diversity (Herring, 2009). This need to "nurture" and encourage diversity as part of the process of translating responsible innovation approaches into industry will also be examined in this field study.

### 6.2.3: Responsible innovation in industry and commerce

The majority of the literature agrees that although RI has never been limited solely to academia, and several commentators have examined its use in industry (eg Martinuzzi et al., 2018), there has not been widespread adoption of RI principles in the industrial context (Nazarko & Melnikas, 2019; van de Poel et al., 2017), However, as shown in the study on the institutionalisation of RI (Chapter 4), the question of how responsible innovation or anticipatory governance concerns can translate into industry is an important one. In that study, for example, CDT Directors were asked about whether they discussed RI with industry partners, or planned to involve industry partners in RI training or

---

<sup>27</sup> <https://www.inc.com/kimberly-weisul/women-venture-capital-ten-years-progress.html>

discussions. Whilst none had previously considered the possibility, several participants agreed that they would actively consider it.

In recent years there have been numerous studies around the translation of RI concerns, methods, or frameworks into industrial contexts – for example Lubberink et al., (2017), whose systematic review of previous work led them to recommend a systems-thinking mindset (as discussed in Chapter 5) that they believed would align more nearly with commercial concerns, and Schönherr et al., (2019), who argued that it was crucial to make the ‘business case’ for RI as a contributor to commercial success. However, van de Poel et al., (2020) suggest that many of these efforts are hindered by their method, which is usually to develop the RI tools and then seek to apply them in commercial contexts. In the view of these authors, one of the most significant constraints on the adoption of RI in industry is the perception that it is a primarily ‘academic’ concern, and an imposition from ‘outside’ that is not sufficiently bespoke for individual circumstances or the needs of a particular company (especially in the case of SMEs<sup>28</sup>). Van de Poel et al (2020) describe an approach that aimed to engage with commercial organisations, and work within and alongside them to assess the responsible-innovation-type activities they were already engaged in. Only after this did they then seek to expand these approaches – documenting the challenges and insights they encountered along the way. This is characteristic of a *de facto* RI approach, as discussed in Chapter 4. In particular, the authors distinguish between ‘ground-up’ activities taking place amongst staff, and ‘top-down’ frameworks originating with management, as seen previously in the discussion on institutionalisation. This analysis supports earlier work by Inzelt & Csonka, (2017) that showed *de facto* RI processes taking place in a variety of corporate enterprises. Indeed, literature and frameworks more directly focused on industry and commercial worlds – for example Porter, (1979) – bear many points of comparison with RI principles in their emphasis on considering context, stakeholders, reflection, and the appropriate response. It is important to note, however, that even a small SME of perhaps 50 employees (the smallest company examined within the van de Poel study) is still many times larger and more stable than a brand-new spinout with as few as four mostly-academic staff. The ability to adopt RI-type activities may be severely limited by restrictions on time, expertise, and resources.

It seems reasonable to suggest, however, that in terms of the actual processes involved, ‘institutionalisation’ and embedding may not be meaningfully different in academia and industry, despite the very different environments and priorities, and that some of the lessons being learned in the effort to embed RI into Centres for Doctoral Training could potentially be translatable into

---

<sup>28</sup> Defined as a business with fewer than 250 employees and turnover of less than EUR50 million  
<https://www.simplybusiness.co.uk/knowledge/articles/2021/05/what-is-an-sme/>

commercial contexts. These might include recognition of *de facto* responsible innovation, and the possibility of building out from existing work. The earlier discussion around the different ‘units of analysis’ (eg the individual, the group, the community, the institution, per Watkins & Marsick, 2019) will also be relevant as it is likely that a small number of employees all committed to responsible or ethical approaches will create a substantively different company than a larger group with more widely varying values.

#### 6.2.4: Research questions

This study therefore focuses on the creation of the new quantum computing sector, how its development is managed and driven, the stimuli that operate upon it, its enablers and constrainers – and the impact of these upon governance. As discussed in 6.1 above, many factors may encourage or inhibit the growth of a new area of the economy, including the availability of finance (and to whom it is available); knowledge transfer mechanisms; the availability of a skilled workforce; access to specialist facilities; and public readiness. Some of these elements are within the control of the policymaking, funding, and research elements of the ecosystem (for example via initiatives such as the Enterprise Investment Scheme, which offers tax breaks for venture capital investors<sup>29</sup>), but there may be significant areas where these are not active; not able to influence developments; or actors are unaware of the possible impacts of their decisions.

In particular there may be lacunae apparent around points of governance – this is of particular concern as a lack of attention to governance may create openings for societally undesirable or damaging developments (however unintended). The previous discussion on the forms and granularity of governance (Chapter 2) argued firstly that questions of governance – anticipatory or otherwise – are key in the growth of a new technology sector because of the ultimate requirement for public trust and societal embedding, and secondly that the difficulty of legislating around novel technologies means that these areas require adaptive, forward-looking forms of governance (such as responsible innovation). This leads to the research question:

*RQ1: What are the major shaping factors in creating a new sector, and where/how is governance manifested in this process?*

A supplementary research question will focus on the ways in which governance techniques may transition from one system to another – in this case from the university research environment to the commercial space:

---

<sup>29</sup> <https://www.gov.uk/guidance/venture-capital-schemes-apply-for-the-enterprise-investment-scheme#how-the-scheme-works>

*RQ1.1 How are governance challenges understood by entrepreneurs moving from academic to commercial contexts in novel technologies?*

This will be investigated through discussions with participants involved in many areas of the nascent quantum sector, including those who are already engaged in quantum companies, but also those in the wider ecosystem: for example professional services that recruit for the sector; sources of funding; consultants and professional organisations, and others. The university knowledge transfer and funding side may be of particular interest, because many of the diversity challenges seen in computer science fields may become amplified at the point of sourcing startup funding.

## 6.3: Methodology

### 6.3.1: Identification and recruitment of participants

In order to review the ways in which the commercial quantum computing sector is developing, a stakeholder-mapping exercise was carried out. This centred the unit of interest – a commercial quantum computing company – at the heart of the map, and then examined and investigated the areas of ecosystem with which it interacted, either directly or indirectly. Some stakeholders were identified through prior knowledge of the systems of spinout and new-venture funding and some through iterative analysis of the commercial context. Individual potential participants were identified through a combination of prior contacts, web-scraping and LinkedIn searches to find participants engaged in the sector in a variety of ways. Those identified through this route were approached directly. Discussions with interview participants also enabled ‘snowball’ sampling through suggesting additional participants.

The result of the stakeholder mapping exercise is included as a diagram below (Figure 4). This does not purport to be a complete map of the ecosystem, but serves to demonstrate the way in which stakeholders can be identified and categorised in terms of their relationship to the focal element. Should the focus change, the stakeholder map would change. Mapping the stakeholders in this way is instructive not least because it is (for this chosen unit of analysis) in many respects also a temporal map. In other words the elements closest to the spinout are those with which the spinout will need to engage in the earliest phases of its life. As it expands and develops, its circle of stakeholders also expands, with some remaining important throughout this process, while others are only critical in the early phases. From an RI perspective this also highlights the expanding scope of the externalities that may need to be considered through anticipatory work.

Stakeholders that were interviewed as participants in the study are highlighted in bold. Others were approached but were not responsive.

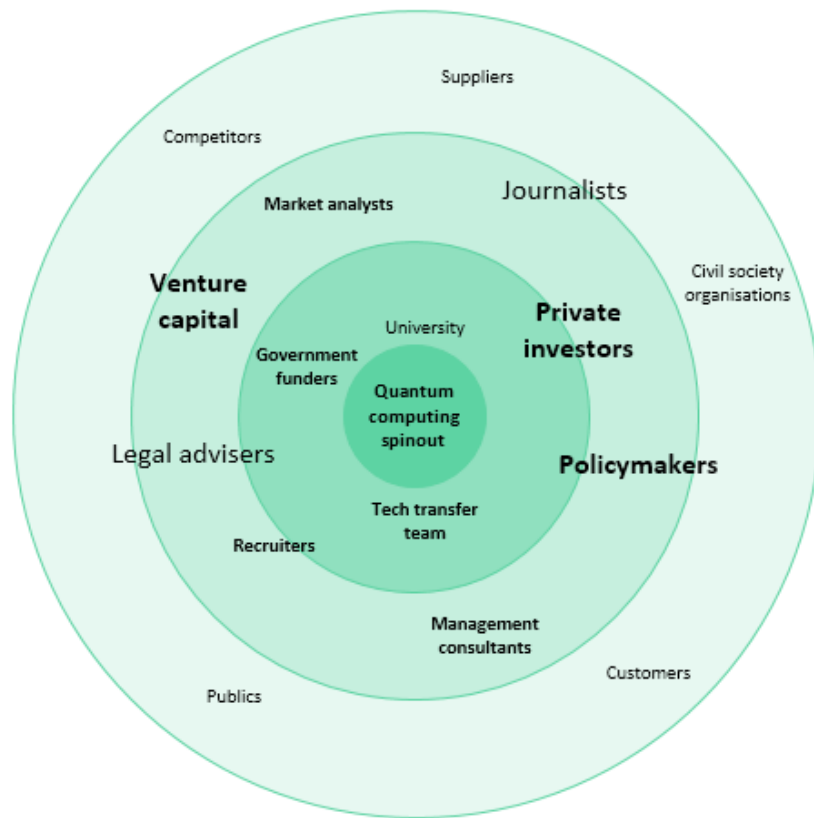


Figure 4: Stakeholder map for a quantum-computing spinout

As with previous studies, semi-structured interviews were identified as the most appropriate method, not least so as to align with those studies and enable potential cross-coding and comparison, but also because the field of research was relatively broad, including recruiters, funders, IP specialists, and others with many different specialisms, albeit a similar purpose. The semi-structured format allows for a wide-ranging discussion around common themes. Participants were asked questions pertinent to their areas of expertise, as well as more general questions about their thoughts on the sector, the enablers and constrainers they see, particular upcoming challenges, and the problems of governance.

### Participants

The table below gives details of the roles of participants and the codes assigned to them.

Role	Code
University technology transfer officer	I50
Professional adviser	I51
Specialist field recruiter	I52
University-linked venture funder	I53
Independent venture capital funding body	I54

Government innovation funder	I55
CEO of quantum computing company	I56
COO of quantum computing company	I57
Licencing specialist	I58

### 6.3.2: Analysis

In order to allow for comparison between studies and assess cross-study themes, the analysis for this study was carried out in the manner described in Chapters 4 and 5. Audio recordings were fully transcribed and corrected, before carrying out a deductive coding analysis in NVivo to draw out themes related to the research questions for this chapter and other points of interest. Interviews from the previous studies were also revisited in order to compare relevant comments from those participants on how the quantum sector might develop, or be developed, and how governance may be expressed. This includes indirect governance, and the operation of enabling and constraining mechanisms at a high level.

## 6.4: Findings

The findings from this set of interviews are shaped by the research questions and grouped broadly into the themes seen in earlier chapters around the relationships between investment, capacity, and ‘responsible’ approaches or other types of governance. In particular, some of the factors that were discussed by earlier interviewees as possible challenges are seen in much sharper focus in the commercial context, such as the problem of hype. As outlined earlier in the literature review (Chapter 2) and the landscape study (Chapter 3), and confirmed here, the speed of development in the last few years in quantum computing has accelerated beyond the predictions of even five years ago, and I57 points out, *“that scares a lot of people.”*

Many factors may be contributing to this acceleration, but it is highly probably that this rapid pace of development is at least partly attributable to the type and scale of investment in recent years – not just in terms of the size of the sums being made available through public and private funding, but also the manner of such investment, which has become structured in more flexible ways. This is a serendipitous development as these forms of investment are well-adapted for the particular needs of a growing quantum computing sector. However, as will be seen, there are significant implications for governance that arise from this.

### 6.4.1: Growth of investment and impact on development

One of the factors supporting the surge in investment seen in the diagrams in 6.2 *Background* was considered by participants to be a significant shift in the funding landscape in recent years. As discussed in 6.2, traditional venture capital funds have in the past had quite strictly defined windows for returns: *“there’s a certain timeline on which people in the private sector are used to funding things”* (I51). The Covid-19 pandemic, however, led to many funds re-examining their mechanisms as a result of so many ventures being impacted by the global slowdown in markets – the effect of this was that ventures with a 5-7 year timeframe to make returns found that extra time was needed. By 2022, participants described a significant increase in the number of ‘evergreen’ funds that can invest on an as-and-when needed basis and over a much longer timeframe than traditional venture funds.

*“it’s not necessarily something new full stop, although it is relatively new to the venture capital industry” I53*

This type of funding is ideal for quantum computing, and may prove to be a significant factor in the development of the sector. The availability of longer-term financing, and the willingness of investors to consider extended timeframes for investment, have added to existing factors around a surging UK private investment market.

#### Influence of private capital on the sector

The venture capital scene in the UK has shifted significantly in the last decade, with ‘megarounds’ of funding (defined as > £80m in one tranche) becoming a more frequent occurrence<sup>30</sup> and creating a more dynamic atmosphere in UK markets.

*“UK’s capital markets have been radically transformed. Now the UK is easily the biggest in Europe.” I50*

This is particularly pertinent for the growth of a quantum computing sector, because of the size of investment that is needed in order to make progress. Hardware, training, and infrastructure are expensive – investment is also needed to manage the capacity challenges that keep competition for talent (and therefore salaries) high.

New methods of structuring these large investments are also a strong driving factor – although ‘evergreen’ funds have been available for use as investment vehicles since the mid 2010s, the long-term element of such funds made them much more attractive under the circumstances created by the Covid-19 pandemic. One key factor about these funds is that, rather than supplying a large tranche of

---

<sup>30</sup> [https://uk.practicallaw.thomsonreuters.com/0-500-8350?contextData=\(sc.Default\)&transitionType=Default&firstPage=true](https://uk.practicallaw.thomsonreuters.com/0-500-8350?contextData=(sc.Default)&transitionType=Default&firstPage=true)

capital in one round of funding, money can be trickled in from the funder as and when needed by the fledgling company – and, crucially, permits for a much longer timeframe for making returns.

*“we have the ability to ... hold our positions in those companies pretty much indefinitely”*  
153

As 153 notes, these new structures do not follow the same pattern as traditional VC investments, discussed in 6.2 *Background* – there is no longer a defined ‘exit window’. This is quite different from the usual, relatively short-cycle nature of VC investment. It is probable that these new, longer term structures reduce the pressure on entrepreneurs to make rapid returns, which may be particularly valuable in a field such as quantum computing where there remain such enormous practical and engineering challenges, and hitting previously-defined technical milestones in development work may simply not be possible.

*“it makes a huge amount of sense ... if you're trying to invest in things that can be massively world changing but realistically you're going to need 10 or 12 years”* 153

This need for patience is crucial in a developing marketplace with such high levels of uncertainty and many remaining ‘unknown unknowns’. There is a requirement for investors to be extremely comfortable with risk, to be sufficiently financially cushioned, and to have a long-term perspective that encompasses the extended timeframes. But many consider the high risk and long wait for returns to be worthwhile.

*“[with] quantum computing ... you have to take a different lens which is like; the market size – it's big. It's really, really big. But the gamble will be big.”* 157

Participants were also quick to point out that these large, steady, and enduring supplies of capital need to be available on a continuing basis, and that without such funding measurable progress in quantum computing cannot be achieved. Some participants take object lessons from other technologies that are perceived to have been persistently under-funded, such as fusion, and have thus not delivered the level of progress that was expected or hoped for.

*“We've never funded it... [so] we have never made the progress that we need”* 151

Whether or not such funding would have resulted in viable fusion technologies, participants are keen to ensure that this does not happen in quantum computing, and perceive that the newly flexible and well-financed investment ecosystem in the UK is in a good position to support technologies such as quantum in the long term. Some of these changes in the ecosystem are also attributed to incentives and tax breaks around investment, *“the EIS<sup>31</sup> and the SEIS<sup>32</sup> ... those are game changers”* (150),

---

<sup>31</sup> Enterprise Investment Scheme [https://en.wikipedia.org/wiki/Enterprise\\_Investment\\_Scheme](https://en.wikipedia.org/wiki/Enterprise_Investment_Scheme)

<sup>32</sup> Seed Enterprise Investment Scheme [https://en.wikipedia.org/wiki/Seed\\_Enterprise\\_Investment\\_Scheme](https://en.wikipedia.org/wiki/Seed_Enterprise_Investment_Scheme)

meaning that there are more readily available, larger, and longer lasting sources of funds for supporting startups.

Such a well-funded group of private investors in turn creates what participants characterised as a virtuous circle of increased profits available for onward investment. In terms of international marketplaces, which compete to host the largest companies on their financial markets, this is viewed by many as a beneficial development.

*“there's also a kind of flywheel effect, where early traction ... increase[s] the visibility of the UK as a place for investment” 150*

It is useful to note at this point that there are substantial overlaps here with policy and any impacts on UK economic security that may be perceived by government. As seen in earlier discussions around the Fusion Doctrine and the National Security & Investment Act 2021, at this level of investment, governments may involve itself in trying to actively shape the sector.

This is not to say, however, that it is now an easy task to persuade private investors – even those accustomed to higher risk investments – to invest in a management team and its quantum computing startup. The early days of a high-risk technology can be offputting even for bullish investors, and companies that are now seen as successful still experienced many struggles in their startup days. One CEO commented that if his company had had to raise funding in its early days, the new enterprise would have foundered.

Although there is substantial capital available at present, even private capital's resources are not bottomless. Investors commented that whatever 'hedging' decisions they may have taken in early investment rounds – choosing to invest in multiple companies working on different versions of hardware, or types of qubit, for example – such broad approaches are necessarily time-limited.

*“as ... companies get more capital-intensive, you're starting to see people are forced or incentivized into picking certain approaches ... and then focusing their money there” 153*

This points up a potentially challenging factor for startup and spinout founders. Investors are clear that, despite the long-term and flexible facility offered by 'evergreen' approaches, or the ready availability of capital, their ultimate goal is that they are looking for returns. Inevitably this will mean some companies are seen as unviable – which may not be a reflection on the quality of the technology or the individuals, simply the commercial realities. Some potentially valuable enterprises may therefore be cut off from funding at a crucial point, and those that do receive funding are likely to experience pressure from VC investors to focus on activities that lead to returns in the shortest timeframe. This may include *not* devoting time and resource to activities that do not directly benefit the bottom line, such as anticipatory governance, stakeholder inclusion, or ethical approaches to

doing business – the possible impacts on governance here are clear. Compliance with legal obligations – for example formal accounting procedures – is necessary, but internal governance approaches that are not a legal requirement may be regarded by such investors as superfluous.

The factors mentioned above are likely to both create gaps in investment and – as a contingent effect – in governance approaches, and a national government looking to grow a societally responsible quantum computing sector may need to consider that these lacunae could potentially be filled by the public sector.

### Public funding – filling the gaps?

As seen in the landscape study (Chapter 3), the desirability of public sector financial support is acknowledged across the board. Several factors that contribute to this desirability were suggested by participants in this study, including the need for government to be seen to be supporting a sector where it has made strong claims for UK technologies; the ability for government to make significant and steady funding available over the long-term; and the likelihood that government may need to be seen to be an early adopter (as indeed has proven to be the case<sup>33</sup>). All these elements lead to wide acceptance of the need for broad and deep government support, as I51 acknowledges.

*“if you want ... an economy that benefits from these technologies, then a government would want to fund that in their sphere of influence.” I51*

The UK government has repeatedly stated its commitment to the development of both the foundational science of quantum computing, and a commercial quantum computing sector<sup>34</sup> and as outlined in the discussion on private funding, study participants viewed the high risk-profile characteristic of young quantum startups as an obvious rationale for making public funding available. It has also been argued in earlier chapters that public funding can provide the opportunity to attach governance and oversight mechanisms to novel technologies, and that it is this aspect that makes public funding especially attractive from a ‘responsibility’ standpoint.

However, the real challenge for making public sector funding available is the question of risk. Risk is part and parcel of innovation – indeed is a constituent element – but as seen in both the policy and the landscape studies, governments are very risk-averse institutions that are not comfortable with the likely repercussions of losing public money (for which there must, and should, be accountability), or with having to try and ‘pick winners’ by choosing where to invest. The pressure to not be seen to be utilising their investment powers poorly inevitably leads to a high level of risk-aversion. This can be

---

<sup>33</sup> <https://www.gov.uk/government/news/government-backs-uks-first-quantum-computer>

<sup>34</sup> Eg <https://www.gov.uk/government/news/government-backs-uks-first-quantum-computer> and <https://www.gov.uk/government/news/70-million-funding-to-secure-uk-position-as-a-world-leader-in-quantum-technology>

particularly acute in a technology such as quantum computing where there are still many different possibilities for hardware development. Governments and policymakers are therefore highly likely to be viewing opportunities through a more cautious lens than private sector investors.

*“you will always be able to find a risk or something that could go wrong, or a way in which this could fail” I53*

Some participants regarded these careful approaches as a significant inhibiting factor on the development of the sector. There were doubts expressed about the abilities of public servants to take the rapid and potentially high-risk decisions widely regarded as necessary to make progress, and therefore the question of whether the UK government’s Innovation Strategy can be effective is regarded as an open one.

*“it's not because people aren't well-meaning, they just don't know how to do it.” I57*

Other participants also saw structural issues in the longer term in wider UK markets, particularly as there is not a good track record in the UK of facilitating longer-term growth – *“if the government wants to build a British Intel ... they need to stop companies exiting when they hit £100 million.”* (I58). Despite these reservations, however, some of the government’s initiatives were widely seen as timely and appropriate, and regarded in a positive light by young companies and investors.

*“The Innovate UK pot of funding is one of the best things that we have in the UK” I52*

As well as the factors mentioned above and the impact on governance, public funders are in a strong position to provide leadership on the ‘systems thinking’ needs that present themselves, especially in terms of managing the integrated processes that can support the expansion of infrastructure and capacity for the sector.

### 6.4.2: Capacity challenges

The capacity question is one that creates some of the most obvious challenges for the development of any new sector but also has implications for governance and responsible innovation approaches. Growing capacity in a sector can lead to engagement with wider groups – this has the potential to support innovation; increase transparency; widen public understanding; expand diversity; and create opportunities for dialogue – therefore efforts to increase resilience and capacity can also support responsible innovation and anticipatory governance.

As seen in the landscape study, one of the biggest problems for sectoral development is that of building capacity at pace, because of the need for deep expertise: *“if you need someone that is an expert in something, we're still looking for PhDs”* (I52). This a challenge partly because the system of funding for new PhDs in the UK may take up to seven years before those new postdoctoral

researchers are moving into the job market, and globally, universities are able to supply only one-third of the current demand (McKinsey & Company, 2022). Although there is UK government support for postgraduate research, there are concerns that undergraduates commence their studies at university with no knowledge of quantum – participants in the Chapter 3 study suggested that the UK’s National Curriculum should be amended to reflect the needs of the sector. Interestingly, quantum-sector recruiters were less focused on the need for doctoral students, understanding that a developing sector will require more than just experts in quantum physics. Some actors were taking pragmatic action to address the shortage of qualified personnel by trying to bring potential employees rapidly up to speed without the benefits of a postgraduate qualification. However, they also acknowledged that familiarising employees with the intricacies of quantum mechanics is not an easy thing to accomplish in a rapid timeframe.

*“if we are talking about people that are completely outside quantum and they need to come to quantum, quantum is hard.” I52*

Another method for rapid upskilling is via hands-on exposure to the business environment, which in some areas of the world (particularly the US) is enabled through internships for doctoral students. These are often popular with companies because the business is able to pay a highly trained, and motivated, short-term employee to look at a particular challenge that the company itself may not have time, capacity, or expertise for.

These capacity challenges are clearly not limited to the UK – they are seen in all countries with quantum programmes, and in all areas of research and development, creating a highly competitive environment for hiring. Even after staff have been recruited, retention may be a challenge while there is such intense competition for talent: *“they are just leaving the UK”* (I52). In particular there may be significant salary differentials between different countries, with US companies having a reputation for paying particularly highly. Unsurprisingly, given that higher salaries are such a factor in recruiting the necessary employees, commercial participants see clear links between the availability of investment, and their ability to solve their capacity challenges.

The ongoing expertise shortages – *“the talent pool is so small!”* (I52) – also create a tension between the desire of national governments to train and retain their own specialists, and the huge technical challenges remaining in the field, which interviewees insisted can only be resolved through international partnerships. The instinct towards protectionism goes against the flow of practice in quantum computing fields, which have traditionally been highly collaborative and international (eg 3.4.2).

*“the limiter right now is the number of people working on problems ... it needs more collaboration” I51*

Within the UK there is also the ongoing impact of Brexit, which is not only limiting the potential talent pool to the UK but also heavily constraining the ability of UK companies to hire staff from overseas – recruiters were particularly aware of the challenges here.

However, it is possible that these capacity constraints may create an environment in which recruiters, trainers and investors are impelled to think more creatively about how to resolve them,

### Capacity and inclusion

As discussed in 6.2 *Background*, questions of capacity should also take account of questions of diversity – as has been shown, quantum physics as a field already has a low proportion of women, for example, and venture capital funding is also more difficult for women to obtain (Bittner & Lau, 2021). This is not least because venture funds are themselves overwhelmingly male-run: *“these are very male environments ... that’s really problematic”* (I50). The intersection of these factors may create multiple pressures for women academics when they are seeking to spin out their own businesses – they are not only already in a male-dominated field, but may find their options restricted, can experience pressure to take particular directions (for example moving into administrative rather than research or technical roles), or will sometimes opt to leave the field altogether. One participant commented on how frequently they had witnessed women moving sideways out of technical and management-track roles, and questioned whether this is always their choice – these pressures will also be experienced by Black and minority academics.

However, there is awareness that this is contributing to capacity shortages (not solely in the quantum sector), and in some areas there are efforts to address it, for example with entrepreneurship programmes for women that can support them in raising investment<sup>35</sup>. There are clear opportunities here to shape the commercial sector in more equitable directions, and there may be the requisite willingness to do so (participants were very aware that current approaches are a long way from ‘enough’).

*“we worked ... really hard on getting that intake... to 50/50. And that had a huge impact on the women involved in the programmes” I54*

Increasing the diversity of stakeholders in the sector would help support responsible innovation approaches for the reasons discussed above. There is also a requirement in these early stages of market-building for individuals with multiple skillsets who can operate in several modes.

*“they want a person that knows about quantum, that has industry experience, business experience, and can work with customers, and good luck finding this kind of person.” I52*

---

<sup>35</sup> <https://quantumvantage.co.uk/quantumwoman>

The capacity/investment issue is especially acute here, because not only are the ‘right’ people rare – such generalist skills being largely contrary to the ever-more-specialist training that doctoral students receive<sup>36</sup> – but they will understandably take their skills to the company that can afford to pay the most. A company’s ability to offer good salaries is, as noted previously, inextricably tied to the ‘investment’ question, especially as there are few saleable products and only a few companies have revenue streams. The smaller the company, the greater the need for investment – this is vital for maintaining the flow of spinouts.

### 6.4.3: Spinning out: academia to industry

The nature of the technology, grounded as it is in complex mathematics and quantum physics, dictates that the genesis of quantum computing startups is most usually academic, and that – as the landscape study showed – academics frequently become involved in commercial startups. One interviewee pointed out that more than half the academic attendees at Google’s first quantum computing conference in 2016 now work at startups.

Tech transfer specialists and venture capital organisations have perceived a general increase in the number of spinouts in all sectors, attributing it partly to the availability of investment, as discussed in 6.4.1. However, they also see increased entrepreneurialism from academics. There may be a variety of motivating factors behind this shift, with one interviewee speculating that *“academics ... want to see their science applied and have some positive impact on society.”* (I58). In the UK, this may also be spurred by the Research Excellence Framework’s emphasis on ‘impact’ from research, and a focus on the calculation of a university’s ‘contribution’ to the UK economy<sup>37</sup>.

*“The wheel is slowly turning to students being excited and interested in the startup world as a theme ... all the universities are working on having an entrepreneurial mindset” I54*

This expansion in spinout numbers may also be influenced by wider societal shifts that have come to regard starting a business as a positive and relatively common endeavour (evidenced by the popularity of reality TV formats such as *Dragons’ Den*<sup>38</sup> – now in its 19<sup>th</sup> year).

*“what’s good is that entrepreneurship as a kind of cultural idea has seeped into researchers now because it’s ... much more common in general in the UK” I50*

This new enthusiasm for spinning out into commercial environments was welcomed by tech transfer participants, but it was also acknowledged that academia and industry are not necessarily comfortable

---

<sup>36</sup> See also the question of T-shaped expertise in 2.3

<sup>37</sup> Eg <https://www.ox.ac.uk/research/recognition/economic-impact> and <https://www.theguardian.com/education/2022/jun/27/university-college-london-generates-10bn-a-year-for-uk-says-report>

<sup>38</sup> <https://www.bbc.co.uk/programmes/b006vq92>

or obvious partners. There are many challenges, including the perception that researchers and businesspeople “*speak a completely different language*” (I52). There are also frequently expectations from academics that they will be able to run or be involved with a startup business on a part-time basis while continuing their academic career – whereas investors by contrast will be expecting them to put all of their time and effort into making the business a success. In addition there are particular skills required by these environments that may conflict with years of academic training, or simply be entirely foreign to their experience.

*“the skills that make you a ... good academic ... don't make you a good entrepreneur” I50*

There is clearly a need for academic research as the basis of the innovation ‘pipeline’ – “*without founders, we don't have the spinouts*” (I58) – but many academics lack the management skills that are required to run a successful business. Investors are very aware that – especially within the University spinout ecosystem – it is not necessarily the best idea or the most innovative piece of research that goes on to create the most successful commercial enterprise.

*“The technology is important, but it's really about ‘do we believe in this team?’” I54*

Instead, success is much more related to the individuals, particularly in terms of how “coachable” they are in learning new skills, and how well they are able to adapt to new environments and requirements. The willingness to learn is key – the experience of being expert in a given academic area can in some cases negatively affect the ability to accept and undertake the learning-curve associated with starting up a commercial business venture.

*“the best founders are the ones who know what they're not good at” I54*

One component of this shortage of business skills is often a lack of market awareness and analysis, including a degree of naivety about what will or will not prove to be a commercial idea (this applies to other sectors as well as quantum computing). As I52 points out, “*not everything that you do in the lab can be a business*” – academics may not understand the difficulty of turning a piece of research into an investment-worthy and profitable enterprise.

*“something that worked once in a lab isn't an investable business proposition” I58*

Academics are also frequently unprepared for the challenges of scaling up their work and there can be a ‘translation gap’ between the proved science and the commercial readiness. There is often an underestimation of how difficult the remaining challenges are, with a tendency to believe, “*Oh, I've made it work once, the rest is just engineering*” (I58). Although tech transfer teams in general welcome the rise in academic entrepreneurialism, they also recognise that there is a balance to be struck between encouraging the trend, and acknowledging that it is not for everybody.

*“you don't want to push anyone who doesn't really want to be an entrepreneur to be an entrepreneur ... it's like pushing water uphill” i50*

However, despite this entrepreneurialistic trend, and the increase in university-originated ventures in recent years, the system of ventures spinning out from UK universities is not necessarily viewed entirely positively in wider venture capital marketplaces, as universities often require a large share of the equity at spinout stage. This can make the business an unattractive proposition for later-stage investors as the available equity becomes more dilute.

*“I have seen insane cap<sup>39</sup> tables from companies that have spun out of universities in the UK ... this destroys your viability from a venture funding standpoint” i51*

Given the dependence on university spinouts for populating the sector, this may act as a constraining factor on the development of a commercial quantum computing sector. Indeed the June 2022 analysis of the sector from McKinsey showed a slowdown in spinout activity, with one hypothesis being that all the available academics were already involved in spinouts (McKinsey & Company, 2022). Relying on the UK's university sector for the creation of young quantum computing or quantum technology companies is also constrained by the fact that there are only a few UK universities with the resources to produce a flow of spinouts. Spinning out is also frequently a lengthy process – the factors mentioned above around a lack of commercial awareness, shortage of management experience, and translational challenges with the technology itself, may combine to make it a relatively drawn-out procedure.

*“‘Agile’ is probably not a word I'd use. ...In six to eight weeks, you could do a spinout if you had to ... Six to eight years has certainly happened.” i58*

It is clear that there are numerous pressures here – associated with a drive to REF-style ‘impact’; difficulty in repeating or scaling-up the product or service; a lack of skills or marketplace know-how – and these pressures are unlikely to be compatible with prioritising questions of governance. The transition from a university environment to a commercial one also brings with it a vast alteration in the applicable governance ecosystem.

#### 6.4.4: Governance in commercial quantum computing

These questions of governance were key to the research focus of this study, and it has been shown how the longstanding themes of all the case studies around investment and capacity have impacts on the way governance can be shaped. When it comes to the day-to-day operations of new companies, there is often a lack of awareness that within the commercial environment there will be no equivalent of

---

<sup>39</sup> A ‘capitalisation table’ shows the breakdown of equity ownership in a company  
<https://www.investopedia.com/terms/c/capitalization-table.asp>

university or funder infrastructure holding academic founders to account on governance matters. Instead, it is now the responsibility of the founders to research and engage with the requisite governance and compliance matters. This lesson is not always learned in time.

*“one in four ...enterprise technology companies fail for avoidable operational reasons. Governance is one of the key limbs of that.” I54*

In particular, there may be no incentive for companies to bring with them any ethical and responsible approaches that had been inculcated or mandated within the research environment.

### Anticipatory governance - translation of RI

*“our mission is to help create a quantum industry that benefits the most people” I51*

The institutionalisation of RI, or any anticipatory governance process, faces its strongest test at the point at which research moves from the university environment into the commercial. The question of whether ethical considerations and anticipatory governance translate from a research lab into the commercial context was not one that any participants had previously considered. There *are* significant governance mechanisms in place during the spinning out of a company, and compliance with these is overseen by technology transfer teams, but these are focused on, for example, the attribution or distribution of intellectual property rights in a given technological innovation. They are not directed towards protection of society, or considerations of appropriate technological ‘fit’.

*“the university has got a lot of process and structure around spinouts that exists ... as reputational protection for the university.” I58*

Instead, spinouts are left to their own devices in terms of the appropriateness of their innovation to society, and to the moral choices of their founders. Participants in this study were doubtful whether these generally influence the development of a company.

*“any kind of personal ethics code and ... training, I don't think fundamentally makes a huge difference. You need regulation.” I50*

There is no equivalent at the point of spinning out of the type of ‘responsible innovation’ requirements that are seen earlier in the pipeline, whether these are embedded in funding calls or stipulated by an institution. Equally there is no governance or oversight of whether technologies might be harmful to society. The decisions are fundamentally commercial, as this is what is rewarded – as I50 points out “[a spinout]’s incentives are to follow the money”. In new marketplaces, there are likely to be few, if any, industry standards or consumer protections, rendering the need for anticipatory governance even more acute.

The question of governance is at the heart of this study, not only because governance is such a strong feature within the research-funding landscape that is the source of most commercial quantum-

computing companies, but also because the point at which a technology translates from the lab into the commercial sector entails a wholesale shift from one governance ecosystem into another. The question of how any responsible innovation-type concerns translate into commercial sectors is therefore a key one.

*“being in a university ecosystem versus being a startup ecosystem, there's ... different metrics ... in terms of how you're measured, how you're incentivized” 154*

These commercial incentives, allied with pressure to perform commercially and to make the necessary returns to investors, may create a temptation to overpromise, or exaggerate progress that has been made.

### Marketplaces and the problem of hype

In a mature marketplace, business development has certain well-defined challenges: market segmentation; dealing with competitors; regulatory compliance; and so on. When building a commercial sector from the ground up, there is as yet no demand for the products that need to be sold, and – in the current instance – no product to sell. Marketplaces do not really exist yet.

*“it's been propped up by by venture capital and by government” 150*

In such an uncertain and novel area, there is a need for vision, and the ability to commit for the long term – but such visions can easily be misinterpreted.

*“we're sort of bouncing around on the hype curve with quantum at the moment” 153*

The tensions that exist around hype have the potential to be extremely challenging – as 152 points out, *“there is a lot of misinformation”*. Much of this misinformation exists in popular culture and elsewhere, but many engaged with or within the quantum computing sector believe they have a responsibility to ensure that they are managing expectations, because of an understanding that *“if the bubble bursts we are left with nothing”* (152). It is therefore in everyone's interest to right-size understandings of what is possible both immediately and in the near future.

*“we've staked a lot of our reputation ... to be as hype free as we possibly can” 157*

However, such intra-industry discussions – “expert-to-expert” conversations – may not be resonating outside of these well-informed groups. Within knowledgeable circles, hype is perhaps less of an issue, but outside of these expert discussions, there is huge potential for unrealistic expectations to skew marketplaces, policy and other shaping factors. In the wider economy, there are hopes and beliefs being created within mainstream and social media that may not reflect reality. It is clear that at a distance from the direct development of the technology, there is less understanding of the number and scale of the challenges remaining.

*“you have some other people ... saying ... next year, we can have a breakthrough in quantum!” 152*

The impact on anticipatory governance here is that there may be an obligation for those within expert quantum computing circles to engage with media, with publics, with industry, and with policy, to create greater understanding of the possibilities, opportunities, and likely timeframes of quantum computing. Participants in this study already perceive overheating in markets.

*“should we have companies that are worth billions already, when they probably don't even have 10 qubits that can communicate with each other?? It's crazy.” 152*

In particular there may be a requirement for those responsible for the shaping of marketplaces, such as policymakers, to have in-depth understanding of the affordances, challenges, and needs of the quantum computing sector, in order that they can convey appropriate messages to communities and to media. There is therefore a need for communication at many levels.

*“We spend a lot of time educating different policymakers in multiple countries” 157*

This need for clearer understanding about the progress of the technology and the sector is a necessary part of governance. It is apparent that simply providing investment and a supportive environment is not sufficient – the problem of hype is that if the development of quantum computing is perceived (rightly or wrongly) to be a ‘bubble’, there may be a significant cool-off period that will slow progress, and a reputation for overpromising will create breaches in public trust. There is already significant concern around this, as some voices express caution about huge company valuations unsupported by assets or revenues (Gourianov, 2022), and others dubious as to whether a ‘quantum computer’ will ever be built (Dyakonov, 2020). It is therefore necessary for the quantum computing sector to engage more closely with governance – this is likely to also involve engaging with the policy shaping the commercial environment.

### **Legislation, regulation, and working with policymakers**

One of the challenges seen in the policy study – that those charged with policymaking in a given area may not have sufficient expertise in relation to the object of that policymaking – is also seen here from the external point of view. In particular, those working on the development of commercial quantum computing are very aware that there is likely to be limited understanding of quantum available within policymaking circles. Consequently, policymakers may lack the tools with which to make informed judgements about questions of hype, capacity, infrastructure, security, governance, funding, and other questions.

*“there will be people who are sitting in government offices, who are incapable of making judgements, but who are allocating capital and resources” 156*

There are also very few knowledgeable voices from the commercial sector advising policymakers. Participants who had considered these questions perceived that they were a long distance from policy and that there is *“a disconnect between policy and what people are doing on the ground.”* (I58)

For those that *are* consulted, there is a clear sense of responsibility in being a ‘trusted adviser’ to policymakers, alongside an awareness that this is very much a privileged position, available only to those with good connections and the resource to nurture them.

*“I think on quantum we probably should do more ... engagement with the relevant parts of the government on what they're thinking.”*153

This reliance of ‘rule-makers’ on potentially a very few ‘rule-takers’ for their information is not necessarily the healthiest for robust policymaking, and there may be a danger of capture of the governance space by a small number of corporate actors.

It is also the case that engaging with policymakers requires a particular combination of sectoral knowledge, understanding of government, and communications skills – a very different skillset from being a quantum physicist, an engineer, or a CEO – and it is important for companies to understand the broader concerns that may be shaping the policy agenda. The work of developing new companies and a new sector from scratch means that there may be little time or scope to focus on the work of engaging with policymakers

*“the people who are building this stuff, they're spending time building it ... [not] in policy conversations arguing that this is a very scary thing that needs to get regulated”* 151

In particular the drafting and passing of the National Security & Investment Act is regarded within the quantum computing business community as an example of an unmindfully low level of engagement. Affected groups – including tech-transfer teams, venture capitalists, and entrepreneurs – were taken by surprise as it passed seemingly without discussion with those who would be impacted by its measures.

*“they may well have consulted with people, but they didn't consult with the people who are actually doing it...”* 158“

However, as has been previously discussed, quantum computing is viewed as a ‘strategic capability’ technology, and one that numerous nation-states believe it is imperative they have access to, for reasons that include national security, prosperity, and prestige.

## 6.4.5: National concerns and protectionist approaches

The connection between security and prosperity for certain technologies, including quantum and AI, that has been set out by the UK government in the Fusion Doctrine (Cabinet Office, 2018), is here demonstrated in the perception that countries need to invest heavily to remain ‘in the race’.

*“The ‘scary thing’ isn’t that you’re not... the winner. The scary thing is you’re not included.”  
I51*

As a result of these perceptions, the nationalistic and protectionist concerns discussed earlier in Chapter 3 and in more detail in Chapter 5 are already being experienced in the developing commercial environment – with a sense of frustration but also with recognition that tensions around questions of economic advantage and national strategy are connected to geopolitical sea-changes.

*“this will be an ongoing issue. Particularly with the fear about technology transfer to ... countries that are considered strategically risky” I50*

Quantum computing companies are emphatic about their need to continue to collaborate across international boundaries, not least because the work itself is very difficult. However, this need is in tension with the evident desire of nation-states to protect what they see as ‘their own’. Companies are rightly concerned that the trend towards ‘national’ concerns will negatively affect their ability to communicate and co-operate with colleagues in other countries. This may be especially true in a technological field that, traditionally, has been extremely open and cross-border.

*“if you’re going to start slicing and dicing along national lines ... you’re just going to slow everything down.” I51*

The companies that are able to engage with policymaking find themselves having to persistently lobby to be able to keep the channels of communication open across national boundaries, arguing repeatedly that barriers to collaboration are necessarily barriers to progress; *“the risk is that we strangle the potential by closing things down”* (I51). This is an additional knock-on effect of the concern seen in 6.4.4 around policymakers’ poor understanding of the technology and the underlying science – if policymakers do not fully grasp the scale of the technical challenges that remain, they are unlikely to be sympathetic to appeals for international co-operation.

*“it’s just a constant education that says ‘... we need to be able to collaborate to advance. You ...put this protection wall up? then it all just goes to zero. No progress gets made.’” I57*

There is therefore considerable wariness within the sector around the increase in, and implications of, protectionist approaches. The National Security & Investment Act, discussed above in terms of how policymakers consulted (or did not consult) with sectoral participants, is considered to be highly likely to impact investment – if not the overall quantity then certainly the speed at which investments can be made. There was widespread unease amongst the tech-transfer and commercial participants about

the Act's implications, and its possible effect on inward investment, in particular where these participants are already operating in risk-averse ecosystems that are likely to view 'notifying' an investment under the Act with a 'better safe than sorry' approach.

*"no university has fully grappled with what the full implications of that are" 150*

The UK is of course not alone in taking steps such as this to try and protect its interests, although the degree and type of protectionism varies across the countries that are working on quantum computing. Companies in the US that are operating in sectors seen as highly significant from a national security standpoint may be subject to an individualised National Security Agreement<sup>40</sup>, which, although it *"adds legal overhead and time"* (151), may enable more tailored approaches than the broad-brush approach of the NS&I Act

It is by no means certain that the NS&I Act is part of a well thought through approach – *"you don't get a clear sense that they really have an idea of what they want"* (150) – and seems to be running counter to what many in the field see as the necessary and inevitable direction of travel in terms of collaboration. For example the UK and US governments have agreed a Joint Statement on collaboration in quantum<sup>41</sup> (though it remains to be seen what the practical outcomes will be), and UK funders are encouraging international collaboration through summer schools and partnerships<sup>42</sup>. There does not necessarily appear to be a coherent governance approach here.

Participants were also concerned about other effects that may stem from these impulses towards 'national' approaches. Such mindsets do not only result in protectionist regulation, they may also shape the motivations for national governments to invest in particular companies, or make decisions on which companies may qualify for ongoing support. Governments may also use their procurement power to support 'local' quantum companies – for example the UK government's purchase of a quantum computer from Orca in 2022.<sup>43</sup>

*"People will make judgments not because [company A] is better than ... [company B], but because [company B] is [their own nationality]" 156*

There were also concerns from a few participants that government sponsorship may reduce effective competition, allowing poor technology or poorly-run companies to survive – because government

---

<sup>40</sup> See <https://www.justice.gov/opa/press-release/file/1457291/download> for an example

<sup>41</sup> <https://www.gov.uk/government/publications/uk-us-joint-statement-on-cooperation-in-quantum-information-sciences-and-technologies/joint-statement-of-the-united-kingdom-of-great-britain-and-northern-ireland-and-the-united-states-of-america-on-cooperation-in-quantum-information-sci>

<sup>42</sup> <https://www.ukri.org/opportunity/host-an-international-summer-school-in-quantum-technologies/>

<sup>43</sup> <https://www.bbc.co.uk/news/technology-61647134>

investment isn't in the business of 'picking winners' – and that this will have the knock-on effect of skewing markets.

Government also has a strong indirect influence in the initial shaping of the technology pipeline – this is perceived later on in the process by venture capital investors, who are aware that the opportunities they see at startup stage have essentially already been selected for at the research-funding stage. As seen previously, the funding stage is itself shaped by multiple factors, and given that the source of funding for much quantum research in the UK is EPSRC, these factors may include the incorporation of responsible innovation approaches within these projects.

*“we see things in our pipeline that have been germinated as a result of either government funding, ... or the government's views on which areas are particularly attractive. So I think it directly influences the pipeline of opportunities” 153*

This study demonstrates extremely clearly the tight linkage between the triple-helix of policy, industry and research in quantum technologies and quantum computing, and how the interrelationships between the strands mean that decisions within one strand are quickly felt elsewhere, sometimes with unpredictable effects.

## 6.5: Discussion

There are several sources of tension evident in this study, many of them concerned with the different varieties of governance in play at different levels. As discussed at the outset, the research question for this study aimed to understand the major shaping factors in creating a new sector, and (given the importance of governance in public confidence around novel technologies) where and how governance was manifested, or where concerns or gaps could be seen in the process. It was also seen as crucial to understand how governance challenges were actually understood and experienced by entrepreneurs during the translation from the academic to the commercial context. This study makes clear that at all scales – from the difficulties being caused by legislation (including the UK's NS&I Act), down to the individual moral choices being made as companies spin out and have to decide what guiding principles they will carry across from academia – there are governance gaps apparent at every point.

### 6.5.1: Governance gaps

It has already been argued that governance in novel technologies is crucial for several reasons. Focusing in this study on the translation of governance at the point of spinning out from the university environment demonstrates that there are serious gaps here that will undoubtedly affect the way the technology – and by extension the sector – is shaped. Almost the only governance-based thinking seen

at the point of spinning out is focused on satisfying the intellectual property concerns of the university and the investors, and compliance with any fiscal responsibilities that may be incurred. Governance of the technology that the new company is seeking to commercialise is considered to be outside the scope of university, tech-transfer, or investor responsibility.

## Investment

As seen above in 6.4.1, there is much attention directed towards both public and private funding for startup companies. In general this is in what may be called a ‘positive’ investment mode, centred around the availability of investment, for example, or the receipt of grant funding – ie a direct injection of capital whether on a once-off or ongoing basis. Such funds, which appear on the positive side of the balance sheet, may be drawn from many possible sources; whether a research council, a venture capital fund, or a government Catapult. Investments from public sector funding are frequently attended by significant due-diligence requirements, reporting, oversight, and the likelihood of future audit on several metrics. However, investment from a venture capital fund is more likely to be directed solely towards commercial ends, creating a governance imbalance.

Additionally, there is what could be termed a ‘negative’ investment mode – the creation of a situation whereby a company has to expend less money or resource to achieve its goals. This affects the other side of the balance sheet, in the form of reduced liabilities, and might take the form of tax breaks, or infrastructure support for which the company may now pay less. In general this form of support would come from the public sector as it would be structured by policymakers. Unlike funds received directly as investment, the availability of support such as this is not generally attended by similar governance and reporting requirements. This creates another governance gap, because there is significantly less oversight of benefits that reduce a company’s liabilities.

Gaps in investment oversight also carry with them assessment gaps, similar to those seen in the assessment of policymaking in Chapter 5 – a shortage of general analysis of effects produced and a lack of investigation as to whether they have achieved objectives other than the strictly commercial. These assessment gaps are also seen in section 6.4.5 in terms of assessing an innovation for either its utility or its potentially damaging effects on society. The problems of hype and overpromising are also linked here.

## Capacity

In every field study, the problems of various types of skills shortages have been seen both in research phases and commercial endeavours – however, the impact of such shortages in terms of governance may not be immediately apparent. Nonetheless, considering these impacts through a responsible innovation lens shows that there will inevitably be effects on an organisation’s ability to embed

anticipatory governance techniques. As has been shown in 6.2.2, diversity (of experience and thought as well as background) is a necessary factor in being able to draw on a range of views and a depth of interdisciplinary expertise – crucial resources in terms of working with stakeholders to anticipate, reflect on, and respond to, possibly poor outcomes. A lack of capacity not only overloads those who are working in a company or project but reduces the likelihood that the required type of expertise will be incorporated or operationalised. The present capacity shortages also represent an opportunity, however, should they be recognised as a way to draw on broader resources than have previously been considered.

## 6.5.2: Promises and realities

The tension here is in the difficult line that must be walked between encouraging eagerness around the possibilities of quantum computing, and the need for realistic understanding of timeframes, difficulties, and costs. A level of excitement around the possibilities of the technology – or hype – is necessary both for drawing investment from capital interests, and gaining attention from policymakers and publics. Developments and possibilities must not be overstated, however, such that investors gain unrealistic expectations around timeframes for returns on investment, or that society becomes concerned about the changes that might be expected from the development of error-corrected quantum computers. The Gartner hype-cycle<sup>44</sup> is a well-known model for the boom-and-bust pathways around investment and development that may be seen if expectations become too inflated. Indeed, as has been seen through the many decades of development in machine learning/AI, the net impact of over-expectation can be a sharp drop in investment and trust from markets. If quantum computing is not to experience a ‘winter’ in the same manner as AI, then this line between hype and reality must be carefully walked. Indeed there are some who believe that a quantum computing ‘winter’ is not only inevitable but already on its way<sup>45</sup>.

There is also a much more negative element to hype – although it is necessary to draw attention to developments in quantum computing from a security perspective in order to support broad understanding of potential threats (for example to cybersecurity), this also creates the potential for widespread misunderstanding and alarm around the size, imminence, and damage-potential of such threats. Cybersecurity experts are frequently at pains to point out that a cryptographically useful quantum computer is even further away than scaleable quantum computing, but these attempts to convey realistic understandings of risk-levels do not make for such attention-grabbing headlines as

---

<sup>44</sup> [https://en.wikipedia.org/wiki/Gartner\\_hype\\_cycle](https://en.wikipedia.org/wiki/Gartner_hype_cycle)

<sup>45</sup> Sabine Hossenfelder “The Quantum Hype Bubble Is About To Burst” <https://www.youtube.com/watch?v=CBLVtCYHVO8>

*“Cryptogeddon is coming for us all”*<sup>46</sup>. This represents another difficult path for the commercial quantum computing world to negotiate.

### 6.5.3: Innovation v risk-aversion

There is further tension between the desire to foster innovation, and the fact that – when it comes to quantum computing – the sums required may be so great that only governments and large companies are able to foot the bill – however innovation by its nature requires agility, a level of comfort with risk, and the ability to fail. These are not, in general, features of government investing. There are clear tensions between the desire to fund and promote innovative small businesses – which are by their very nature extremely risky – and the distaste of almost all large organisations, including governments, for losing their investment. There is an understanding that these two positions are incompatible, and organisations including the government choose various methods of trying to reconcile the problem, recognising that there is a great deal at stake – as I53 commented, *“your problem isn't investing in things that fail. Your problem is not investing in the billion pound company that works.”*

## 6.6: Conclusions

The most obvious differentiating feature between this study and the closely-linked study of the policymaking environment is the sharp contrast in pace between the cumbersome development of policy, governance or any form of regulation around quantum computing, and the speed of progress in the commercial quantum computing world. Policymaking is designedly a measured and thought-through process, but with such a fast-moving sector these traditional approaches may seriously hamper the ability of policymakers to react in relevant and appropriate ways. This creates a risk that the future could hold challenges similar to those seen in the AI space, where lack of expertise and timely regulation has created many deleterious outcomes for both individuals and communities, and lawmakers are struggling to respond with appropriate legislation. Policymaking in the UK is also restricted by the limitations on Parliamentary time that frequently serve to create a bottleneck. There is no mechanism for the removal of such frictions, and very little will to address the problem – considerable amounts of policymaking are therefore delegated to secondary regulatory and rule-making bodies. By contrast, commercial companies, investors, and other market-builders actively seek ways to eliminate obstacles to progress, and remove or avoid friction. This can only increase the difference in pace between the commercial sector and the policymaking context.

---

<sup>46</sup> <https://www.ft.com/content/a8204a7d-2922-4944-bdff-5449a8f3aee9>

This difference underlines the importance of a ‘prospective’ or anticipatory approach to some of the governance gaps described above in order to try and ensure that quantum computing is deployed in a manner that takes due regard of responsibilities to societal concerns and planetary boundaries. This is even more acute because of the levels of investment in quantum computing – particularly the surge of available capital in recent years – that are driving innovation and development. With this level of investment comes an expectation not just of rapid technological progress but also of potentially high-level returns, and competitive advantage – and preferably sooner rather than later. Anticipatory governance approaches are required in order to try and mitigate the effects of these pressures, but it is unlikely that such prospective governance approaches will be sufficient in and of themselves. There will need to be governance at varying levels of granularity, as discussed in section 2.1, that can address some of the gaps enumerated above in 6.5.1 *Governance gaps*. The research questions set out at the beginning of this field study – as one of their concerns – were focused on the question of whether, or to what degree, responsible-innovation type understandings and other forms of governance were carried from the research environment into the commercial world. It is clear from this study that at the point of any such translation the onus for retaining any governance framework rests with the entrepreneurs themselves. If concepts such as responsible innovation approaches or ethical frameworks are going to become embedded into their new endeavour, that is entirely their own decision. It is even possible that – should such approaches not find favour with investors – there will be pressures on them to actively avoid constraining themselves in such a way.

The huge surge in both the quantity and the type of private investment represents a further challenge for governance. Government support has been called for in previous work to ensure that private capital does not ‘capture’ the quantum computing sector (eg Ten Holter et al., 2021a), but the degree to which public sector funding is now outweighed by private funding represents not just a loss of public representation in the sector, but a reduction of the proportion of funding that is attended by relatively stringent governance requirements around research. However, there is potentially an opportunity to improve governance structures in privately-funded work due to the use of ‘evergreen’ funds that reduce the pressure for rapid returns on investment – there may be time and space to permit of greater reflection and anticipatory governance approaches. This would be worthy of deeper investigation.

### 6.6.1: Implications for the Framework

This study demonstrates not only the requirement for some form of anticipatory governance in the development of commercial quantum computing, but also the need for it to be relatively tightly structured (as compared to the traditional understanding of RI, for example, as a relatively nebulous

concept). There is a clear need both for a bridging framework at the point of translation from the academic into the commercial sector, and also for guiding principles for the in-practice phases of company growth. Any framework must take account of the commercial realities of running a company and be relatively straightforward to operationalise if it is to be of use.

However, it is also clear that it is not solely companies themselves that will need to engage with any such framework. The attitude of investors, tech-transfer teams, and universities will also be key in adoption, and there may need to be some form of mandatory element to counter the incentives and disincentives of a marketplace focused solely on economic concerns and rates of return.

---

# CHAPTER 7: A FRAMEWORK FOR ANTICIPATORY GOVERNANCE

---

*“The philosophy is that you push the power of decision making out to the periphery and away from the center. You give people the room to adapt, based on their experience and expertise. All you ask is that they talk to one another and take responsibility.”*

*Atul Gawande, The Checklist Manifesto*

This chapter reviews the case studies carried out to date and draws on the lessons and directions gained from them to create a novel Framework to support processes of responsible innovation. It then describes the evaluation of the Framework in various contexts.

## 7.1: Introduction

Thus far in the research for this thesis, four field studies have been carried out and analysed both for significant themes and for specific challenges exhibited in their area, which could inform a novel Framework and supporting indicators. Together, such a Framework and indicators form one of the objectives of this thesis. Each study has disclosed particular questions and guiding factors for the Framework, reflecting the priorities and context of its environment, and these can now be developed into an anticipatory governance tool useful for quantum computing, but also valuable for novel technologies and research projects more generally. The Framework will be shaped and informed by the following considerations identified by the field studies:

### 7.1.1: Factors from the landscape study:

1. This work drew attention to the need for **tensions that are surfaced in stakeholder and anticipatory work to be given persistence** throughout an iterative process. The Framework will need to provide a mechanism to retain context on the sources of conflict, challenges, and the trade-offs that may need to be made.
2. The landscape study also highlighted the importance of understanding the various perspectives on the challenges and concerns faced in each thread of the triple-helix, which

**may be rooted in similar considerations, but will be experienced differently** and will require different responses or a different application.

3. Finally, this study revealed significant gaps in current methods of anticipation, engagement, and responsiveness, and highlighted the need for a **more programmatic approach** that can reassure publics, create agile governance, and support the academics, policymakers and commercial organisations working to develop useful, safe quantum computing.

#### 7.1.2: Factors from the institutionalisation study:

1. This study drew attention to the importance of examining responsible innovation at **several levels of magnification** – the individual, the research-group or small community, and the wider organisation. All will be relevant in terms of assessing how responsible innovation is understood, actioned, and assessed. This need for different perspectives was also seen elsewhere.

#### 7.1.3: Factors from the policy study:

1. The inherent ‘messiness’ of the policymaking process indicates the **value of retaining elements of informality and redundancy** that allow a system to adapt rapidly to new situations. It demonstrated that attempting to adhere to a linear process may be an unrewarding endeavour, as well as not reflecting the reality of decision-making processes.
2. The study also supported the **value of asking difficult questions** and – even if these cannot be addressed immediately – retaining those questions for future iterations. Taking complicated questions out of the equation acts as a type of filtration, and can further amplify the effects of a lack of diversity and inequitable power-balances.
3. The importance of assessment is also key, and is linked to the question of linearity – there must be space for the decisions and outcomes of policymaking process to be reviewed and assessed. **Building in iteration**, with its necessary reviewing and recombining of different elements, as well as the process of repeatedly ‘going through’, can provide an opportunity for these evaluative processes.

#### 7.1.4: Factors from the industry study:

1. This study demonstrated not only the requirement for anticipatory governance in the development of commercial quantum computing, but also the **need for it to be clearly structured** (as compared to the traditional understanding of RI, for example, as a problematically nebulous concept).

2. Any framework must also take account of the commercial realities of running a company and be **relatively straightforward to operationalise** if it is to be of use.
3. The attitude of investors, tech-transfer teams, and universities will also be key in adoption, and there may need to be **some form of mandatory element** to counter the incentives and disincentives of a marketplace focused solely on economic concerns and rates of return.

### 7.1.5: Additional factors from the literature review

As seen in Chapter 2, some of the critiques of responsible innovation frameworks include challenges such as that they are too nebulous, too flexible, too adaptive, with a lack of guidance as to what constitutes ‘successful’ responsible innovation, and a shortage of agreed foundational principles such as a delineation of the values that are being expressed. There is concern that these characteristics raise challenges such as:

1. Responsible innovation being **insufficiently concrete** to be easy to apply in practice, raising questions such as ‘how do you know when you’ve done it?’ or ‘how much responsible innovation is ‘enough’?’
2. The **difficulty of comparing** even like-for-like projects. In particular the necessary reflexive elements may encourage a decision in a certain direction in one project and in the reverse direction in another, without any means of establishing how this is established or justified.
3. Another challenge was seen in the study of the Sciencewise consultations, where stakeholders’ input was often deemed to be ‘out of scope’ because it **could not specifically be addressed within the available remit**. This can undermine the credibility of stakeholder engagement and create disillusionment in those who have contributed.
4. The above points have contributed to the challenge of **‘assessing’ responsible innovation** in practice, identified earlier in the thesis as a possible roadblock in terms of institutionalising it as a method.

Although one tool cannot be expected to address every challenge in a fully satisfactory way, the Framework described in this chapter responds to the concerns raised by stakeholders, anticipates both current needs and future developments, and reflects on these elements together with designing a response. In this way the development of the Framework has itself followed a responsible innovation process.

### 7.1.6: Underlying principles

Overall, there are six key factors underlying the specific requirements above. Any framework that attempts to address these requirements must be:

1. **Transparent** – it is critical that the responsible innovation process can be comprehensible from beginning to end, that the deliberations and results can be interrogated by users and third-parties, and that those concerned can be aware of why particular responses were chosen.
2. **Accountable** – this does not only mean that the project or task may be accountable to a funder, or a governing body, but also refers to accountability to the stakeholders concerned in the process, a key concern in participatory work. This is particularly true for the type of situation discussed above, where stakeholder concerns may be considered to be ‘out of scope’.
3. **Comparable** – projects must be able to establish baselines and comparisons for anticipatory governance work, in order to cumulatively gather evidence, and also to enable decisions to be challenged if necessary. There is also a need to allow for like-for-like comparison between projects.
4. **Adaptable** – the framework will need to be sufficiently flexible that it can be utilised around many different kinds of technology or development process, as well as being translatable across disciplinary and sector boundaries.
5. **Repeatable** – this is allied to comparability, in the sense that if the same process with the same stakeholders was to take place again, then it would be possible to test the outcomes of the earlier exercise. This does not argue that the outcomes would be the same, but any changes would themselves be illuminating.
6. **Documented** – this element underlies many of the others stipulated here and represents a novel addition to responsible innovation work. Documentation of processes is key in many industries and sectors, as has been seen, because of the support it provides not only for decision-making processes, but for principles such as repeatability and transparency.

These are the shaping elements that the Framework will need to satisfy.

## 7.2: What is a framework?

A framework can take many forms, and indeed conceptually is fundamentally different in different disciplines. For social scientists the key idea is of a ‘frame’ that can comprise a set of concepts and theoretical perspectives. In software engineering, however, a framework is a much more concrete building-block – almost a template – that is used as a basis from which to build out a new software development, while in management literature and process engineering, a framework might take the form of a graphical model that describe processes, elements, and their relationships to each other. Elsewhere it can be a decision-making tool, a form of checklist, a matrix, or a process. In most cases,

however, the common thread is that it forms a scaffold, which has been designed to assist with guiding thinking and structuring a way forward for a given procedure. The objectives will vary, but may prioritise consistency between different instances of the procedure, whatever it may be, or aim to focus attention on any gaps that may thus be perceived. It is in some ways a mental model – a simulation that allows users to think and reason in a particular way. As Forrester, (1971) points out, such mental models are useful for identifying weaknesses or lacunae that can then be addressed, strengthening the overall procedure. Additionally, he suggests that use of such models can force assumptions to be examined, and any internal contradictions to be surfaced for either acceptance or resolution. Some of the more common types of framework from different disciplines, as well as their value for construction of this new framework, are discussed below.

### 7.2.1: Process models

It has been argued earlier in this thesis that the use of ‘process models’ and the potential for drawing some of these techniques into the deployment of RI might provide a useful pathway for analysis, or help to resolve elements of some of the critiques of RI.

Process models are diagrammatic representations that “*convey tasks, sequencing, decisions, participation, and information*” (Long, 2014 p1). The diagrams are accompanied by explanatory text but it is the diagrams themselves – following carefully worked out representations of tasks, decision points and so on – that are the key element. As with other types of framework, they provide structure and create a representation that can then be analysed for gaps or refined, depending on the requirements. Shaping the process in this way can provide elements of formality and repeatability. However, like all models, they are useful but limited: they cannot take account of the context in which the process is used, or the external interfaces; they also cannot represent what may be regarded as the ‘human’ elements (the importance of acknowledging this was seen in the policy field study, where the human factors frequently override the linear processes that are nominally in place). With due regard to these limitations, elements of process modelling may offer useful support in the development of an anticipatory governance framework.

### 7.2.2: Policy frameworks

The concept of a framework as it is understood in policy has a certain amount in common with the process model approach described above, but will be more influenced by the ‘social science’ type of perspective. This means that the process is not purely functional, but instead is guided by certain underlying principles and values. The paths and processes through which policy is developed tend to be less widely understood than, for example, the ways strategy is developed in commercial

companies. This may be due, perhaps, to policy-creation taking place in relatively closed-off and self-sufficient institutions that can appear arcane and impenetrable from an external perspective. They are also, naturally, narrower in their field of applicability. Policymaking frameworks are therefore less well known than comparable work in industry, but like all complex processes that need to be relatively repeatable, policymaking has its own best-practice structures. The work of Harold Lasswell (Lasswell, 1956), who described a five-stage cycle of policy-creation, remains influential and continues to be discussed today (for example Strassheim, 2019). Lasswell's process included agenda-setting, policy formulation, decision-making, policy implementation, and policy evaluation. Other frameworks include the ROAMEF process of Rationale, Objectives, Appraisal, Monitoring, Evaluation, Feedback (HM Treasury, 2022). Both of these examples follow linear models, though as demonstrated in the policy study this is frequently not reflected in practice, revealing a significant gap between the hypothetical norm and the lived experience in practice of those developing policy.

### 7.2.3: Frameworks in industry

Commercial organisations and the management literature are familiar with, and have produced, a wide range of tools and models which are frequently employed in industry and consultancy. These range from the well-known (eg SWOT) to the more rarefied (such as Six-Sigma practices). Their purposes can be extremely diverse, reflecting the number and variety of concerns on which a commercial organisation may need to impose some structure in order to achieve some certainty. They can include marketplace modelling through a tool such as Porter's Five Forces, which allow a company to analyse its business environment through the lens of competitors, customers, suppliers and replacements (Porter, 1979). Horizon-scanning exercises, using tools such as PESTLE – which assesses the political, economic, social, technical, legal and environmental context for a business (Day, 1990) – are strongly aligned with the type of anticipatory work seen in RI, and indeed sometimes cross over into research (eg Nandonde, 2019). Strategic decision-making tools such as SWOT analysis – strengths, weaknesses, opportunities, threats – are also in wide use, providing as they do a blending of the macro/external and micro/internal environment (Wehrich, 1982). Less well-known outside management consultancy, but still in widespread use, are tools such as the Balanced Scorecard (Kaplan & Norton, 1992). Consultancy firms and professional advisers use tools and frameworks like these to focus and structure their own thinking and that of their clients, as well as to surface implicit knowledge, assumptions, and understandings.

Many of these tools will also be familiar to those in policy and academia, but it is particularly valuable (although outside the scope of this thesis) to compare some of them to responsible innovation-type frameworks. As discussed previously in Chapter 6, translation of academic-generated models such as

RI can often fail to recognise anticipatory, participatory work that is already taking place in commercial companies – it may not be the case that businesses are unwilling to carry out these activities, but rather that the communication of potential benefits is insufficiently persuasive, or that there is a failure to recognise work that is already taking place (in an echo of the *de facto* RI work discussed in Chapter 4).

#### 7.2.4: Academic frameworks

The divergence, specialisation, and variety of academic fields makes it less likely that any given tool enjoys wide acceptance or understanding (as, for example, the SWOT model could be said to do in the commercial context). However, the principles of creating and using a framework in academia are somewhat similar to those used in industry – academia has long recognised the value for shaping thinking that frameworks and models provide. For example the version of RI used in the UK is based around the AREA Framework, but there are also additional or alternative frameworks such as the 4P set of indicators (Stahl et al., 2017), and numerous RI or RRI tools. Some of these have their foundations in European projects, and may include physical prompts such as ideation cards to spur debate and facilitate discussion.

#### 7.2.5: The Double Diamond

The British Design Council's Double Diamond model (illustrated in Figure 7a) has become one of the most well-known and widely-adapted frameworks in design over the last two decades (Ball, 2019). Its versatility has meant that its principles and processes are utilised in numerous countries and across widely diverse fields including healthcare in Australia (Banbury et al., 2021); management and business analysis in the European Union (Molendowski & Źmuda, 2013); and economic and environmental sustainability in the UK and Germany (Andrews et al., 2021). It is shaped around the partnered mechanisms of *diverging* and *converging* – within the model these processes of opening up and then closing down again are repeated, to try and ensure both that a multiplicity of voices are incorporated, but also that movement forward does not become 'frozen' in a morass of competing requirements. It thus implicitly recognises that trade-offs play a necessary role in any decision-making process. There are therefore clear parallels with not only the AREA framework for RI, but also some of the other frameworks from industry and policy discussed above, such as the PESTLE analysis. These processes of 'opening up' and then 'closing down' – taking on board new ideas, information, and viewpoints, then absorbing, reflecting on, and moving forward with improved understanding and responsiveness – then iterating – provide one useful way of clarifying and scaffolding the type of anticipatory governance mechanisms that RI frameworks seek to operationalise.

This Double Diamond approach is therefore a useful model because it provides a means to allow for tensions, meanings, assumptions, multiple inputs and stakeholder nuance to be gathered, while also recognising that not all of these inputs can be instrumental in the final response.

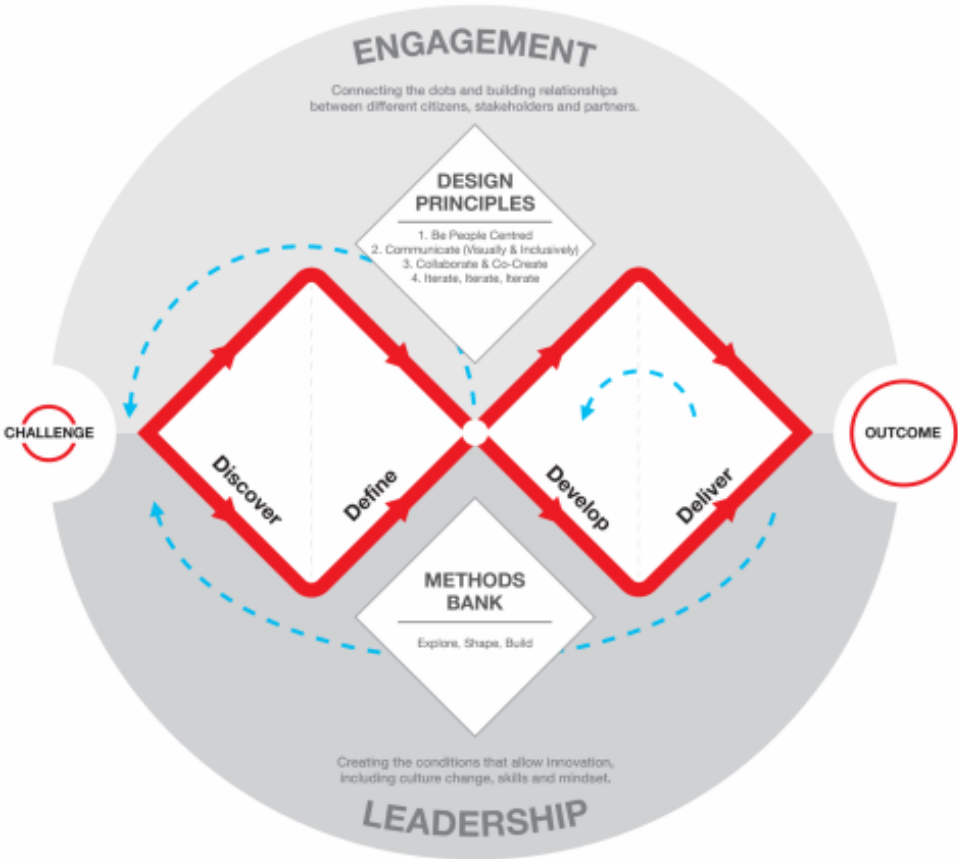


Figure 5: The Design Council’s Double Diamond (©Design Council)

However, as outlined above, it is also important to *retain* and *revisit* even the elements that could not be incorporated into the first iteration of the process model and to continue to allow them to influence the outputs. This is a factor that the Double Diamond model lacks, and one that will need to be developed for the present framework.

### 7.3: Creation of the Beehive Framework

The preceding case studies and the discussion above on the nature of the frameworks used by the various strands of the triple helix (as well as others) have all contributed shaping factors to the design of the Framework. Additionally, the research has demonstrated the following important themes.

1. As seen in the policymaking fieldwork, although policymaking processes are not quite as linear as they are deemed to be, there are still very few evaluative processes built in – in general once a policy is ‘made’ it is seldom revisited by its makers, and metrics for success or failure are not established at the point of policy formulation. This demonstrates **the importance of incorporating mechanisms for evaluation** at key stages. In the SPICE project described by Stilgoe, Watson, et al., (2013), one of the conceptually important papers for responsible innovation, Stilgoe et al describe the concept of ‘stage-gates’ to act as pause points at given stages of technological development. Although these were not ultimately built into the AREA Framework, they remain a useful shaping concept as part of a reviewing and reflecting process. Such ‘pause points’ are also seen in project management methodologies such as PRINCE2<sup>47</sup>.
2. One of the problems for the development of any innovation – whether of technology, or policy – is **how it manages divergence of views**, balances different values, trades off one value or perspective against another. The way in which these divergences are handled is crucial in terms of relationships with both stakeholders and society. It is essential to retain these tensions around different values and divergent views, and to represent them in the Framework, for it to have a way of *holding* – retaining – such diverse opinions and challenges and taking them forward into the design of the work.
3. Within the participatory design literature, explored in Chapter 2 as a potentially valuable contributor to the responsible innovation discourse, commentators discuss the **politics of participation**, and who is enabled to be a stakeholder. This theme raises challenging questions, such as *whose views are considered to have value?* and *whose arguments carry the day?* These political questions, centring as they do on the power-relations at play, are – perhaps unsurprisingly – very present for policymakers but much less so in other arenas. Participants in the policy study were aware that the careful choice of stakeholders, as well as framing the question in certain ways, may provide the answers that are most convenient – meaning that diverse views can potentially be subsumed into a response that has the appearance of consensus. Such questions around the sourcing of different opinions were viewed differently in the industry study, where a lack of diversity within a company (perhaps because it is very small) may be a serious limitation on its ability to find creative responses to challenge. Creative and diverse responses are often not the goal in policymaking consultations.

---

<sup>47</sup> <https://en.wikipedia.org/wiki/PRINCE2>

### 7.3.1: The Beehive Framework



Figure 6: The Beehive Framework

The Beehive Framework, shown above, is based around a ‘process model’ approach to the methods of responsible innovation, which means that it formalises one way in which responsible innovation can be carried out. There is no implication of normativity, and it is not proposed that this Framework is the way in which responsible innovation ‘should’ be operationalised, rather that this creates a layer of process that can structure the methods of responsible innovation. This is designed for ease of use; for translation purposes into commercial environments; and to address some of the challenges identified earlier in the case studies and outlined in 7.1 above. It draws on the Double Diamond approach in order to illustrate the diverging and converging elements, as this sense of ‘opening up’, then reflecting, and ‘closing down’ is vital. However, unlike the Double Diamond it is a single process that is then iterated upon, rather than, as seen in section 7.2.5, different processes that follow a similar

pattern. This means that the central figure only requires one repetition rather than two, because the intention is that this process is revisited. The purpose of this repetition is to re-reflect on and re-process the considerations and discussions from earlier rounds, as well as potentially widen discussion, for example with a new or expanded set of stakeholders.

In another departure from the Double Diamond, the Beehive Framework draws upon the strategic and management frameworks discussed in 7.2.3 for one of its most important features, which is the *recording* of information. This is discussed in more detail in section 7.4.4, and serves to address several of the challenges identified during the fieldwork. The purpose of the areas grouped around the central 'process' cell is to provide a workspace that is intended to be populated – to act as a storage space for the stakeholder-mapping, reflective, political questioning, and consensus or conflict elements. Populating the model in this way addresses the challenge of both identifying and *retaining* the tensions, trade-offs, concerns and other information gleaned from stakeholder work, anticipation, and reflection as these move forward through the decision-making process, with the intention that these can be revisited at a later stage as the process iterates. This was a key principle brought through from the case studies – that information, competing priorities, and challenges should not be lost or discarded even if they cannot influence the process in its present iteration. It is not possible to say definitively what will be of value in future iterations, or how a particular technology, process, or policy, may develop in the future. Accordingly, these factors need to be retained and revisited. A crucial difference between this and other frameworks, therefore, is that the Beehive Framework acts as a living document for continuing use, in the same way as a project management plan serves as an ongoing outline for guidance and referral. The stages of the process are discussed in more detail below, but the objective is to ensure that, unlike a decision-tree process, or stage-gates, previous work is not lost as decisions are made, but instead is built upon, and insights gathered and developed during the 'opening up' process are retained.

A further purpose for this recording of information is to provide transparency around decision-making, and comparability between projects – the Beehive Framework can act as a record that is available for similar projects to consult, and perhaps to critique.

For the purposes of design, and the evaluation described in section 7.5 below, the online whiteboard tool 'Miro' has been utilised as a mechanism for creating templates, recording, and storing the various elements that populate the surrounding hexagons – however, any widespread propagation of the Framework and its toolkit would require a more robust and permanent delivery and storage system. Miro's functionality is particularly valuable for several reasons. It is shareable online, meaning that multiple parties and stakeholders are able to collaborate, improving accessibility. It is simple to use,

but provides features that can add layers of complexity. It also has the potential to provide a vast workspace and supports the ability of users to add increased levels of detail by zooming in to create more workspace. A Miro board also allows for workshoping – whether online or in person – and the capture of ideas during discussion that can then be used for populating the workspaces at a later point.

In these ways, the overall Framework meets the following concerns identified in the case studies, as well as the underlying principles identified in 7.1.6:

- A more programmatic approach that systematises (to some degree) the process of working through responsible innovation and allows for simpler operationalising during processes of transfer (eg between academia and industry).
- Multiple iterations allow the use of different ‘lenses’ of magnification for governance, as well as being useable both for fine-grained projects and for high-level strategic discussion.
- The creation of a space for evaluation and assessment, as well as repeatability.

The ways in which the elements of the Beehive Framework answer the other challenges raised by the field studies are elaborated below in the discussion of how it is used in practice.

## 7.4: Using the Beehive Framework

### 7.4.1: Diverging – anticipation and engagement

The process element of the Beehive Framework is illustrated by the arrowed pathway through the central hexagon. As can be seen from Figure 7, when the process pathway (in purple) enters the central hexagon, there is an initial diverging phase. This is the ‘opening up’ point - both to stakeholders and to anticipatory processes. It is during this anticipatory stage of the Framework that stakeholders are selected and invited to participate. Bearing in mind the political challenges that were discussed at length in Chapter 2, this is the opportunity to engage in a broad range of activities including speculation, envisioning, and use of moral imagination, to think through questions and challenges in a collaborative way. The related hexagons are then used to record both the stakeholders invited, and the questions they raised, the scenarios that were suggested, the challenges, possible solutions, roadblocks, tensions. The users can populate the framework with the details of those who were involved, and the work that was undertaken.

This phase addresses the following challenges identified in the fieldwork:

- Different perspectives on issues – this is achieved through the inclusion of stakeholders.
- Retaining a variety of viewpoints – these are recorded in the model.

- Retaining ‘messiness’ and redundancy – ensuring that concerns are not lost
- Identifying ‘difficult’ questions – even if these cannot be answered
- Creating a space for imaginative understandings of possible outcomes and paths for development

These elements also reflect the concern from the policy study about the need to record decision-making processes in case of review – as this Framework is designed to be used iteratively, such review and evaluation is built-in and therefore the recording element allows for progressive construction of a collaborative knowledge-base.

### Stakeholder engagement

One of the challenges for the ‘diverging’ stage of the framework is the stakeholder engagement process. However, although these participatory processes may be unfamiliar, as seen in the review of the concept of participation in Chapter 2 (Section 2.3.2), they are essential. This is not just for the practical expertise to be gained by drawing on a multiplicity of viewpoints, but for the purposes of testing ideas, investigating impacts on diverse populations, and engaging with societal actors to improve trust. This is a key element, not least because in this as in many contexts, increasing the weight of numbers adds depth and complexity to these cultural and societal questions (Powell et al., 2009). There are many methods of doing this – for example full co-creation; or participatory design; or consultation – and the particular method chosen will be context-specific, but the process is essential. The stakeholder mapping exercise, including the ‘circles of engagement’ approach (described in Section 6.3.1), can also provide a useful guide as to who may potentially be involved in this phase of the process (or indeed in later iterations).

### Anticipation

As discussed in Chapter 2 (section 2.3.1) there are many challenges for anticipatory work. These can include epistemological questions such as considering whether we should be attempting to govern the future, and the unlikelihood that we can ‘predict’ outcomes in any meaningful way. However, as argued in that section, the anticipatory element of RI does not seek to predict – rather, it has more in common with the strategic frameworks seen in industry (and indeed in other sectors such as the military), where such exercises as envisioning, collective imagination, and exploration of possibilities, are used to ‘prepare’ for possible outcomes.

Figure 7c shows this part of the Framework, populated with data from a project. It can be seen that re-iterations of the process suggested different stakeholders (denoted by different coloured notes), and

that those stakeholders were associated with bringing to the fore different concerns around the research questions.



Figure 7: Diverging – anticipation and stakeholder engagement

#### 7.4.2: Reflecting – surfacing concerns

It is inevitable that not all stakeholder considerations and apprehensions can be acted upon. As has been seen in discussions around policy and the difficulties of engagement, some concerns may conflict or outweigh each other, priorities may compete, or some elements may need to be subservient to other factors due to external pressures or unavoidable realities. A wide variety of elements can influence outcomes. Although the Framework shows the reflexive activities – reflecting on the stakeholders and reflecting on the anticipatory work – taking place in two different areas (Figure 8), it is likely that in practice there is significant overlap here, especially where a particular group of stakeholders expresses a particular view. This connection between the set of concerns or ideas, and the groups that raised them should be made clear, in order that the anticipatory work carried out by individual groups is not lost.

The reflection phase of the work addresses the need for evaluation not only by requiring a conscious reflective process, but by requiring the recording of this process. As has been explained in the earlier

discussion of professionalism (section 2.2.2, and the discussion on Schön, 1983) reflective practice is part of continuing personal and professional development in many professions – it creates a bounded space where reflection and consideration not only can but *must* take place, and must be recorded, before moving on to a next stage.



Figure 8: Reflecting – surfacing concerns

### 7.4.3: Converging – response

Any decision-making, strategic, or policy process needs to move forward, however, and in order to do so after the opening up and reflecting phases, it must ‘converge’ again in order to make choices and decisions. This inevitably also means closing off some of the available options. But before doing so, the rationales for these choices should be carefully considered, and should be recorded. If a particular path is chosen, the reasoning for that should be documented in the Framework – it may be that these choices can be revisited during further iterations, may be affected by the development of the project, or could gain additional weight from discussions with new stakeholders. It is important to note again that although the Framework’s demarcated spaces may suggest separation of the various ideas and processes, these are not hard boundaries. For example it is not crucial that an issue be recorded in the southern hexagon of the ‘closing down’ stage as opposed to the northern. At the point of the process

where the arrowed pathways come together and rejoin, therefore, the various concerns and challenges may be assumed to be part of one single process. In particular it is key at this point to record tensions and trade-offs. These are inevitable in any decision-making process as there are always compromises that need to be made.

An example of such a trade-off is illustrated in the work briefly described in Figures 7-9. Within this project, our research showed that citizens require, or expect, that autonomous vehicles will be recording data as they deploy in order to facilitate fast and accurate accident-investigation. However, this data-recording, if carried out at the required level of detail, would create terabytes of data every time the vehicle was operating. This would have significant knock-on implications for the privacy of other citizens who may be recorded as the vehicle drives around, problems of data-provenance if cases need to be tried in court, security and storage of colossal volumes of data, (and therefore the carbon-footprint of this), as well as technical challenges for manufacturers. There is therefore a necessary trade-off to be made, and this must be noted, particularly given the likelihood that as this particular technology develops, solutions to some of these challenges may be found. Unless this tension and the expectations of society in relation to this issue are recorded, it may not be regarded as a question for further research. (Indeed, our research in this project discovered that these data-recording expectations were not being considered as a priority by the manufacturer that participated.)



Figure 9: Converging - response

#### 7.4.4: Recording and reporting

As noted previously, the element of recording, or ‘narrating’ the process is key for the Beehive Framework, and sets it apart from many other frameworks – this ‘recording’ element draws on diverse antecedents, such as the participatory design ‘diary’ described by Pilemalm (2018) from one project and discussed in Chapter 2, and classical project management methodologies drawn from the management literature that emphasise the importance of clear communication and scoping (for example Sargeant et al., 2010; Špundak, 2014; Tinoco et al., 2016). It also reflects the point raised by policy participants in the Chapter 5 study, about the importance of reporting the factors taken into consideration in a decision, in case of later audit or judicial review, and is similarly informed by the systems of precedent within legal frameworks, which likewise explain their reasoning. These systems all emphasise the necessity of documenting the processes undertaken and the considerations acknowledged, partly as a means of providing transparency and accountability, but also to enable processes of comparison and assessment to be carried out.

The ‘recording’ element thus addresses many of the challenges identified in the preceding chapters, as well as supporting other aspects of responsible innovation – for example:

1. The process of documenting can help to illuminate any gaps in the process, such as stakeholders that may not yet have been consulted.
2. The requirement to elucidate and list the various parties consulted and their inputs, can also act as a prompt for reflection and thus as a form of data analysis. In this way, the Framework can also act as a support for the processes of insight and consideration that are such a significant part of responsible innovation work.
3. The creation of a record of the responsible innovation carried out, in a format applicable across many different project types and disciplines, will allow for simpler processes of comparison and assessment, identified as necessary for embedding responsible innovation in to institutions and projects.

However, as stressed earlier, it is important to note that the tool is designed to be used as a ‘living’ document that can be readily referred to and shared, and therefore the requirements for its completion must not be too onerous. There is a necessary trade-off here between usability and the provision of a ‘complete’ record, but the balance must lean in favour of the former in order to ensure the Framework’s ongoing utility as a record and a reference. This ease of use factor will be particularly important for the iterative element and for translation into commercial environments, such as being used within spinout companies.

#### 7.4.5: Iterating – a circular process

The importance of iteration has been noted previously, and accordingly the Beehive Framework is designed specifically to be re-used and iterated upon during any development or decision-making process. This is for several reasons:

1. The work to identify and document stakeholders; the ‘opening-up’ process that foregrounds anticipatory concerns, tensions, and multi-disciplinary inputs; the reflective and responsive processes that create a decision – all these elements form part of the groundwork for the decision. It is critical not to leave behind these understanding and rationales once a decision is made or a development processed, but to retain and take them forward – revisiting them through iteration serves as a reminder of the previous opening-up and reflective phases.
2. A strong rationale for iteration is not just to revisit earlier decisions, but to *assess* those decisions. The policy fieldwork in particular demonstrated that once a process (whether messy or linear) is complete, there is often a lack of time, willingness, and resource to evaluate its effectiveness, and metrics or indicators for success are frequently not considered or built in. Iterating on the Beehive process provides a means of evaluation by re-assessing and potentially re-opening the process.

- The reflective and converging stages of the process are likely to recognise gaps that were not identified in the first phase, and will offer the opportunity to revisit and enrich the initial considerations.

There is not automatically a logical end-point for these iterative processes, but it is suggested that two to three iterations will offer sufficient opportunities for analysis.

### 7.4.5: Framing the process

Figure 10 shows the completed Beehive Framework for the small research project referred to above. This has been through three iterations, indicated by different notepad colours. The colour-coding has been retained to indicate the concerns raised by given stakeholders. Should more space for recording be required, additional cells could be added – the Beehive Framework could also be enlarged by adding wider circles of cells as an alternative way to denote later iterations, as it could reflect widening groups of stakeholders.

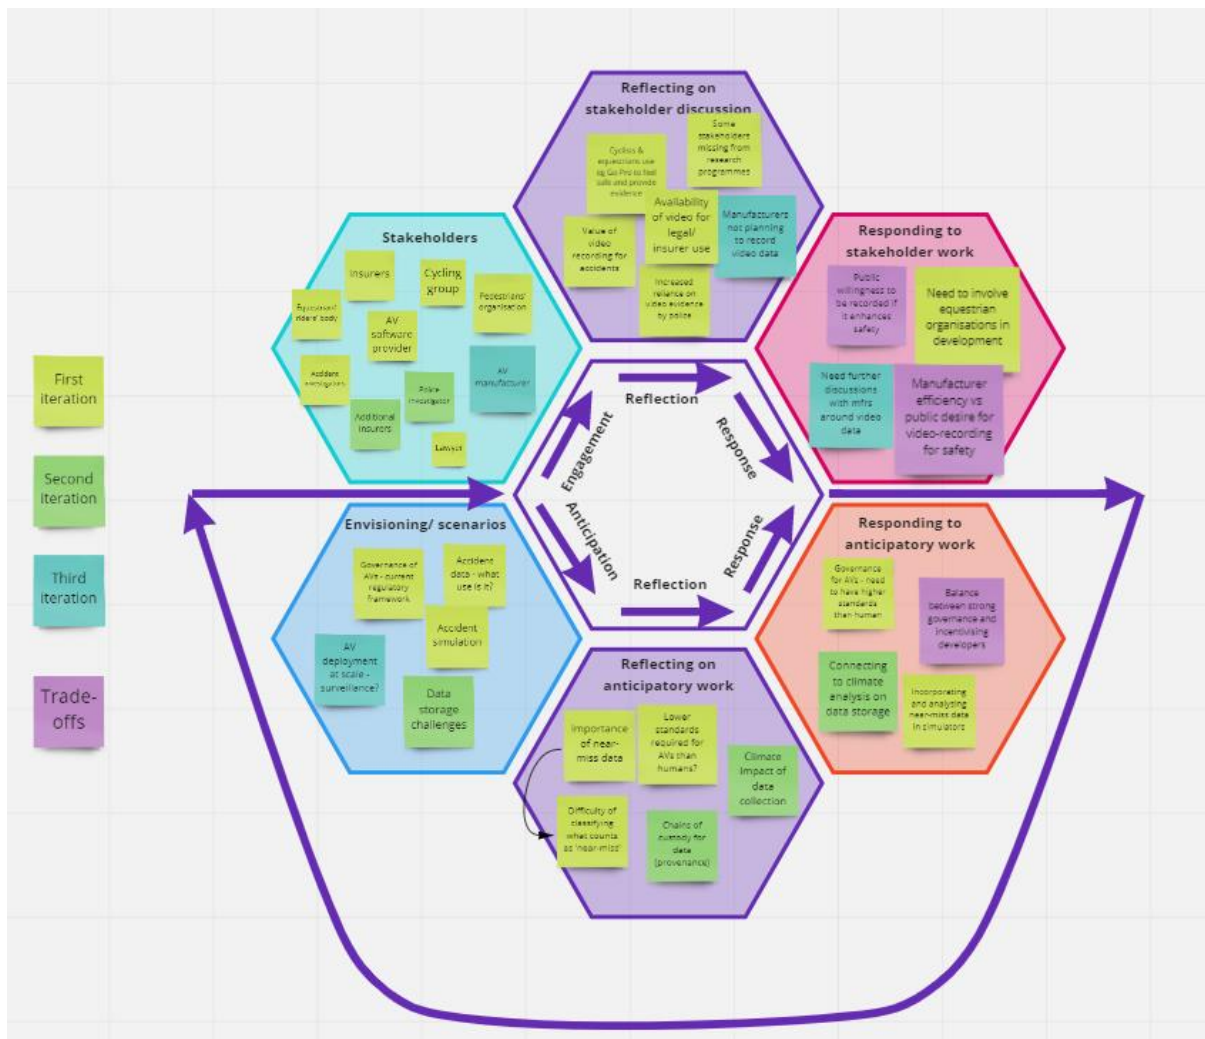


Figure 10: Sample populated Framework

## 7.5: Validating the Beehive Framework

As outlined above, the Beehive Framework was developed as a result of the five preceding chapters of the thesis - the literature review (Chapter 2) and the four case studies (Chapters 3-6). However, in order to establish its utility, create a form of peer-review process, and potentially note elements for further development, it was necessary to evaluate both the underpinning theory and assumptions, and its use in practice. This was undertaken as a three stage process:

1. Delphi study with academic, policy and industry participants (the three strands of the triple-helix, that have been investigated in parallel throughout the thesis) to canvas feedback and opinion on the design.
2. Presenting the Framework to various groups for critique and discussion – these included the author’s wider research group, as well as associated groups that expressed interest in a method of operationalising RI.
3. Field trials with ongoing research projects to evaluate its use in practice, including establishing whether the correct balance had been struck between simplicity of use and explanation, as set against its ability to incorporate complexity and nuance.

### 7.5.1: Delphi study

The traditional Delphi method was originally developed in the 1950s for the RAND Corporation, with the aim of assessing the impact of technology on warfare (Dalkey & Helmer, 1963). The method is based around iterative circulation of work to a panel of experts who respond anonymously, and see each other’s responses in further iterations. It is therefore a group interaction method, but one that – by distancing the participants – aims to *“remove ... the effect of powerful personalities, individual’s status, social pressures of conformity, and peer pressure”* (Chalmers & Armour, 2019 p717). The method thus attempts to address some of the ‘political’ challenges of group work (cf *The politics of inclusion* Chapter 2). In the classic Delphi study, each iteration aims to move closer to expert consensus – this iterative process, and the feedback mechanism make it unique in terms of method (Chalmers & Armour, 2019; Lopez-Perez & Olvera-Lobo, 2018).

There is also a variant of the Delphi study, however, that does *not* seek consensus, the ‘policy’ Delphi. In this variant, a diverse panel is consulted in order to draw out as many perspectives as possible. It is particularly valuable in *“investigating problems that require inputs from multiple, different, and often conflicting points of view”* (de Loë, Melnychuk, Murray, & Plummer, 2016, p79). Its value in the policy context comes from its ability to enquire into complex areas and reveal multiple challenges that may not have one normative solution. It draws on multiple sources of expertise in order to illustrate the landscape of possible concerns or solutions. This type of Delphi may consequently be seen to have

elements in common with other stakeholder consultation methodologies – including responsible innovation and participatory design – in its seeking-out of a range of views. The policy Delphi format therefore specialises in sourcing divergent opinion and this was anticipated to be a valuable tool in evaluating the Framework.

### Participants

Given the small and relatively tightly-knit community of quantum computing researchers, policymakers and industrialists in the UK, and the previous engagement with them for this research, the entire field of people who had been approached for the various case studies in this research were approached for a second time to ask them if they would be willing to review and potentially evaluate the Framework. A research Report that explained the genesis, grounding, and use of the Framework was therefore written up and sent, along with explanatory text, to all the participants of previous studies. This was accompanied by a link to a short questionnaire, with a final free-text response component, as well as the opportunity to comment by email. The responses were collated both quantitatively (for the questionnaire element) and qualitatively (for the free response element).

### 7.5.2: Feedback

The comments received on the Beehive Framework through the survey responses were almost entirely positive. Although this study encountered the usual challenge of Delphi studies – that many potential participants do not have the time to engage fully with the work – several encouraging responses were received. These included an invitation to present the Framework to the Directors of the UK Quantum Computing & Simulation Hub, and potentially thereafter to the Management Board of the Hub. The response from policy participants was also helpful, with offers to forward the Report on to other interested policymakers, and CDT Directors who expressed gratitude for a tool that they could use with students. It is possible that those who found the Framework to be less helpful for their work did not respond, so further work would be necessary to establish reasons for lack of response and any constructive critiques.

### 7.5.3: Field-testing

During the course of the research on which this thesis is based, the researcher was also either working on or consulting on various other ongoing projects within their wider research group. These included:

- i) a project providing responsible innovation training and consulting services (ORBIT)
- ii) a project examining the utility and use of the data collected by autonomous vehicles, including its ethical and legal implications (RoAD)

- iii) an investigation into the ethical, legal and societal implications of the deployment of autonomous vehicles (RAILS)
- iv) an adolescent mental health project, including digital interventions (AMHDM)
- v) the creation of a design framework for novel technologies that could ensure they were designed to be compliant with Paris Climate Agreement targets (PARIS-DE)
- vi) a responsible innovation integration project within the Trustworthy Autonomous Systems Hub (TAS-RRI)
- vii) a project focused on improving responsible-innovation based engagement and communication within the quantum computing community in the UK (ResQCCom)

All these projects had a strong grounding in responsible innovation and engaged with challenges raised by novel technologies. It was therefore considered that they could be useful testing-grounds for the Framework. After discussion with the principal investigators of the projects, the Framework was presented to relevant members of the projects' research teams.

## Method

Although the Beehive Framework could potentially be instantiated in various forms (for example it could simply be printed out and filled in, if an offline version was required), for the above-described evaluations and presentations the Framework was created as a Miro board template to which the participants were given access. Miro was used for several reasons. As outlined above it was considered to be ideal for this process as it offers a large area for working and the ability to zoom in and out to various levels of magnification. A key benefit is also that, just as the responsible innovation process is designedly consultative and collaborative, a Miro board is specifically set up to be used in this way – its shareable nature means it can be worked on by several people at once. The 'sticky-note' facility provides a means for multiple participants to populate the 'recording' portions of the Framework.

## Findings

No framework can fully meet all possible responsible innovation requirements (for example, it cannot interrogate the values underpinning a particular design approach, only record challenges for it, and possible improvements), and the use of the Beehive Framework will therefore be highly context-dependent. This means its use will need to be situated within a wider understanding of responsible innovation processes and commitment to engaging with the challenges of participatory work. However, in the above-mentioned projects it was found to be easy to understand, straightforward to use, and readily comprehensible for those without an in-depth understanding of responsible innovation. This latter group also included industry partners in some of the projects, although it was

not possible to introduce it to all partners. It can thus serve as an introduction to the concepts and approaches of RI, and may potentially be valuable as an aid in institutionalising its processes.

## 7.6: Conclusions

The design of this Framework has been informed and shaped by the preceding case studies, and by a review of the literature on responsible innovation and novel technologies (specifically quantum computing). Its aim has been to try to solve for some of the challenges identified both as part of the responsible innovation/ novel technologies discourse and those additionally identified during the fieldwork studies.

These challenges included: the difficulty of knowing in what order to carry out RI activities, scoping those activities, directing them to obtain the most useful and societally aligned results, and evaluating or assessing the RI work carried out in order to ascertain whether it was of sufficient quality and quantity. The underlying themes that emerged demonstrated that RI work needed to be transparent, accountable, comparable, adaptive, repeatable, and documented. These were the key challenges that the Beehive Framework was designed to address. The work on evaluating the Framework has demonstrated a positive response from users in terms of usability and utility, as well as potential avenues for development and deployment. The Beehive should therefore provide a model that can be used and adapted both in quantum computing and other computer science fields – it is hoped that its worth will be particularly apparent in fields where traditional governance is not yet a significant part of the ecosystem, for example responsible machine learning.

However, it is important to stress that the Beehive is a necessary – but not sufficient – response to the challenges discussed in the preceding chapters. Many of the difficulties identified during the course of this research in developing governance for novel technologies, institutionalising responsible innovation, or creating new policy, are *not* due to a lack of tools, but rather are rooted in behaviours that respond to the prevailing incentives. For the Beehive to be a useful addition to the anticipatory governance of novel technologies, it must be grounded in changed or changing behaviours that understand it is not a tick-box exercise, but rather is designed to encourage developers, policymakers, and researchers to think in new directions about the way in which they design, develop, and deploy technology, and to support those new directions. It is also important to add that the Beehive Framework cannot guide decisions in one direction or another – as discussed earlier, it does not interrogate the values underpinning given choices, rather, there is an underlying assumption that the adoption of such a Framework itself indicates a commitment to responsible approaches. However, the methods described here can increase the chances that a responsive, considered, and informed path can be chosen, assuming there are incentives to do so. Those who work in novel technologies will

need to be both empowered and encouraged to consider additional factors in their development, otherwise any tool such as the Beehive will fail to make a difference.

Finally, it is believed that the Beehive offers a novel contribution to the RI discourse and can act as a catalyst for discussion around the need for granularity in responsible innovation, as well as providing a possible template for future development work.

---

# CHAPTER 8: CONCLUSIONS

---

*“We keep moving forward, opening new doors, and doing new things, because we're curious and curiosity keeps leading us down new paths.”*

*Walt Disney*

This chapter reviews the programme of research completed for the thesis and discusses not only the outcomes of the various studies, but also the underlying themes that have been drawn out through the work. It closes with an assessment of the contribution of the thesis, along with recommendations and directions for future research, in terms of both broader engagement and deeper investigation.

## 8.1: Introduction

At the outset of the thesis, the early scoping work and the literature review (Chapters 1 and 2) made the argument that achieving a well-informed, holistic, and pre-emptively governed grasp on the development of quantum computing was vital for several reasons. It was necessary firstly to secure public trust, particularly when publics had made clear their requirement for good governance, but also, crucially, to ensure that such trust was justified – ie that the sector merited the public’s trust. One way for such public confidence to be justified is through evidence from the entire sector of making every effort to ensure that the development and deployment of quantum computing is managed in an equitable, environmentally responsible, and societally responsive manner. It was further argued that one of the largest challenges for ensuring such a responsible development and deployment trajectory is the speed with which the growth of a quantum computing sector is progressing, and the likelihood that, without additional efforts being made, governance efforts will fall behind where they need to be.

As was seen and discussed in Chapter 2, much ‘traditional’ governance – based in policymaking techniques designed for less accelerated technologies – is slow-moving, experiences major bottlenecks during its creation, and frequently lacks nuance. The speed of development of quantum computing, it was argued, is therefore a huge challenge for such *post hoc* forms of governance that require both stabilisation of a deployed technology and, frequently, evidence of harm before legislation or regulation can be put in place. Chapter 2 argued that alternative forms of governance, characterised by adaptivity and agility, are needed – in fact that *anticipatory*, or pre-emptive governance, is indicated as a necessary first step. This call for anticipatory governance, that was seen

in the literature, is also borne out by the number of countries independently investigating such approaches. A recent informal survey of international ‘responsible quantum computing’ researchers reveals that at present Australia, Germany, the Netherlands, the USA, Canada, Germany, and Finland, as well as the UK, are producing research into responsible innovation approaches in quantum computing. High-level investigations in international bodies such as the World Economic Forum are also making progress. As seen in Chapter 6, the speed of development over the last few years has surprised even experienced academics and industrialists in quantum computing fields, and there is unease evident in several areas about a widespread lack of preparedness for what the technology may bring<sup>48</sup>.

It is also important to note that there are still many who are sceptical about the possibilities of quantum computing, who deplore the amount of investment being poured into the technology, and insist that there needs to be much greater management of expectations (for example Gourianov, 2022). They regard the sky-high valuations of quantum computing companies (discussed later in Chapter 6) as deeply concerning evidence of poorly managed communications and an excess of hype. Any anticipatory governance approach may need to recognise that a key role for ‘responsible’ approaches in quantum computing could be to act as a handbrake or a counter to the overpromising that can come from “*quantum enthusiasts*” (152) in policy and marketplaces.

Chapter 2 went on to argue that the creation of a decision-support scaffold, or framework, could assist in such processes of anticipatory governance, and that if such pre-emptive governance processes could be institutionalised across research and – crucially – the commercial sector, then this would serve several purposes. Firstly, it would create a body of evidence that researchers, developers, and spinout companies were indeed making pre-emptive efforts to self-regulate and be reflexive about their activities. Secondly, a framework that could cross the boundaries inside the ‘triple-helix’ could enable anticipatory approaches to be carried across from research into commercial development, as well as help with engagement between research and policy. Thirdly, it would enable like-for-like comparison between projects, and enable collaborative building of such approaches, so that projects and spinout companies were not having to ‘reinvent the wheel’ every time they sought to embed responsible innovation approaches.

---

<sup>48</sup> Eg “*The dawning of the quantum age*”: Financial Times <https://www.ft.com/content/3282f918-59b7-4c3a-8a78-9346d7915159>

The triple-helix approach, blending industry, academic, and policy processes across traditional boundaries, was argued to be necessary because of the close influence each strand has on the others, meaning that an action in, for example, academia, is likely to have effects also in policy and industry.

The case studies were therefore devised around an assessment of the development of quantum computing through an investigation of this triple-helix structure, dividing the quantum computing landscape according to infrastructure locale – policy, research, and industry. It is worth noting in passing that it would also be possible to use different frames for analysis of the landscape – for example an investigation that divided the quantum computing world along the lines of the quantum computing ‘stack’ would demonstrate that some elements of the stack are significantly more advanced than others and thus potentially require an approach that takes account of higher Technology Readiness Levels (TRLs) (Olechowski et al., 2015). Such an alternative framing might require alternative approaches to responsible innovation (eg de Heaver et al., 2020). The ‘triple-helix’ model was used in this case because, as discovered in the literature review, the policy context was seen to be instrumental in shaping the development of the technology and the sector, as well as the governance ecosystem. Additionally, the translation of research into commercial environments was seen as a key point at which ‘responsible’ approaches may be lost. The interplay of these three areas was therefore seen as critical for developing anticipatory governance approaches.

The first case study, described in Chapter 3, therefore addressed questions stemming from the literature review around whether there were particular features exhibited by the quantum computing landscape that differentiated it from other novel technologies, and whether there were more broadly transferable guiding principles that could be gained from a study of how quantum computing is developing and commercialising. Comparisons were drawn with other novel technologies, both those that exhibit significant oversight, such as synthetic biology, and also those that experience a low level of governance, such as machine learning. The latter is, perhaps not coincidentally, experiencing challenges in both public trust and perceptions of trustworthiness. The case-study discussed the value of taking an approach that examined the whole triple-helix, and helped to shape the topics of investigation for the remaining case studies.

Chapter 4, the second case study, investigated the topic of change in organisations, looking at how a relatively novel approach, such as responsible innovation, could become accepted, embedded, and normalised – in this case within research-based institutions. This is particularly valuable to study given the relatively high degree of support that RI receives from some funders and within the EU research landscape. The study examined the question of *How different fields and disciplines are addressing the challenge of incorporating responsible innovation into their work, and what this can tell us about how*

*new norms such as responsible innovation become institutionalised.* This study also investigated *What the institutionalisation (or not) of responsible innovation can tell us about the use of anticipatory governance methodologies in novel technologies.* Given the investigation of organisational change that can illuminate such questions of institutionalisation, the analysis used a ‘learning organisation’ lens to address these questions, finding utility in particular in the ‘granularity’ suggested by Watkins & Marsick, (2019), who recommend a dimensional focus on individuals, small groups, larger communities and finally the institution as a whole. This insight also proved valuable in the remainder of the case studies.

The third case study (Chapter 5) examined the policymaking context for quantum computing, investigating the way in which the sector is being developed, the enablers and constrainers that are either occurring accidentally or being deliberately put in place. Its focus was on answering the main research question around *how policy is developed for novel technologies such as quantum computing.* The study also incorporated questions on both the adaptivity and granularity of policymaking processes, encapsulated as the query *How can policymakers retain the ability to respond to changing circumstances?* The study found that policymaking does not necessarily follow the patterns and processes that are usually claimed for it, but frequently depends on jurisdictional issues, and is also heavily reliant on individual relationships and networks. The study also discussed the problems that can arise from inadequately joined-up responses to ‘enabling’ technologies, as these are generally cross- or multi-jurisdictional – this is a challenge that has persisted for decades and also applies to technologies now considered as fundamental infrastructure, such as electricity and communications. It is possible that quantum computing may one day be regarded in similar terms, but more cohesive policymaking would be valuable in creating a joined-up response.

The final fieldwork study, discussed in Chapter 6, investigated the quantum computing sector. Its questions focused on: *the shaping factors in creating a new sector, and where or how governance is manifested.* It also investigated *How new governance challenges are framed and understood by entrepreneurs in novel technologies.* It was shown how the creation of a new sector is shaped as much by luck as by judgement, and that this also applied to governance. It was serendipitous, for example, that attractive tax incentives happened to be available at the time that potential quantum computing enterprises were looking to spin out of universities. It was also noted that the nascent sector has a large number of university-generated spinouts (perhaps inevitably given the type of expertise that is required for this technology). This may mean that researchers who become Chief Technical Officers or CEOs of new enterprises are to a certain degree familiar with responsible innovation approaches and other anticipatory governance work – however, it was also seen that such approaches may not be prioritised if they are regarded as a distraction from the need to make a return on investment. When

novel technologies are ‘breaking ground’, in areas that lack standards or other regulation, there is a clear gap here – a market sector where there is no relevant applicable governance (other than relating to IP or financial compliance) and no incentive for founders to adopt responsible approaches. The role of public funding was highlighted as a means to enhance oversight of the developing sector, despite some of the risks of such funding that were noted by participants (such as the undermining of competition), as it can offer incentives to founders other than the purely financial.

Chapter 7 drew together what were seen as the most significant challenges from the four case studies, and developed the Beehive Framework as a response to these. It was acknowledged that although it could not address all of them, it could potentially solve for some of the underlying issues, especially those particularly associated with the triple-helix boundaries. Hypothetically it may form a helpful ‘boundary object’ (Star & Griesemer, 1989) that can assist with translation of responsible innovation approaches both within projects and between research and industry. It provides a concrete, actionable, easily useable, collaborative tool that can help researchers and those in the commercial sector utilise responsible innovation approaches in their work.

## 8.2: Themes from the work

### 8.2.1: Granularity: individuals, teams, and institutions

One theme that has recurred throughout the case studies and chapters of this thesis is the importance of using several levels of magnification to examine a given object of study. This was seen initially in the discussion on institutionalisation, per the analysis from Watkins & Marsick, (2019) of how organisations ‘learn’, but has also been demonstrated throughout the case studies, from the need to view responsible innovation within institutions as both macro and micro process, through to the ways in which individual decisions and personalities combine to form networks and mutually supportive groups within policymaking circles. It applies to the ways in which governance processes can be understood, and is an important tool for any responsible innovation approach in general because the use of different perspectives affords the opportunity to review impacts at different levels of granularity, from the national down to the individual. The ability to alter perspective to examine different impacts in this way is a critical tool in the armoury of a responsible innovation practitioner, regardless of their field or discipline.

The most valuable use for this insight around changing the units of analysis, however, is in the challenge – frequently experienced by practitioners – of the ‘location’ of responsible approaches. For example, the problem “*how to carry out responsible innovation in quantum computing*” is too broad and too multi-faceted for any one group to be able to create meaningful and valuable responses –

except potentially at the level of strategic policy-making and extremely broad-brush governance. Altering the 'unit of analysis', or the dimension at which the question is posed – perhaps locating it around an individual project, research question, or engineering challenge – enables more significant and useful responses.

This 'granularity' approach is an important part of the design of the Beehive Framework at the stakeholder-engagement phase. The identification of stakeholders is key to the process of anticipatory governance because of the need for multiple perspectives bringing different understandings and offering unique insights – if the unit of analysis is too broad, then no significant stakeholder-mapping can be carried out, much less useful engagement. Instead, a broad question needs to be re-examined as a series of smaller, more detailed enquiries that can be meaningfully engaged with by stakeholders relevant to that particular enquiry. It is only once the unit of analysis has been properly identified that tools such as the Beehive and the 4P method (Stahl et al., 2017) can be utilised to best effect.

### 8.2.2: 'Messiness' and procedure

Chapter 5 discussed the tension in processes between the 'tidy' approach (procedure) and the on-the-ground reality (messiness). Well-ordered procedures are created to ensure that a given task is carried out in an ordered, logical, and 'known' fashion, but this is often not how the task is carried out in practice, which may be in an *ad hoc*, overlapping, 'organic' manner that depends to a large degree on the personalities and management styles of those carrying it out. It is important to note that there are strong arguments to be made for both approaches: systematisation encourages certainty; allows for concepts such as chain of command; and endeavours to ensure repeatability, while the organic approach can potentially shortcut unnecessary steps; create new connections; and encourage initiative. There are many analogies in management literature for these styles – for example the US military has traditionally followed a systematised approach, also referred to as 'managerial', while the UK army tends to favour the 'initiative' model, known as 'mission command' (Shamir, 2010). It is likely that the tension between these two models is irresolvable, but an important first step is to recognise and acknowledge that processes are not, in fact, following a 'textbook' procedure, and that messy approaches have much to recommend them. Identifying and allowing for messiness can provide a more realistic view than one that tries to ignore such inconvenient elements, and recognises the strengths of such disorder. There are useful guides to be found in the mixed-methods literature (Sakata, 2022; Sanscartier, 2020) as well as in qualitative research methods, but the importance for responsible innovation approaches is in understanding that many questions or issues may initially be perceived to be 'out of scope' and messy during development work. The Beehive allows for the recording of these, as well as the recognition of trade-offs, but the iterative nature of responsible

innovation ensures that messy questions can be refocused – perhaps at a later stage of a design process, or from a different perspective. Such approaches can be particularly helpful in developing adaptive governance for novel technologies.

### 8.2.3: Governance and adaptability

This work has drawn on investigations of governance – how it is created and how it is experienced – in multiple domains, and drawn together threads from numerous sources to create a framework for *anticipatory* governance and a method for its use. This process has also revealed gaps in the governance – and perhaps more critically in the structures for creating governance – that presently exist in relation to quantum computing technologies but also novel technologies in general. To seek to resolve such challenges entirely is beyond the scope of this thesis, but the analysis produced through the research suggests useful directions for investigating and improving governance-development methodologies.

The use of a responsible innovation lens, and the application of varying perspectives as described in 8.2.1, have also illustrated a range of urgent needs in policymaking in general. These have been discussed in detail in the preceding chapters, but include a substantially improved ‘line of sight’ from the places where governance is experienced back up to the places where it is instituted, suggesting a requirement for better connection between rule-makers and rule-takers. There is also a need for more widely available, adaptive forms of governance, potentially following a model of ‘agile’ development and using work such as policy sandpit exercises, or ‘zones of exception’ to enable experimentation and investigation. Such new models for adaptive governance may prove to be crucial for oversight of research and development in a field that is unfolding on an almost daily basis, and frequently in unpredictable directions, which lends a degree of urgency to such investigations.

As discussed in Chapter 2, governance ‘types’ can be conceptualised as being almost on a continuum, or spectrum of governance that shades from international regulation down to individual responsibility – with correspondingly graded degrees of response-time, from slow to rapid. However another framing is that different governance forms can be thought of as building up an overall picture. In this analogy a broad scheme of primary legislation may have finer detail supplied through secondary legislation, regulations, and on an ongoing basis through more responsive governance modes such as industry or professional standards, or case-law. More widespread understanding of different varieties of governance could support rapid investigation and evaluation of possible interventions in a fast-moving field such as quantum computing.

## 8.2.4: Diversity vs homogeneity

The fundamental position of a responsible innovation approach is that inclusivity is not just a ‘good’ in and of itself, but is a critical element of a high-quality process. As will be recalled, *“Responsible innovation should not just welcome diversity; it should nurture it.”* (Stilgoe et al., 2013 p1573), and the most recent developments in responsible innovation argue for the foregrounding of an equality, diversity and inclusion narrative (Smith et al., 2022). This is for numerous reasons, not only the ethical standpoint of including those who may be affected by a technology (*“nothing about us without us”*<sup>49</sup>), but also a pragmatic approach of drawing on all available resources in order to ensure that alternatives have been considered and possible outcomes reflected upon.

This position is linked to the work of Chapters 4 and 5, where it was seen that individuals tend to coalesce into networks and institutions that exhibit considerable inertia, as well as trending towards homogeneity and self-reproduction. However, the perceived need for stakeholder participation in research, development, innovation and policymaking ought to provide a counterbalancing drive towards diversity. These trends are in constant tension but it is clear from the findings and discussion in this research that greater diversity – whilst bringing its own challenges – improves processes, outcomes, and societal engagement. Diversity should therefore be constantly and deliberately sought-out in order to counterbalance conscious and unconscious leanings towards homogeneity.

These questions of diversity, and its impact on processes, also provide an interesting perspective on the *“move fast and break things”* model of ‘disruptive’ technology, where in practical terms this often involves deploying a Minimum Viable Product (MVP) model into the marketplace and then relying on customer and societal feedback to iterate and improve upon it. Insufficiently diverse feedback prior to deployment can prevent developers being able to adequately test a product. However, a responsible innovation approach – that involves stakeholders, anticipation and reflection in the pre-deployment stage – is at the heart of the argument that RI may assist in precluding many of the harms that can arise from launching inadequately-considered or poorly-developed products into society.

## 8.2.5: Institutions and institutionalisation

One of the themes throughout this thesis and research has been the role of institutions in the governance of novel technologies – the power they have to decide on matters as diverse as research priorities, allocation of funding, creation of a new industrial sector, provision of research infrastructure, and many other enablers or constrainers. As was seen in the discussion of the degree

---

<sup>49</sup> [https://en.wikipedia.org/wiki/Nothing\\_About\\_Us\\_Without\\_Us](https://en.wikipedia.org/wiki/Nothing_About_Us_Without_Us)

to which responsible innovation is becoming embedded in funding bodies and universities, and again in the policy case study, the role played by institutions is fundamental to understandings of governance and policymaking. Any question of governance therefore needs to take account of the influence of institutional shaping, trends, and alterations, and such an understanding of the slow shifts that take place – but also the weight of inertia – is useful in analysing meta-trends such the growth of responsible innovation approaches. Using lenses such as Complex Adaptive Systems theory, the work of Thelen on the way institutions can change without appearing to, and ‘learning organisation’ theory, can provide helpful ways to understand the ebb and flow of policy. This also provides an additional layer of understanding to the discussion around granularity in 8.2.1 and the ripple-effects that can be caused by individual decisions – it is necessary to understand both perspectives.

### 8.2.6: Linearity and repetition

As discussed early on in the literature review, processes of responsible innovation are by definition iterative – they require a reassessment of areas that have already been visited. This provides not only the opportunity for deepening insight, but can also act as part of an evaluative mode, understanding the importance of critique and re-appraisal. Even if such processes are not, strictly speaking, a form of assessment, they provide an opportunity for such assessment to take place. These iterative processes are designed-in – an essential part of the reflective element of anticipatory governance. Processes that are designed to be linear, by contrast, often tend to exclude this process of revisiting and re-evaluation, and there is therefore a tension between the practice of responsible innovation and such linearity. This can be seen in the policymaking practices discussed in Chapter 5, and also in the spinning-out of new enterprises seen in Chapter 6 – these are both key processes in the development of novel technologies in general, as well as in quantum computing, and in both instances a set of linear stages is in place, with little opportunity for revisiting. It could be argued that these processes have been ‘streamlined’ – but any winnowing-down of necessary steps such as this in the name of efficiency is misdirected, as it represents a lost opportunity to course-correct. Any deeper embedding of anticipatory governance in policymaking, or in processes such as technology-transfer, would need to address this tension.

## 8.3: Responsible innovation in practice

Throughout this research, which has examined the need for, value of, and deployment of anticipatory governance processes within a particular field, there has been a conscious process of utilising responsible innovation methods within the case studies and applying their principles in practice. Qualitative research necessitates ‘stakeholder engagement’ as an integral part of its process, ‘reflection’ is incorporated as part of the data analysis, and ‘response’ takes place as the result of this

work. Meta-reflection on the process makes clear how many individual choices combine to shape every part of this, from research questions through to choice of interviewees, and ultimately the drawing of conclusions. It could be argued that this makes it even more imperative to find a means of rendering transparent the responsible innovation work carried out in a project, and enabling the comparison of like-for-like work, as well as offering the opportunity to build on, rather than repeat, efforts undertaken elsewhere.

## 8.4: Contribution of the work

### 8.4.1: Contributions to responsible innovation

At the outset of the thesis, it was shown that despite its many advantages, there were many and varied critiques of existing theories and practices of responsible innovation as a form of anticipatory governance. These ranged from the difficulty of applying it consistently; the question of what indicators can or should be used to measure success; the problem of how to decide on ‘enough’ responsible innovation; and many others. The factors discussed above in section 8.2 can contribute to new understandings and approaches for responsible innovation. For example the ‘granularity’ theme addresses one of the challenges of carrying out responsible innovation in a given area by providing support and evidence for the value of breaking the task down into relevant, smaller, stakeholder groups. Using this method, it is not necessary (indeed this researcher would argue that it is not possible) to successfully carry out responsible innovation in “quantum computing”, which – depending on the perspective – can be a sector, a field of research, or several fields of research. Using a variety of lenses, it is clear that this should be broken into smaller units – for example, *responsible innovation in policymaking for capacity and investment in quantum-computing associated research fields*. Selecting the appropriate unit for analysis is critical in order to be able to identify stakeholders and discuss anticipatory issues. The variance in perspective is also necessary to ensure that a holistic overview can be taken where issues overlap (as they will).

Another example of a contribution to the RI discourse comes in the strongly argued case for taking lessons from project management and engineering fields in the approach to RI. The value of documenting processes, stakeholders, reflections and responses has been demonstrated to be critical – not just in terms of the transparency and comparability this approach can provide, but in the way it can support reflective and iterative work, and the effect this then has on assessment and evaluation. This method can also support the retention of tensions and trade-offs throughout any further iterations, as well as ensuring these are carried across to comparable projects, demonstrating the value to stakeholders in participating and contributing their considerations.

### 8.4.2: Contributions to computer science

The literature review demonstrated the degree to which computer science fields (and ICT more generally) have suffered from an absence of sufficiently adaptable governance. Such governance is important as it can shape paths of development in directions that are responsive to stakeholder concerns, as well as helping to support societal trust and acceptance. The right types of governance can also support creativity. Development approaches that incorporate adaptive and anticipatory governance processes can ensure better ‘fit’ with their required specifications by incorporating more development phases prior to deployment. This form of governance means that before a technology creates societal impact (either negative or positive), its development process is able to incorporate more and greater evaluation and assessment loops. The fieldwork has demonstrated the way in which an anticipatory governance framework such as responsible innovation can be highly adaptive to new fields and in new ways, allowing the requisite degree of tailoring to accommodate particular concerns, supporting creativity, and engaging with questions of stakeholder concern and societal fit. This thesis therefore makes a theoretical contribution to computer science through improving both the understanding of the substance of RI and how to apply it in practice.

The most important contribution of this research, however, lies in the Beehive Framework and its associated method. Grounded in the robust qualitative research carried out for the case studies, and evaluated across several different projects, the Framework provides a method for carrying out anticipatory governance in quantum computing technologies (subject to the need for selecting the correct unit of analysis), and is also, in line with other forms of anticipatory governance, adaptable to numerous forms of other technology. As was seen in Chapter 7, its use in several different domains serves to illustrate the amount of flexibility it incorporates. The Framework offered by this thesis therefore provides a methodological contribution to computer science that will provide computer science and associated disciplines with an evaluated methodology for engaging stakeholder input at upstream phases of research in order to improve outcomes for both research and applied work.

### 8.4.3: Contributions to quantum computing

The empirical work that was outlined in Chapter 3 has resulted in the publication of two journal articles, one position paper, and one report, demonstrating the value of this work to the field. The themes of the work described in the thesis have also contributed to a successful funding application for a programme of engagement work aimed at building connections among disparate groups in the quantum computing ecosystem.

The studies described in the thesis have shown that researchers in quantum computing fields frequently struggle with the connection of their work to society. As well as the empirical contribution

described above, the Beehive anticipatory governance framework provides a methodological bridge between academic quantum disciplines and industrial quantum contexts that will assist academic entrepreneurs as they move from research to commercial domains.

An additional challenge for researchers in the field and responsible quantum innovators is that quantum computing continues to expand in myriad ways – developments in the near-term include co-processor models<sup>50</sup>, hybrid quantum-HPC machines<sup>51</sup> and analogue quantum computers<sup>52</sup>, which are already advancing beyond the capabilities of classical computers. These may not represent ‘steps’ in the direction of the large-scale fault-tolerant quantum computers that were generally considered to be the goal, so much as providing systems that offer new affordances in much shorter timeframes than have previously been envisaged. Such developments may provide fresh opportunities in commercial terms, but are likely to create corresponding externalities and potential risks. The methods of stakeholder identification and online, rapid collaboration discussed in the thesis provide one means of responding to these ongoingly-unfolding developments. The thesis contributes deeper understanding of governance-types that can pivot quickly to operationalise anticipation and horizon-scanning, and use of the Beehive framework can enable identification of possible responses and potential interventions.

#### 8.4.4: Contributions to policy

This programme of research identified lacunae in the practical understanding of novel technologies and adaptive governance available to policymakers. It has provided an empirical contribution and substantive knowledge to policymakers at many levels in the form of a deeper understanding of the oversight possibilities of anticipatory governance and ways in which these can be operationalised. The value of this impact has already been demonstrated at a policy event in November 2021 based on the work of Chapter 3. The insights from Chapter 5 have also led to engagement with policymakers through University networks and to events planned to take place in 2023.

#### 8.4.5: Contributions to funders

The work described here will afford funders such as EPSRC insight into the application of responsible approaches throughout their funding remit. The work detailed in Chapter 4 has been written up into a paper published in 2022 in *The Learning Organisation* and will contribute to a report to EPSRC.

---

<sup>50</sup> <https://newsroom.ibm.com/2022-11-09-IBM-Unveils-400-Qubit-Plus-Quantum-Processor-and-Next-Generation-IBM-Quantum-System-Two>

<sup>51</sup> <https://www.fujitsu.com/global/about/resources/news/press-releases/2022/1108-01.html>

<sup>52</sup> <https://phys.org/news/2023-01-analog-quantum-previously-unsolvable-problems.html>

Additionally, the insight gained during the interviews and analysis for that Chapter are currently contributing to a proposed ‘rubric’ for the assessment of responsible innovation in funding proposals.

## 8.5: Recommendations for future work

This thesis has shown that there are significant gaps in the governance of one particular novel technology and proposed a new approach for operationalising better governance in quantum computing fields. However, it would be unrealistic to suppose that in this respect quantum computing is entirely unique – especially given the prior examples of machine learning and others discussed at the outset and throughout the thesis. It would therefore be valuable to expand on the work undertaken here in a way that addresses some of the challenges and gaps identified. This might include:

1. a focus on investigating ways to increase international consensus and co-operation on the need for responsible innovation in quantum computing and other technologies. This might usefully also apply, for example, to the huge global datasets that are frequently used for training machine-learning models.
2. an engagement process that would work on improving cross-departmental communication within policymaking. It was seen in Chapter 5 that there are major challenges around boundary-keeping and boundary-crossing within government and policymaking teams – this is especially the case for particular types of technologies. Clear communication and a common approach is essential to prevent the actions of one team negatively impacting the objectives of another. This would be particularly useful in assessing the impacts of the National Security and Investment Act.
3. research into addressing a key juncture at which responsible innovation approaches may be left behind, the point of spinning out from a university lab into a business. This would need to investigate how ‘responsibility’ concerns could be translated from the particular concerns that apply in the university context, into a different but no less specific set of concerns that may apply in the public context. It is hoped that the Beehive Framework could be of value here, but significantly more support would be needed for entrepreneurs seeking to incorporate responsible innovation approaches.

## 8.6: Final thoughts

During the period of research that has formed the basis of this thesis, broader shifts have been occurring in political, commercial, and societal opinion about the need to take account of concerns often described under the acronym ESG – Environmental, Social, and Governance. Perhaps partly in

response to escalating climate crisis and sustainability pressures, these ESG factors have undergone a shift from being perceived as simply another form of risk management, to being seen as drivers of value (Farmer, 2021). This – critically – includes a change in perspective from investors who understand that businesses adopting such principles are often in a stronger competitive position. Global consulting firms such as McKinsey, Accenture, and EY, have formed groups that specialise in advising clients on their ESG questions, and governments around the world have begun to roll out plans to move to net zero. The development of new business models such as B Corporations<sup>53</sup> also allows commercial organisations to adopt principles other than the purely financial. The principle of the ‘guiding hand’ of markets, which drove much 20<sup>th</sup> century economics, appears to be receding to some degree, perhaps to be replaced by an understanding that governance is an essential requirement for social trust and that new models are needed. This may provide a receptive environment for some of the concerns that have been discussed throughout the thesis, and offer an opportunity for anticipatory governance approaches such as responsible innovation to become the new norm.

Good governance in quantum computing will therefore help to establish its value for society, in a way that ensures citizens can rely on its safety, utility, and ‘fit’ with the kind of future that – collectively – society wants to shape.

---

<sup>53</sup> “companies verified by B Lab to meet high standards of social and environmental performance, transparency, and accountability” <https://bcorporation.uk/b-corp-certification/what-is-a-b-corp/>

## References

- Aaronson, S. (2014). Quantum randomness. *American Scientist*, 102(4), 266–271. <https://doi.org/10.1511/2014.109.266>
- Abbott, A. D. (1988). *The system of professions : an essay on the division of expert labor*. University of Chicago Press.
- Abdulla, A., Vaishnav, P., Sergi, B., & Victor, D. G. (2019). Limits to deployment of nuclear power for decarbonization: Insights from public opinion. *Energy Policy*, 129(March), 1339–1346. <https://doi.org/10.1016/j.enpol.2019.03.039>
- ACM. (1992). ACM Code of Ethics and Professional Conduct. *Communications of the ACM*, 35(5), 94–99. <https://doi.org/10.1145/3274591>
- Aguilar, F. (1967). *Scanning the Business Environment*. Macmillan.
- Andrews, D., Newton, E. J., Adibi, N., Chenadec, J., & Bienge, K. (2021). A circular economy for the data centre industry: Using design methods to address the challenge of whole system sustainability in a unique industrial sector. *Sustainability (Switzerland)*, 13(11). <https://doi.org/10.3390/su13116319>
- Arendt, H. (1963). *Eichmann in Jerusalem: A Report on the Banality of Evil*. Viking Press.
- Aynsley, C. (2020). The Online Harms Bill Cannot Wait. *Infosecurity Magazine*. <https://www.infosecurity-magazine.com/opinions/online-harms-bill-wait/>
- Baker, B. (2022). *Japan Targets 10M Quantum Computer Users by 2030*. Quantum Business News. <https://www.quantumbusinessnews.com/research/japan-targets-10m-quantum-computer-users-by-2030>
- Ball, J. (2019). *The Double Diamond: A universally accepted depiction of the design process*. Design Council. <https://www.designcouncil.org.uk/our-work/news-opinion/double-diamond-universally-accepted-depiction-design-process>
- Banbury, A., Pedell, S., Parkinson, L., & Byrne, L. (2021). Using the Double Diamond model to co-design a dementia caregivers telehealth peer support program. *Journal of Telemedicine and Telecare*, 27(10), 667–673. <https://doi.org/10.1177/1357633X2111048980>
- Banks, C. (2016). Focusing upstream: supporting scholarly communication by academics. *Insights*, 29(1), 37–44. <https://doi.org/10.1629/uksg.292>
- Beck, E. (2002). P for Political: Participation is Not Enough. *Scandinavian Journal of Information Systems*, 14(1), 77–92. <https://aisel.aisnet.org/sjis/vol14/iss1/1>
- BEIS. (2019). *Public attitudes to science 2019*. <https://doi.org/10.1093/spp/4.1.38>
- Bernstein, D. J. (2009). Introduction to Post-Quantum Cryptography. In D. J. Bernstein, J. Buchmann, & E. Dahmen (Eds.), *Post-Quantum Cryptography*. Springer-Verlag Berlin Heidelberg.
- Bernstein, D. J., & Lange, T. (2017). Post-quantum cryptography. *Nature*, 549(7671), 188–194. <https://doi.org/10.1038/nature23461>
- Biamonte, J., Wittek, P., Pancotti, N., Rebentrost, P., Wiebe, N., & Lloyd, S. (2017). Quantum machine learning. In *Nature* (Vol. 549, Issue 7671, pp. 195–202). <https://doi.org/10.1038/nature23474>
- Bittner, A., & Lau, B. (2021). Women-Led Startups Received Just 2.3% of VC Funding in 2020. *Harvard Business Review Digital Articles*, 1–4. <https://hbr.org/2021/02/women-led-startups-received->

- Bjorklund, G. (1988). Planning for uncertainty: a case study. *Technological Forecasting & Social Change*, 33, 119–148.
- Blok, V. (2019). From participation to interruption: toward an ethics of stakeholder engagement, participation and partnership in corporate social responsibility and responsible innovation. In R. von Schomberg & J. Hankins (Eds.), *International Handbook on Responsible Innovation* (pp. 243–257). Edward Elgar. <https://doi.org/10.4337/9781784718862.00024>
- Blue, G., & Dale, J. (2016). Framing and power in public deliberation with climate change: Critical reflections on the role of deliberative practitioners. *Journal of Public Deliberation*, 12(1), 2. <https://doi.org/10.16997/jdd.242>
- Bødker, S., & Kyng, M. (2018). Participatory Design that Matters-Facing the Big Issues. *ACM Transactions on Computer-Human Interaction*, 25(1), 31. <https://doi.org/10.1145/3152421>
- Boenink, M., & Kudina, O. (2020). Values in responsible research and innovation: from entities to practices. *Journal of Responsible Innovation*, 7(3), 450–470. <https://doi.org/10.1080/23299460.2020.1806451>
- Botterill, L. C., & Fenna, A. (2019). *Interrogating Public Policy Theory: A political values perspective*. Edward Elgar. [https://www.google.co.uk/books/edition/Interrogating\\_Public\\_Policy\\_Theory/m46ODwAAQBAJ?hl=en&gbpv=1&dq=Lasswell+policy+making+model&pg=PA77&printsec=frontcover](https://www.google.co.uk/books/edition/Interrogating_Public_Policy_Theory/m46ODwAAQBAJ?hl=en&gbpv=1&dq=Lasswell+policy+making+model&pg=PA77&printsec=frontcover)
- Braun, V., & Clarke, V. (2006). Using thematic analysis in psychology. *Qualitative Research in Psychology*, 3(2), 77–101. <https://doi.org/10.1191/1478088706qp063oa>
- Braun, V., Clarke, V., Hayfield, N., & Terry, G. (2019). Thematic analysis. In *Handbook of Research Methods in Health Social Sciences* (pp. 843–860). Springer Singapore. [https://doi.org/10.1007/978-981-10-5251-4\\_103](https://doi.org/10.1007/978-981-10-5251-4_103)
- Brouneus, F., Lindholm, M., & Bohlin, G. (2019). Telling it straight - a focus group study on narratives affecting public confidence in science. *Journal of Science Communication*, 4(1), 75–84. <https://doi.org/10.22323/2.18050203>
- Burke, D. (2004). GM food and crops: what went wrong in the UK? *EMBO Reports*, 5(5), 432–436. <http://www.ncbi.nlm.nih.gov/pubmed/15184970><http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=PMC1299063>
- Burke, F. (2020). Qubit to get ahead: Germany is racing to catch up with the quantum revolution. *Science Business*. <https://sciencebusiness.net/news/qubit-get-ahead-germany-racing-catch-quantum-revolution>
- Cabinet Office. (2018). *National Security Capability Review*. <http://www.nationalarchives.gov.uk/doc/open-government-licence/oremail:psi@nationalarchives.gsi.gov.uk>
- Cao, Y., Romero, J., & Aspuru-Guzik, A. (2018). Potential of quantum computing for drug discovery. *IBM Journal of Research and Development*, 62(6). <https://doi.org/10.1147/JRD.2018.2888987>
- Chalmers, J., & Armour, M. (2019). The Delphi technique. In *Handbook of Research Methods in Health Social Sciences* (pp. 715–735). [https://doi.org/10.1007/978-981-10-5251-4\\_99](https://doi.org/10.1007/978-981-10-5251-4_99)
- Charmaz, K. (2014). *Constructing grounded theory* (2nd ed.).
- Charmaz, K., & Henwood, K. (2017). Grounded Theory Methods for Qualitative Psychology. In *The*

- SAGE Handbook of Qualitative Research in Psychology*. <https://doi.org/10.4135/9781526405555>
- Chatfield, K., Iatridis, K., Stahl, B. C., & Paspallis, N. (2017). Innovating responsibly in ICT for ageing: Drivers, obstacles and implementation. *Sustainability (Switzerland)*, 9(6). <https://doi.org/10.3390/su9060971>
- Coenen, C., Grinbaum, A., Grunwald, A., Milburn, C., & Vermaas, P. (2022). Quantum Technologies and Society: Towards a Different Spin. *NanoEthics*, 0123456789. <https://doi.org/10.1007/s11569-021-00409-4>
- Coenen, C., & Grunwald, A. (2017). *Responsible research and innovation (RRI) in quantum technology*. 19, 277–294. <https://doi.org/10.1007/s10676-017-9432-6>
- Collingridge, D. (1980). *The Social Control of Technology*. Pinter.
- Conley, S. N., Foley, R. W., Gorman, M. E., Denham, J., & Coleman, K. (2017). Acquisition of T-shaped expertise: an exploratory study. *Social Epistemology*, 31(2), 165–183. <https://doi.org/10.1080/02691728.2016.1249435>
- Corner, A., Venables, D., Spence, A., Poortinga, W., Demski, C., & Pidgeon, N. (2011). Nuclear power, climate change and energy security: Exploring British public attitudes. *Energy Policy*, 39(9), 4823–4833. <https://doi.org/10.1016/j.enpol.2011.06.037>
- Cox, A. M., & Corral, S. (2013). Evolving academic library specialties. *Journal of the American Society for Information Science and Technology*, 64(8), 1526–1542. <https://doi.org/10.1002/ASI.22847>
- Coyle, D. (2018). *Not the new oil: the real value of data*. Bennett Institute for Public Policy. <https://www.bennettinstitute.cam.ac.uk/blog/new-oil-real-value-data/>
- Cross, R. (1998). Managing for knowledge: managing for growth. *Knowledge Management*, 1(3), 9–13.
- da Silva, L. M., Bitencourt, C. C., Faccin, K., & Iakovleva, T. (2019). The role of stakeholders in the context of responsible innovation: A meta-synthesis. *Sustainability (Switzerland)*, 11(6). <https://doi.org/10.3390/su11061766>
- Dalkey, N., & Helmer, O. (1963). An Experimental Application of the DELPHI Method to the Use of Experts. *Management Science*, 9(3), 458–467. <https://doi.org/10.1287/mnsc.9.3.458>
- Dalquist, S. (2004). A Chronology of Public Opinion on Nuclear Power in the United States and United Kingdom. In *Timeline* (Issue April).
- Davies, S. R., & Selin, C. (2012). Energy futures: Five dilemmas of the practice of anticipatory governance. *Environmental Communication*, 6(1), 119–136. <https://doi.org/10.1080/17524032.2011.644632>
- Day, G. S. (1990). *Market Driven Strategy: Processes for Creating Value*. The Free Press. <https://doi.org/10.2307/1251961>
- de Heaver, M., Jirotko, M., Nulli, M., Stahl, B. C., & Ten Holter, C. (2020). Assessment of Responsible Innovation. In E. Yaghmaei & I. van de Poel (Eds.), *Assessment of Responsible Innovation: Methods and Practices*. Routledge. <https://doi.org/10.4324/9780429298998>
- de Loë, R. C., Melnychuk, N., Murray, D., & Plummer, R. (2016). Advancing the State of Policy Delphi Practice: A Systematic Review Evaluating Methodological Evolution, Innovation, and Opportunities. *Technological Forecasting and Social Change*, 104, 78–88. <https://doi.org/10.1016/j.techfore.2015.12.009>
- de Saille, S. (2015). Innovating innovation policy: the emergence of ‘Responsible Research and Innovation.’ *Journal of Responsible Innovation*, 2(2), 152–168.

<https://doi.org/10.1080/23299460.2015.1045280>

- De Wolf, R. (2017). The potential impact of quantum computers on society. *Ethics and Information Technology*, 19, 271–276. <https://doi.org/10.1007/s10676-017-9439-z>
- Delgado, A., & Åm, H. (2018). *Experiments in interdisciplinarity: Responsible research and innovation and the public good*. <https://doi.org/https://doi.org/10.1371/journal.pbio.2003921>
- Department for Business Energy and Industrial Strategy. (2021). *UK innovation strategy: Leading the future by creating it*.
- Dihal, K. (2017). The stories of quantum physics: quantum physics in literature and popular science, 1900-present. *PQDT - UK & Ireland*. <https://search.proquest.com/docview/2189047121?accountid=8113>
- DiVincenzo, D. P. (2017). Scientists and citizens: getting to quantum technologies. *Ethics and Information Technology*, 19(4), 247–251. <https://doi.org/10.1007/s10676-017-9435-3>
- Dreyer, M., Chefneux, L., Goldberg, A., von Heimburg, J., Patrignani, N., Schofield, M., & Shilling, C. (2017). Responsible innovation: A complementary view from industry with proposals for bridging different perspectives. *Sustainability (Switzerland)*, 9(10), 1–25. <https://doi.org/10.3390/su9101719>
- Dyakonov, M. I. (2020). *Will We Ever Have a Quantum Computer?* Springer Nature. <http://www.springer.com/series/8902>
- Easterby-Smith, M., Crossan, M., & Nicolini, D. (2000). Organizational learning: Debates past, present and future. *Journal of Management Studies*, 37(6), 783–796. <https://doi.org/10.1111/1467-6486.00203>
- Ehn, P. (2002). Neither Bauhäusler nor nerd: Educating the interaction designer. *Proceedings of the Conference on Designing Interactive Systems: Processes, Practices, Methods, and Techniques, DIS*, 19–23. <https://doi.org/10.1145/778712.778717>
- EPSRC. (2014). *National Network of Quantum Technology Hubs*. <https://epsrc.ukri.org/research/ourportfolio/themes/quantumtech/activities/nationalnetworkofquantumtechnologyhubs/>
- EPSRC. (2017). *Quantum Technologies Public Dialogue Report*. [https://www.nqit.ox.ac.uk/sites/www.nqit.ox.ac.uk/files/2018-07/Quantum Technologies Public Dialogue Full Report\\_0.pdf](https://www.nqit.ox.ac.uk/sites/www.nqit.ox.ac.uk/files/2018-07/Quantum%20Technologies%20Public%20Dialogue%20Full%20Report_0.pdf)
- EPSRC. (2018). *2018 CDT Exercise - EPSRC website*. <https://epsrc.ukri.org/skills/students/centres/2018-cdt-exercise/>
- EPSRC. (2020). *UK National Quantum Technologies Programme*. <https://uknqt.epsrc.ac.uk/>
- European Commission. (2011). *A renewed EU strategy 2011-14 for Corporate Social Responsibility* (Vol. 16, Issue 22).
- European Commission. (2020). *Science with and for Society (SWAFS)*. EC Website. <https://ec.europa.eu/research/swafs/index.cfm>
- European Commission. (2022). *European Chips Act*. Priorities 2019-2024. [https://ec.europa.eu/info/strategy/priorities-2019-2024/europe-fit-digital-age/european-chips-act\\_en](https://ec.europa.eu/info/strategy/priorities-2019-2024/europe-fit-digital-age/european-chips-act_en)
- European Environment Agency. (2001). Late lessons from early warnings: the precautionary principle 1896-2000. In *Environmental Issue Report no 22*.

- Farmer, R. (2021). *How to respond to rising investor expectations around ESG*.  
[https://www.ey.com/en\\_uk/assurance/how-to-respond-to-rising-investor-expectations-around-esg](https://www.ey.com/en_uk/assurance/how-to-respond-to-rising-investor-expectations-around-esg)
- Farthing, M., Horton, R., & Smith, R. (2000). Research misconduct: Britain's failure to act. *British Medical Journal*, *321*, 1485. <https://doi.org/10.1136/bmj.1.3101.777>
- Feynman, R. (1988). Mr Feynman Goes To Washington: Investigating the Space Shuttle Challenger Disaster. In *What do you care what other people think?* Penguin.
- Feynman, R. P. (1982). Simulating physics with computers. *International Journal of Theoretical Physics*, *21*(6–7), 467–488. <https://doi.org/10.1007/BF02650179>
- Fingleton, J. (2021). *Britain must rethink its 'national security' law*. Financial Times.  
<https://www.ft.com/content/8606231f-5660-4553-a263-31b16faf4071>
- Forrester, J. W. (1971). Counterintuitive behavior of social systems. *Theory and Decision*, *2*(2), 109–140. <https://doi.org/10.1007/BF00148991>
- Forsberg, E. M., Quaglio, G. L., O’Kane, H., Karapiperis, T., Van Woensel, L., & Arnaldi, S. (2015). Assessment of science and technologies: Advising for and with responsibility. *Technology in Society*, *42*, 21–27. <https://doi.org/10.1016/j.techsoc.2014.12.004>
- Frow, E., & Calvert, J. (2013). Opening up the future(s) of synthetic biology. *Futures*, *48*, 32–43. <https://doi.org/10.1016/j.futures.2013.03.001>
- Funtowicz, S. O., & Ravetz, J. R. (1993). Science for the post-normal age. *Futures*, *25*(7), 739–755. [https://doi.org/10.1016/0016-3287\(93\)90022-L](https://doi.org/10.1016/0016-3287(93)90022-L)
- Garvie, C. (2019). *Statement of Clare Garvie Senior Associate , Center on Privacy & Technology at Georgetown Law Before the U . S . House of Representatives Committee on Oversight and Reform Hearing on Facial Recognition Technology ( Part 1 ): Its Impact on Our Civil Rights*.  
<https://docs.house.gov/meetings/GO/GO00/20190522/109521/HHRG-116-GO00-Wstate-GarvieC-20190522.pdf>
- Genus, A., & Iskandarova, M. (2018). Responsible innovation: its institutionalisation and a critique. *Technological Forecasting and Social Change*, *128*(September 2017), 1–9. <https://doi.org/10.1016/j.techfore.2017.09.029>
- Genus, A., & Stirling, A. (2018). Collingridge and the dilemma of control: Towards responsible and accountable innovation. *Research Policy*. <https://doi.org/10.1016/j.respol.2017.09.012>
- Gerber, A. (2018). RRI: How to “mainstream” the “upstream” engagement. *Journal of Science Communication*, *17*(3). <https://doi.org/10.22323/2.17030306>
- Gibney, E. (2019). Quantum gold rush: the private funding pouring into quantum start-ups. *Nature*, *574*(7776), 22–24. <https://doi.org/10.1038/d41586-019-02935-4>
- Gordon, A. V., Ramic, M., Rohrbeck, R., & Spaniol, M. J. (2020). 50 Years of Corporate and Organizational Foresight: Looking Back and Going Forward. *Technological Forecasting & Social Change*, *accepted*(January), 119966. <https://doi.org/10.1016/j.techfore.2020.119966>
- Gourianov, N. (2022). *The quantum computing bubble*. Financial Times.  
<https://www.ft.com/content/6d2e34ab-f9fd-4041-8a96-91802bab7765?desktop=true&segmentId=d8d3e364-5197-20eb-17cf-2437841d178a#myft:notification:instant-email:content>
- Gray, B., & Stites, J. (2013). Sustainability through partnerships: Capitalizing on Collaboration.

*Wageningenportals.Nl*, 109. <http://www.theguardian.com/sustainable-business/paul-polman-unilever-sustainable-living-plan>

- Grimpe, B., Hartswood, M., & Jirotko, M. (2014). Towards a closer dialogue between policy and practice: Responsible design in HCI. *Conference on Human Factors in Computing Systems - Proceedings*, 2965–2974. <https://doi.org/10.1145/2556288.2557364>
- Grimpe, B., Stahl, B. C., Ten Holter, C., Inglesant, P., Eden, G., Patel, M., & Jirotko, M. (2020). From collaborative to institutional reflexivity: Calibrating responsibility in the funding process. *Science and Public Policy*, 47(5), 720–732. <https://doi.org/10.1093/scipol/scaa038>
- Grinbaum, A. (2017). Narratives of quantum theory in the age of quantum technologies. *Ethics and Information Technology*, 19(4), 295–306. <https://doi.org/10.1007/s10676-017-9424-6>
- Grinbaum, A., & Groves, C. (2013). What Is “Responsible” about Responsible Innovation? Understanding the Ethical Issues. In R Owen, J. Bessant, & M. Heintz (Eds.), *Responsible Innovation: Managing the Responsible Emergence of Science and Innovation in Society* (pp. 119–142). Wiley Online Library. <https://doi.org/10.1002/9781118551424.ch7>
- Grove-White, R., Macnaghten, P., & Wynne, B. (2000). *Wising Up: the public and new technologies*. July 2014, 1–48.
- Grunwald, A. (2014). Technology assessment for responsible innovation. In *Responsible Innovation 1: Innovative Solutions for Global Issues* (pp. 15–31). [https://doi.org/10.1007/978-94-017-8956-1\\_2](https://doi.org/10.1007/978-94-017-8956-1_2)
- Guest, D. (1991). The Hunt is on for the Renaissance Man of Computing. *The Independent*.
- Guston, D. H. (2014). Understanding “anticipatory governance.” *Social Studies of Science*, 44(2), 218–242. <https://doi.org/10.1177/0306312713508669>
- Habermas, J. (2008). *Between Naturalism and Religion: Philosophical Essays*. Wiley.
- Habermas, J. (2011). “The Political”: The Rational Meaning of a Questionable Inheritance of Political Theology. In *The Power of Religion in the Public Sphere* (pp. 15–33). Columbia University Press.
- Handy, C. B. (1976). So You Want to Change Your Organisation? Then First Identify its Culture. *Management Learning*, 7(2), 67–84. <https://doi.org/10.1177/135050767600700202>
- Hansen, J. Ø., Jensen, A., & Nguyen, N. (2020). The responsible learning organization: Can Senge (1990) teach organizations how to become responsible innovators? *Learning Organization*, 27(1), 65–74. <https://doi.org/10.1108/TLO-11-2019-0164>
- Hansen, N. B., Dindler, C., Halskov, K., Iversen, O. S., Bossen, C., Basballe, D. A., & Schouten, B. (2019). How participatory design works: Mechanisms and effects. *ACM International Conference Proceeding Series*, 12, 30–41. <https://doi.org/10.1145/3369457.3369460>
- Harrison, R. (1972). *Understanding your organisation’s character*. Harvard Business Review.
- Hart, H. L. A. (1968). *Punishment and Responsibility*. Oxford University Press.
- Hekler, A., Utikal, J. S., Enk, A. H., Hauschild, A., Weichenthal, M., Maron, R. C., Berking, C., Haferkamp, S., Klode, J., Schadendorf, D., Schilling, B., Holland-Letz, T., Izar, B., von Kalle, C., Fröhling, S., Brinker, T. J., Schmitt, L., Peitsch, W. K., Hoffmann, F., ... Thiem, A. (2019). Superior skin cancer classification by the combination of human and artificial intelligence. *European Journal of Cancer*, 120, 114–121. <https://doi.org/10.1016/j.ejca.2019.07.019>
- Hemphill, T. A. (2020). “The innovation governance dilemma: Alternatives to the precautionary principle.” *Technology in Society*, 63(August). <https://doi.org/10.1016/j.techsoc.2020.101381>

- Hern, A. (2022). *Microsoft limits access to facial recognition tool in AI ethics overhaul*. The Guardian. <https://www.theguardian.com/technology/2022/jun/22/microsoft-limits-access-to-facial-recognition-tool-in-ai-ethics-overhaul>
- Herring, C. (2009). Does diversity pay?: Race, gender, and the business case for diversity. *American Sociological Review*, 74(2), 208–224. <https://doi.org/10.1177/000312240907400203>
- Hillier, J. (2020). *CUWiP UK 2020 Evaluation Report*. <https://www.hesa.ac.uk/data-and-publications/cuwip-uk-2020-evaluation-report>
- Public Sector (Social Value) Act, (2012). <https://www.legislation.gov.uk/ukpga/2012/3/enacted>
- HM Government. (2021). *Global Britain in a competitive age: The Integrated Review of Security, Defence, Development and Foreign Policy*. [www.gov.uk/official-documents](http://www.gov.uk/official-documents)
- HM Treasury. (2022). *The Green Book: Central Government Guidance Appraisal and Evaluation*. [www.gov.uk/official-documents](http://www.gov.uk/official-documents)
- Hoe, S. L. (2019). The Topicality of the Learning Organization: is the Concept Still Relevant Today? In *The Oxford Handbook of the Learning Organization* (pp. 18–32). Oxford University Press (OUP). <https://doi.org/10.1093/oxfordhb/9780198832355.013.2>
- Holland, J. (1992). Complex adaptive systems. *Daedalus*, 121(1), 17–30. [https://doi.org/10.1007/3-540-26869-3\\_16](https://doi.org/10.1007/3-540-26869-3_16)
- Holland, S. J. (2018). *The Time To Talk About The Ethics Of Quantum Computing Is Now*. QCEthics.Org. [https://img1.wsimg.com/blobby/go/a44e2290-7b86-4e61-b364-de5d16d93cbb/downloads/The Time To Talk About The Ethics Of QC Is Now.pdf?ver=1580511922757](https://img1.wsimg.com/blobby/go/a44e2290-7b86-4e61-b364-de5d16d93cbb/downloads/The%20Time%20To%20Talk%20About%20The%20Ethics%20Of%20QC%20Is%20Now.pdf?ver=1580511922757)
- House of Commons Science and Technology Committee. (2019). *Quantum technologies* (Issue February).
- House of Lords Science & Technology Select Committee. (2000). *Science and Society*. <https://publications.parliament.uk/pa/ld199900/ldselect/ldsctech/38/3801.htm>
- House of Lords Science & Technology Select Committee. (2022). *“Science and technology superpower”: more than a slogan?*
- Hsu, S. W., & Lamb, P. (2020). Still in search of learning organization?: Towards a radical account of The Fifth Discipline: The Art and Practice of the Learning Organization. *Learning Organization*, 27(1), 31–41. <https://doi.org/10.1108/TLO-09-2019-0142>
- IEEE. (2019). *IEEE Quantum Initiative Support for Standards* (pp. 1–2). <https://quantum.ieee.org/images/files/pdf/ieee-support-for-standards.pdf>
- Inglesant, P., Hartswood, M., & Jirotko, M. (2016). *Thinking Ahead to a World with Quantum Computers The Landscape of Responsible Research and Innovation in Quantum Computing*.
- Inglesant, P., Ten Holter, C., Jirotko, M., & Williams, R. (2021). Asleep at the wheel? Responsible Innovation in quantum computing. *Technology Analysis and Strategic Management*, 0(0), 1–13. <https://doi.org/10.1080/09537325.2021.1988557>
- Inzelt, A., & Csonka, L. (2017). The approach of the business sector to responsible research and innovation (RRI). *Foresight and STI Governance*, 11(4), 63–73. <https://doi.org/10.17323/2500-2597.2017.4.63.73>
- Jasanoff, S. (2006). *States of Knowledge: The Co-Production of Science and the Social Order*. Routledge.
- Jirotko, M., Grimpe, B., Stahl, B., Eden, G., & Hartswood, M. (2017). Responsible research and innovation in the digital age. *Communications of the ACM*, 60(5), 62–68.

<https://doi.org/10.1145/3064940>

- Johnson, W. G. (2019a). Governance Tools For The Second Quantum Revolution. *Jurimetrics*, 59(4), 487–521. <https://search.proquest.com/docview/2308458004?fromopenview=true&pq-origsite=gscholar>
- Johnson, W. G. (2019b). *Quantum Supremacy and the Regulation of Quantum Technologies*. The Regulatory Review. <https://www.theregview.org/2019/12/30/johnson-quantum-supremacy-regulation-quantum-technologies/>
- Kaplan, R. S., & Norton, D. P. (1992). The balanced scorecard - measures that drive performance. *Harvard Business Review*, 70(1), 71–79. <https://hbr.org/1992/01/the-balanced-scorecard-measures-that-drive-performance-2>
- Khan, I. (2021). *Will Quantum Computers Truly Serve Humanity?* ScientificAmerican.Com. <https://www.scientificamerican.com/article/will-quantum-computers-truly-serve-humanity/>
- Koberg, C. S., & Chusmir, L. H. (1987). Organizational culture relationships with creativity and other job-related variables. *Journal of Business Research*, 15, 397–409.
- Konig, A. M. (2022). *Weathering the First Quantum Short*. Weekly Quantum World Detangled. <https://quantumcomputing.substack.com/p/weathering-the-first-quantum-short?s=r>
- Krabbenborg, L., & Mulder, H. A. J. (2015). Upstream Public Engagement in Nanotechnology: Constraints and Opportunities. In *Science Communication* (Vol. 37, Issue 4). <https://doi.org/10.1177/1075547015588601>
- Krishna, A. (2020). *IBM CEO's Letter to Congress on Racial Justice Reform*. <https://www.ibm.com/blogs/policy/facial-recognition-sunset-racial-justice-reforms/>
- Kuhlmann, S., Stegmaier, P., & Konrad, K. (2019). The tentative governance of emerging science and technology—A conceptual introduction. *Research Policy*, 48(5), 1091–1097. <https://doi.org/10.1016/j.respol.2019.01.006>
- Lasswell, H. D. (1956). The Political Science of Science: An Inquiry into the Possible Reconciliation of Mastery and Freedom. *American Political Science Review*, 50(4), 961–979. <https://doi.org/10.2307/1951330>
- Lehoux, P., Miller, F. A., & Williams-Jones, B. (2020). Anticipatory governance and moral imagination: Methodological insights from a scenario-based public deliberation study. *Technological Forecasting and Social Change*, 151, 1–13. <https://doi.org/10.1016/j.techfore.2019.119800>
- Lester, S. (2015). On professions and being professional. In *Stan Lester Developments* (Vol. 2007, Issue June, pp. 1–12). <http://devmts.org.uk/profnal.pdf>
- Levinthal, D. A. (2018). From Strategy to Strategic Organization. *Advances in Strategic Management*. <https://doi.org/10.1108/S0742-332220180000039005>
- Liberatore, A., & Funtowicz, S. (2003). Introduction “Democratizing” expertise, “expertizing” democracy: what does this mean, and why bother? In *Science and Public Policy* (Vol. 30). <https://academic.oup.com/spp/article/30/3/146/1628309>
- Lindner, R., Kuhlmann, S., Randles, S., Bedsted, B., Gorgoni, G., Griessler, E., Loconto, A., & Mejlgaard, N. (2016). *Navigating towards shared responsibility in research and innovation: Approach, process and results of the Res-AGorA project*. Fraunhofer Institute for Systems and Innovation Research ISI.
- Long, J. (2014). *Process Modeling Style*. Morgan Kaufmann.

<https://doi.org/https://doi.org/10.1016/C2013-0-13505-2>

- Lopez-Perez, L., & Olvera-Lobo, M. D. (2018). Public engagement in science via web 2.0 technologies: Evaluation criteria validated using the Delphi method. *Journal of Science Communication, 17*(2), 1–21. <https://doi.org/10.22323/2.17020208>
- Lu, D. (2019). ... but what, actually, is a quantum computer. *New Scientist*. [https://doi.org/https://doi.org/10.1016/S0262-4079\(19\)32038-X](https://doi.org/https://doi.org/10.1016/S0262-4079(19)32038-X)
- Lubberink, R., Blok, V., Ophem, J. van, & Omta, O. (2017). Lessons for responsible innovation in the business context: A systematic literature review of responsible, social and sustainable innovation practices. In *Sustainability (Switzerland)* (Vol. 9, Issue 5, p. 721). <https://doi.org/10.3390/su9050721>
- Lynch, M. (2000). Against Reflexivity as an Academic Virtue and Source of Privileged Knowledge. *Theory, Culture & Society, 17*(3), 26–54. <https://doi.org/10.1177/02632760022051202>
- Macnaghten, P. (2020). Towards an anticipatory public engagement methodology: deliberative experiments in the assembly of possible worlds using focus groups. *Qualitative Research*. <https://doi.org/10.1177/1468794120919096>
- Maeda, K., Sasaki, T., & Koashi, M. (2019). Repeaterless quantum key distribution with efficient finite-key analysis overcoming the rate-distance limit. *Nature Communications, 10*(1). <https://doi.org/10.1038/s41467-019-11008-z>
- Mahoney, J., & Thelen, K. (2010). A theory of gradual institutional change. In *Explaining Institutional Change: Ambiguity, Agency, and Power* (pp. 1–37). <https://doi.org/10.1017/CBO9780511806414.003>
- Malecki, C. (2018). *Corporate Social Responsibility: perspectives for sustainable corporate governance*. Edward Elgar.
- Marquardt, M., & Reynolds, A. (1994). *The Global Learning Organization: Gaining Competitive Advantage Through Continuous Learning*. McGraw Hill.
- Martinuzzi, A., Blok, V., Brem, A., Stahl, B., & Schönherr, N. (2018). Responsible Research and Innovation in industry-challenges, insights and perspectives. *Sustainability (Switzerland), 10*, 702. <https://doi.org/10.3390/su10030702>
- McKinsey & Company. (2021). *Quantum computing: An emerging ecosystem and industry use cases*.
- McKinsey & Company. (2022). *Quantum Technology Monitor* (Issue June).
- Mertens, M. (2018). Liminal innovation practices: questioning three common assumptions in responsible innovation. *Journal of Responsible Innovation, 5*(3), 280–298. <https://doi.org/10.1080/23299460.2018.1495031>
- Mittelstadt, B. D., Stahl, B. C., & Fairweather, N. Ben. (2015). How to Shape a Better Future? Epistemic Difficulties for Ethical Assessment and Anticipatory Governance of Emerging Technologies. *Ethical Theory and Moral Practice, 18*(5), 1027–1047. <https://doi.org/10.1007/s10677-015-9582-8>
- Molendowski, E., & Żmuda, M. (2013). Changes In Competitiveness Among The Visegrad Countries After Accession To The European Union: A Comparative Analysis Based On A Generalized Double Diamond Model. *Comparative Economic Research, 16*(4), 121–153. <https://doi.org/10.2478/cer-2013-0031>
- Möller, M., & Vuik, C. (2017). On the impact of quantum computing technology on future

- developments in high-performance scientific computing. *Ethics and Information Technology*, 19(4), 253–269. <https://doi.org/10.1007/s10676-017-9438-0>
- Moon, J. A. (2004). *Reflection in Learning and Professional Development: Theory and Practice*. Routledge Falmer.
- Murgia, M. (2021). *Quantum computing comes out of shadows into public markets*. Financial Times. <https://www.ft.com/content/72b71d34-5804-4346-a10d-66f583e6300e>
- Murgia, M., & Waters, R. (2019). Google claims to have reached quantum supremacy. *Financial Times*. <https://www.ft.com/content/b9bb4e54-dbc1-11e9-8f9b-77216ebe1f17>
- Murphy, P. (2010). Nanotechnology: Public Engagement with Health, Environmental and Social Issues. STRIVE Report (2007-FS-EH-1-M5). In *STRIVE Report Series No.61*. www.epa.ie
- Nandonde, F. A. (2019). A PESTLE analysis of international retailing in the East African Community. *Global Business and Organizational Excellence*, 38(4), 54–61. <https://doi.org/10.1002/joe.21935>
- Nazarko, L., & Melnikas, B. (2019). Operationalising Responsible Research and Innovation - tools for enterprises. *Engineering Management in Production and Services*, 11(3), 21–28. <https://doi.org/10.2478/emj-2019-0017>
- Nielsen, M. V. (2016). The concept of responsiveness in the governance of research and innovation. *Science and Public Policy*, scv078. <https://doi.org/10.1093/scipol/scv078>
- Nordmann, A. (2014). Responsible innovation, the art and craft of anticipation. *Journal of Responsible Innovation*, 1(1), 87–98. <https://doi.org/10.1080/23299460.2014.882064>
- Nowotny, H. (2003). Democratising expertise and socially robust knowledge. *Science and Public Policy*, 30(3), 151–156. <https://doi.org/10.3152/147154303781780461>
- Olechowski, A., Eppinger, S. D., & Joglekar, N. (2015). Technology readiness levels at 40: A study of state-of-the-art use, challenges, and opportunities. *Portland International Conference on Management of Engineering and Technology, 2015-Septe*, 2084–2094. <https://doi.org/10.1109/PICMET.2015.7273196>
- Owen, Richard, 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/https://doi.org/10.1093/scipol/scs093>
- Owen, Richard, Pansera, M., Macnaghten, P., & Randles, S. (2021). Organisational institutionalisation of responsible innovation. *Research Policy*, 50(1), 104132. <https://doi.org/10.1016/j.respol.2020.104132>
- Owen, Richard, Stilgoe, J., Macnaghten, P., Gorman, M., Fisher, E., & Guston, D. (2013). A Framework for Responsible Innovation. In Richard Owen, J. Bessant, & M. Heintz (Eds.), *Responsible Innovation: Managing the Responsible Emergence of Science and Innovation in Society* (pp. 27–50). Wiley. <https://doi.org/10.1002/9781118551424.ch2>
- Owen, Richard, von Schomberg, R., & Macnaghten, P. (2021). An unfinished journey? Reflections on a decade of responsible research and innovation. *Journal of Responsible Innovation*, 1–17. <https://doi.org/10.1080/23299460.2021.1948789>
- Ozmen Garibay, O., Winslow, B., Andolina, S., Antona, M., Bodenschatz, A., Coursaris, C., Falco, G., Fiore, S. M., Garibay, I., Grieman, K., Havens, J. C., Jirotko, M., Kacorri, H., Karwowski, W., Kider, J., Konstan, J., Koon, S., Lopez-Gonzalez, M., Maifeld-Carucci, I., ... Xu, W. (2023). Six Human-Centered Artificial Intelligence Grand Challenges. *International Journal of Human-Computer Interaction*. <https://doi.org/10.1080/10447318.2022.2153320>

- Pagel, C. (2022). *Physics: Do girls avoid it because it's too hard?* BBC Science Focus Magazine. <https://www.sciencefocus.com/news/physics-do-girls-avoid-it-because-its-too-hard/>
- Pandza, K., & Ellwood, P. (2013). Strategic and ethical foundations for responsible innovation. *Research Policy*, 42(5), 1112–1125. <https://doi.org/10.1016/j.respol.2013.02.007>
- Pansera, M., & Owen, R. (2018). *Report from national case study: United Kingdom*. [www.rri-practice.eu](http://www.rri-practice.eu)
- Pansera, M., Owen, R., Meacham, D., & Kuh, V. (2020). Embedding responsible innovation within synthetic biology research and innovation: insights from a UK multi-disciplinary research centre. *Journal of Responsible Innovation*, 384–409. <https://doi.org/10.1080/23299460.2020.1785678>
- Pellizzoni, L. (2004). Responsibility and environmental governance. *Environmental Politics*, 13(3), 541–565. <https://doi.org/https://doi.org/10.1080/0964401042000229034>
- Pereira, Â. G., & Saltelli, A. (2017). Post-normal institutional identities: Quality assurance, reflexivity and ethos of care. *Futures*, 91, 53–61. <https://doi.org/10.1016/j.futures.2016.11.009>
- Perrier, E. (2021). *Ethical Quantum Computing: A Roadmap*. [www.aaai.org](http://www.aaai.org)
- Pilemalm, S. (2018). Participatory design in emerging civic engagement initiatives in the new public sector: Applying PD concepts in resource-scarce organizations. *ACM Transactions on Computer-Human Interaction*, 25(1). <https://doi.org/10.1145/3152420>
- Porcari, A., Pimponi, D., Borsella, E., Klaassen, P., Maia, M., & Mantovani, E. (2021). Supporting RRI uptake in industry: A qualitative and multi- criteria approach to analysing the costs and benefits of implementation. In E. Yaghmaei & I. van de Poel (Eds.), *Assessment of Responsible Innovation: Methods and Practices*. Routledge.
- Porter, M. J. (1979). How Competitive Forces Shape Strategy. *Harvard Business Review*, March. <https://hbr.org/1979/03/how-competitive-forces-shape-strategy>
- Powell, A., Shennan, S., & Thomas, M. G. (2009). Late Pleistocene demography and the appearance of modern human behavior. *Science*, 324(5932), 1298–1301. <https://doi.org/10.1126/science.1170165>
- Preskill, J. (2018). *Quantum Computing in the NISQ era and beyond*. <https://doi.org/10.22331/q-2018-08-06-79>
- Quantum Tech. (2021). *The Road to Quantum Supremacy*. <https://doi.org/10.1126/science.1109554>
- Randles, S., & Laasch, O. (2016). Theorising the Normative Business Model. *Organization and Environment*, 29(1), 53–73. <https://doi.org/10.1177/1086026615592934>
- Randles, S., Laredo, P., Loconto, A., Walhout, B., & Lindner, R. (2016). Framings and frameworks: six grand narratives of de facto RRI. In *Navigating Towards Shared Responsibility in Research and Innovation: Approach, Process and Results of the Res-AGorA Project* (pp. 31–36). <https://hal.archives-ouvertes.fr/hal-01320462>
- Rawluk, A., Beilin, R., Bender, H., & Ford, R. (2019). Practices in social ecological research: Interdisciplinary collaboration in “adaptive doing.” In *Practices in Social Ecological Research: Interdisciplinary Collaboration in “Adaptive Doing.”* Palgrave Macmillan. <https://doi.org/10.1007/978-3-030-31189-6>
- Reece, R. P., & Stahl, B. C. (2015). The professionalisation of information security: Perspectives of UK practitioners. *Computers and Security*, 48, 182–195. <https://doi.org/10.1016/j.cose.2014.10.007>
- Reeves, M., & Levin, S. (2021). Chapter 2. Think Biologically: Messy Management for a Complex World.

- In *Mastering the Science of Organizational Change* (Issue July 1953, pp. 13–26).  
<https://doi.org/10.1515/9783110697834-002>
- Reeves, S. (2012). Envisioning ubiquitous computing. *Conference on Human Factors in Computing Systems - Proceedings*, 1573–1582. <https://doi.org/10.1145/2207676.2208278>
- Reiher, M., Wiebe, N., Svore, K. M., Wecker, D., & Troyer, M. (2017). Elucidating reaction mechanisms on quantum computers. *Proceedings of the National Academy of Sciences of the United States of America*, 114(29), 7555–7560. <https://doi.org/10.1073/pnas.1619152114>
- Reynolds, C. (2020). UK Increases Quantum Investment with Funding of 38 New Projects. *Computer Business Review*. <https://www.cbronline.com/news/quantum-investment>
- Ribeiro, B., Bengtsson, L., Benneworth, P., Bühner, S., Castro-Martínez, E., Hansen, M., Jarmai, K., Lindner, R., Olmos-Peñuela, J., Ott, C., & Shapira, P. (2018). Introducing the dilemma of societal alignment for inclusive and responsible research and innovation. *Journal of Responsible Innovation*, 5(3), 316–331. <https://doi.org/10.1080/23299460.2018.1495033>
- Ribeiro, B. E., Smith, R. D. J., & Millar, K. (2017). A Mobilising Concept? Unpacking Academic Representations of Responsible Research and Innovation. *Science and Engineering Ethics*, 23(1), 81–103. <https://doi.org/10.1007/s11948-016-9761-6>
- Ribeiro, B., & Shapira, P. (2019). Anticipating governance challenges in synthetic biology: Insights from biosynthetic menthol. *Technological Forecasting and Social Change*, 139, 311–320. <https://doi.org/10.1016/j.techfore.2018.11.020>
- Rip, A. (2010). De Facto Governance of Nanotechnologies. In M. E. A. Goodwin, E. J. Koops, & R. E. Leenes (Eds.), *Dimensions of technology regulation* (pp. 287–310). Wolf Legal Publishers.
- Rohrbeck, R., & Schwarz, J. O. (2013). The value contribution of strategic foresight: Insights from an empirical study of large European companies. *Technological Forecasting and Social Change*, 80(8), 1593–1606. <https://doi.org/10.1016/j.techfore.2013.01.004>
- Rovelli, C. (2021). *Helgoland*. Allen Lane.
- RRI Tools. (2013). <https://rri-tools.eu/>
- Ruddick, N., & Heisenberg, B. (2001). *The Search for a Quantum Ethics: Michael Frayn's "Copenhagen" and Other Recent British Science Plays*. <https://about.jstor.org/terms>
- Sakata, N. (2022). Embracing the Messiness in Mixed Methods Research: The Craft Attitude. *Journal of Mixed Methods Research*, 2022(0), 155868982211085. <https://doi.org/10.1177/15586898221108545>
- Sanscartier, M. D. (2020). The Craft Attitude: Navigating Mess in Mixed Methods Research. *Journal of Mixed Methods Research*, 14(1), 47–62. <https://doi.org/10.1177/1558689818816248>
- Sargeant, Richard, Bambang, Coffey, V., & Kraatz, J. (2010). *Creating value in project management using PRINCE2*. [www.prince-officialsite.com/](http://www.prince-officialsite.com/)
- Savulescu, J., Gyngell, C., & Kahane, G. (2021). Collective Reflective Equilibrium in Practice (CREP) and controversial novel technologies. *Bioethics*, 35(7), 652–663. <https://doi.org/10.1111/BIOE.12869>
- Schön, D. A. (1983). *The reflective practitioner: How professionals think in action*. Basic Books. <https://doi.org/10.4324/9781315237473>
- Schönherr, N., Martinuzzi, A., & Jarmai, K. (2019). Towards a Business Case for Responsible Innovation. In K. Jarmai (Ed.), *Responsible Innovation: Business Opportunities and Strategies for Implementation*. Springer Netherlands. <https://doi.org/10.1007/978-94-024-1720-3>

- Schubotz, D. (2019). *Participatory Research: Why and how to involve people in research*. Sage Publications Ltd.
- Schwarz-Plaschg, C. (2018). The Power of Analogies for Imagining and Governing Emerging Technologies. *NanoEthics*, 12(2), 139–153. <https://doi.org/10.1007/s11569-018-0315-z>
- Sen, B. A. (2010). Reflective writing: a management skill. *Library Management*, 31(1/2), 79–93. <https://doi.org/10.1108/01435121011013421>
- Senge, P. M. (1990). *The Fifth Discipline: The art and practice of the learning organisation*. Doubleday.
- Shamir, E. (2010). The Long and Winding Road: The US Army Managerial Approach to Command and the Adoption of Mission Command (Auftragstaktik). [Http://Dx.Doi.Org/10.1080/01402390.2010.498244](http://Dx.Doi.Org/10.1080/01402390.2010.498244), 33(5), 645–672. <https://doi.org/10.1080/01402390.2010.498244>
- Sharma, M., Choudhary, V., Bhatia, R. S., Malik, S., Raina, A., & Khandelwal, H. (2020). Leveraging the power of quantum computing for breaking RSA encryption. *Cyber-Physical Systems*. <https://doi.org/10.1080/23335777.2020.1811384>
- Smallman, M. (2018). Science to the rescue or contingent progress? Comparing 10 years of public, expert and policy discourses on new and emerging science and technology in the United Kingdom. *Public Understanding of Science*, 27(6), 655–673. <https://doi.org/10.1177/0963662517706452>
- Smallman, M. (2019). ‘Nothing to do with the science’: How an elite sociotechnical imaginary cements policy resistance to public perspectives on science and technology through the machinery of government. *Social Studies of Science*. <https://doi.org/10.1177/0306312719879768>
- Smith, H., Manzini, A., & Ives, J. (2022). Inclusivity in TAS research: an example of EDI as RRI. *Journal of Responsible Technology*, 100048. <https://doi.org/10.1016/J.JRT.2022.100048>
- Špundak, M. (2014). Mixed Agile/Traditional Project Management Methodology – Reality or Illusion? *Procedia - Social and Behavioral Sciences*, 119, 939–948. <https://doi.org/10.1016/J.SBSPRO.2014.03.105>
- Stahl, B. C., Chatfield, K., Ten Holter, C., & Brem, A. (2019). Ethics in corporate research and development: can responsible research and innovation approaches aid sustainability? *Journal of Cleaner Production*, 239, 118044. <https://doi.org/10.1016/j.jclepro.2019.118044>
- Stahl, B. C., Timmermans, J., & Mittelstadt, B. D. (2016). The Ethics of Computing. *ACM Computing Surveys*, 48(4), 1–38. <https://doi.org/10.1145/2871196>
- Stahl, B., Obach, M., Yaghmaei, E., Ikonen, V., Chatfield, K., & Brem, A. (2017). The Responsible Research and Innovation (RRI) Maturity Model: Linking Theory and Practice. *Sustainability*, 9(6), 1036. <https://doi.org/10.3390/su9061036>
- Star, S. L., & Griesemer, J. R. (1989). Institutional Ecology, ‘Translations’ and Boundary Objects: Amateurs and Professionals in Berkeley’s Museum of Vertebrate Zoology, 1907–39. *Social Studies of Science*, 19(3), 387–420. <https://doi.org/10.1177/030631289019003001>
- Steen, M. (2013). Co-design as a process of joint inquiry and imagination. *Design Issues*, 29(2), 16–28. [https://doi.org/10.1162/DESI\\_a\\_00207](https://doi.org/10.1162/DESI_a_00207)
- Stelzer, H. (2020). Responsible Innovation and Climate Engineering. A Step Back to Technology Assessment. *Philosophy of Management*, 19(3), 297–316. <https://doi.org/10.1007/s40926-020-00127-z>

- Stilgoe, J. (2015). *Experiment Earth: Responsible innovation in geoengineering*. Routledge.
- Stilgoe, J., Lock, S. J., & Wilsdon, J. (2014). Why should we promote public engagement with science? *Public Understanding of Science*, 23(1), 4–15. <https://doi.org/10.1177/0963662513518154>
- Stilgoe, J., Owen, R., & Macnaghten, P. (2013). Developing a framework for responsible innovation. *Research Policy*, 42(9), 1568–1580. <https://doi.org/10.1016/j.respol.2013.05.008>
- Stilgoe, J., Watson, M., & Kuo, K. (2013). Public Engagement with Biotechnologies Offers Lessons for the Governance of Geoengineering Research and Beyond. *PLoS Biology*, 11(11). <https://doi.org/10.1371/journal.pbio.1001707>
- Strand, R., Spaapen, J., Bauer, M., Hogan, E., Revuelta, G., & Stagl, S. (2015). *Indicators for promoting and monitoring Responsible Research and Innovation*. <https://doi.org/10.2777/9742>
- Strassheim, H. (2019). Lasswell's legacy revisited: critical perspectives on psychological and behavioural public policy. In *Critical Policy Studies* (Vol. 13, Issue 2, pp. 221–225). <https://doi.org/10.1080/19460171.2019.1625650>
- Suddaby, R., & Laasch, O. (2020). Responsible Management as Re-enchantment and Retrovation. In O. Laasch, R. Suddaby, D. Jamali, & E. Freeman (Eds.), *The research handbook of responsible management*. Edward Elgar. <https://www.researchgate.net/publication/335727449>
- Ten Holter, C. (2020). The repository, the researcher, and the REF: “It’s just compliance, compliance, compliance.” *Journal of Academic Librarianship*, 46(1). <https://doi.org/10.1016/j.acalib.2019.102079>
- Ten Holter, C. (2022). Participatory design: lessons and directions for responsible research and innovation. *Journal of Responsible Innovation*. <https://doi.org/10.1080/23299460.2022.2041801>
- Ten Holter, C., Inglesant, P., & Jirotko, M. (2021a). *Creating a Responsible Quantum Future*. <https://ora.ox.ac.uk/objects/uuid:8e175e19-f879-4827-9cbd-c4849cb6bd60>
- Ten Holter, C., Inglesant, P., & Jirotko, M. (2021b). Reading the road: challenges and opportunities on the path to responsible innovation in quantum computing. *Technology Analysis and Strategic Management*. <https://doi.org/10.1080/09537325.2021.1988070>
- Ten Holter, C., Inglesant, P., Srivastava, R., & Jirotko, M. (2022). Bridging the quantum divides: a chance to repair classic(al) mistakes? *Quantum Science and Technology*. <https://doi.org/10.1088/2058-9565/ac8db6>
- Ten Holter, C., Stahl, B., & Jirotko, M. (2022). Built in, not bolted on: responsible innovation in UK Centres for Doctoral Training – the new normal? *The Learning Organization*. <https://doi.org/10.1108/tlo-09-2021-0117>
- Thelen, K. (2002). The Explanatory Power of historical Institutionalism. In R. Mayntz (Ed.), *Akteure — Mechanismen — Modelle*. Springer Link. <https://doi.org/10.1007/978-1-4020-4676-6>
- Thelen, K. (2004a). *How Institutions Evolve*. Cambridge University Press.
- Thelen, K. (2004b). The Political Economy of Skills in Comparative-Historical Perspective. In *How Institutions Evolve*. <https://doi.org/10.1017/cbo9780511790997.002>
- Thorstensen, E. (2020). *Responsible Assessments Frameworks for a Value-Based Governance of Assistive Technologies*.
- Timmermans, J., Yaghmaei, E., Stahl, B. C., & Brem, A. (2017). Research and innovation processes revisited – networked responsibility in industry. *Sustainability Accounting, Management and Policy Journal*, 8(3), 307–334. <https://doi.org/10.1108/SAMPJ-04-2015-0023>

- Tinoco, R. A., Sato, C. E. Y., & Hasan, R. (2016). Responsible project management: beyond the triple constraints. *Journal of Modern Project Management*, 4(1), 81–93. <http://sro.sussex.ac.uk>
- National Security and Investment Act, (2021). <https://www.legislation.gov.uk/ukpga/2021/25/contents>
- UK NQTP. (2019). *Establishing the National Quantum Computing Centre (NQCC) Town meeting 21 August 2019*.
- Uyarra, E., Ribeiro, B., & Dale-Clough, L. (2019). Exploring the normative turn in regional innovation policy: responsibility and the quest for public value. *European Planning Studies*, 27(12), 2359–2375. <https://doi.org/10.1080/09654313.2019.1609425>
- van de Poel, I., Asveld, L., Flipse, S., Klaassen, P., Kwee, Z., Maia, M., Mantovani, E., Nathan, C., Porcari, A., & Yaghmaei, E. (2020). Learning to do responsible innovation in industry: six lessons. *Journal of Responsible Innovation*, 7(3), 697–707. <https://doi.org/10.1080/23299460.2020.1791506>
- van de Poel, I., Asveld, L., Flipse, S., Klaassen, P., Scholten, V., & Yaghmaei, E. (2017). Company strategies for responsible research and innovation (RRI): A conceptual model. *Sustainability (Switzerland)*, 9(11), 1–18. <https://doi.org/10.3390/su9112045>
- van de Poel, I., & Sand, M. (2018). Varieties of responsibility: two problems of responsible innovation. *Synthese*. <https://doi.org/10.1007/s11229-018-01951-7>
- van der Kolk, B. (2022). Numbers Speak for Themselves, or Do They? On Performance Measurement and Its Implications. *Business & Society*, 000765032110684. <https://doi.org/10.1177/00076503211068433>
- van Oost, E., Kuhlmann, S., Ordóñez-Matamoros, G., & Stegmaier, P. (2016). Futures of science with and for society: towards transformative policy orientations. *Foresight*, 18(3), 276–296. <https://doi.org/10.1108/FS-10-2014-0063>
- van Oudheusden, M. (2011). Questioning “Participation”: A Critical Appraisal of its Conceptualization in a Flemish Participatory Technology Assessment. *Science and Engineering Ethics*, 17(4), 673–690. <https://doi.org/10.1007/s11948-011-9313-z>
- van Oudheusden, M. (2014). Where are the politics in responsible innovation? European governance, technology assessments, and beyond. *Journal of Responsible Innovation*, 1(1), 67–86. <https://doi.org/10.1080/23299460.2014.882097>
- Vermaas, P. E. (2017). The societal impact of the emerging quantum technologies: a renewed urgency to make quantum theory understandable. *Ethics and Information Technology*, 19(4), 241–246. <https://doi.org/10.1007/s10676-017-9429-1>
- Vermeer, M. J. D., & Peet, E. D. (2020). *Securing Communications in the Quantum Computing Age: Managing the Risks To Encryption*. [https://www.rand.org/pubs/research\\_reports/RR3102.html](https://www.rand.org/pubs/research_reports/RR3102.html)
- Volante, L. (2004). *Teaching To the Test: What Every Educator and Policy-maker Should Know*. <http://www.umanitoba.ca/publications/cjeap/articles/volante.html>
- von Schomberg, R. (2013). A Vision of Responsible Research and Innovation. In Richard Owen, J. Bessant, & M. Heintz (Eds.), *Responsible Innovation* (First edit, pp. 51–74). John Wiley & Sons, Ltd. <http://onlinelibrary.wiley.com/doi/10.1002/9781118551424.ch3/summary>
- Watkins, K. E., & Marsick, V. J. (2019). Conceptualizing an Organization That Learns. In *The Oxford Handbook of the Learning Organization* (pp. 50–66). <https://doi.org/10.1093/oxfordhb/9780198832355.013.3>

- WEF. (2022). Quantum Computing Governance Principles. In *World Economic Forum* (Issue January).
- Weirich, H. (1982). The TOWS Matrix. A Tool for Situational Analysis. *Long Range Planning*, 15(2), 54–66.
- Wesselink, A., & Hoppe, R. (2011). If post-normal science is the solution, what is the problem?: The politics of activist environmental science. *Science Technology and Human Values*, 36(3), 389–412. <https://doi.org/10.1177/0162243910385786>
- Wilson, T. D. (2002). The nonsense of “knowledge management.” *Information Research*, 8(1). <https://doi.org/10.1108/13673270510590227>
- Winner, L. (1980). Do Artifacts have Politics? *Daedalus*, 109(1), 121–136.
- Wynne, B. (1993). Public uptake of science: A case for institutional reflexivity. *Public Understanding of Science*, 2(4), 321–337. <https://doi.org/10.1088/0963-6625/2/4/003>
- Wynne, B. (2014). Further disorientation in the hall of mirrors. *Public Understanding of Science*, 23(1), 60–70. <https://doi.org/10.1177/0963662513505397>
- Wynne, B. (2006). Public engagement as a means of restoring public trust in science - Hitting the notes, but missing the music? *Community Genetics*, 9(3), 211–220. <https://doi.org/10.1159/000092659>
- Yan, B., Tan, Z., Wei, S., Jiang, H., Wang, W., Wang, H., Luo, L., Duan, Q., Liu, Y., Shi, W., Fei, Y., Meng, X., Han, Y., Shan, Z., Chen, J., Zhu, X., Zhang, C., Jin, F., Li, H., ... Long, G.-L. (2022). Factoring integers with sublinear resources on a superconducting quantum processor. *ArXiv*, 1–32. <https://arxiv.org/pdf/2212.12372.pdf>

## Appendix A: Papers and other publications

**Ten Holter, C.**, Kunze, L., Pattinson, J.-A., Salvini, P., Attias, J., Jirotko, M. *What's missing from this picture? Ethical, legal, and practical challenges for autonomous-vehicle data-recorders* (In review at *Technology in Society*) [https://papers.ssrn.com/sol3/papers.cfm?abstract\\_id=4310714](https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4310714)

Pattinson, J. A., **Ten Holter, C.**, Grieman, K., Salvini, P., Kunze, L., Jirotko, M. *I can do anything with my AV data (but I won't do that): Public attitudes towards data recorders in self-driving cars* (In review at *Journal of Urban Technology*)

Widdicks, K., Lucivero, F., Samuel, G., Somavilla, L., Tavares Smith, M., **Ten Holter, C.**, Berners-Lee, M., Blair, G. S., Jirotko, M., Knowles, B., Sorrell, S., Börjesson Rivera, M., Cook, C., Coroamă, V. C., Foxon, T., Hardy, J., Hilty, L. M., Hinterholzer, S., Penzenstadler, B. (2023). Systems Thinking and Efficiency under Emissions Constraints: Addressing Rebound Effects in Digital Innovation and Policy. *Patterns* (in press)

Inglesant, P., Webb, H., **Ten Holter, C.**, Patel, M., Jirotko, M., The Responsible Innovation of Disruptive Technologies. (2022) In *The SAGE Handbook of Digital Society* (in press)

Portillo, V., Dowthwaite, L., Creswick, H., Pérez-Vallejos, E., Koene, A., **Ten Holter, C.**, Jirotko, M., Zhao, Z.: *How, Why, What and Who? Young people's views and recommendations for a more transparent online world* (in review at *Information, Communication and Society*)

Ozmen Garibay, O., Winslow, B., Andolina, S., Antona, M., Bodenschatz, A., Coursaris, C., Falco, G., Fiore, S., Garibay, I., Grieman, K., Havens, J. C., Jirotko, M., Kacorri, H., Karwowski, W., Kidera, J., Konstan, J., Koon, S., Lopez-Gonzalez, M., Maifeld-Carucci, I., McGregor, S., Salvendy, G., Shneiderman, B., Stephanidis, C., Strobel, C., **Ten Holter, C.**, Xu, W. (2023) Six Human-Centered Artificial Intelligence Grand Challenges *International Journal of Human-Computer Interaction* (2023) <https://doi.org/10.1080/10447318.2022.2153320>

**Ten Holter, C.**, Stahl, B., & Jirotko, M. (2022). Built in, not bolted on: responsible innovation in UK Centres for Doctoral Training – the new normal? *The Learning Organization*. <https://doi.org/10.1108/tlo-09-2021-0117>

**Ten Holter, C.**, Inglesant, P., Srivastava, R., & Jirotko, M. (2022). Bridging the quantum divides : a chance to repair classic(al) mistakes? *Quantum Science and Technology*, 4–11. <https://doi.org/https://doi.org/10.1088/2058-9565/ac8db6>

**Ten Holter, C.**, Kunze, L., Pattinson, J.-A., Salvini, P., & Jirotko, M. (2022). Responsible innovation; responsible data. A case study in autonomous driving. *Journal of Responsible Technology*, 11(July), 100038. <https://doi.org/10.1016/j.jrt.2022.100038>

**Ten Holter, C.** (2022). Participatory design: lessons and directions for responsible research and innovation. *Journal of Responsible Innovation*. <https://doi.org/10.1080/23299460.2022.2041801>

**Ten Holter, C.**, Inglesant, P., & Jirotko, M. (2021). Reading the road: challenges and opportunities on the path to responsible innovation in quantum computing. *Technology Analysis and Strategic Management*. <https://doi.org/10.1080/09537325.2021.1988070>

Inglesant, P., **Ten Holter, C.**, Jirotko, M., & Williams, R. (2021). Asleep at the wheel? Responsible Innovation in quantum computing. *Technology Analysis and Strategic Management*, 0(0), 1–13. <https://doi.org/10.1080/09537325.2021.1988557>

Widdicks, K., Knowles, B., Blair, G., **Ten Holter, C.**, Jirotko, M., Lucivero, F., ... Webb, H. (2021). Anticipatory Governance in the Technology Sector: Processes, Critiques and Principles for Addressing Grand Challenges in Computing. In *Conference on Human Factors in Computing Systems - Proceedings*. Association for Computing Machinery. <https://doi.org/10.1145/3411763.3441314>

**Ten Holter, C.**, Inglesant, P., & Jirotko, M. (2021). *Creating a Responsible Quantum Future*. (Report)

Grimpe, B., Stahl, B. C., **Ten Holter, C.**, Inglesant, P., Eden, G., Patel, M., & Jirotko, M. (2020). From collaborative to institutional reflexivity: Calibrating responsibility in the funding process. *Science and Public Policy*, 47(5), 720–732. <https://doi.org/10.1093/scipol/scaa038>

de Heaver, M., Jirotko, M., Nulli, M., Stahl, B. C., & **Ten Holter, C.** (2020). Assessment of Responsible Innovation. In E. Yaghmaei & I. van de Poel (Eds.), *Assessment of Responsible Innovation: Methods and Practices*. Routledge. <https://doi.org/10.4324/9780429298998>

**Ten Holter, C.** (2020). The repository, the researcher, and the REF: “It’s just compliance, compliance, compliance.” *Journal of Academic Librarianship*, 46(1). <https://doi.org/10.1016/j.acalib.2019.102079>

**Ten Holter, C.**, Stahl, B. (2020) *Questions of Responsibility: How is RRI being embedded into CDTs?* (Report)

Stahl, B. C., Chatfield, K., **Ten Holter, C.**, & Brem, A. (2019). Ethics in corporate research and development: can responsible research and innovation approaches aid sustainability? *Journal of Cleaner Production*, 239, 118044. <https://doi.org/10.1016/j.jclepro.2019.118044>

# Appendix B: Sample coding framework (policy study)

Nodes				
Name	Files	Referenc		
Addressing challenges		0	0	
Agility		2	3	
Capacity		3	4	
Consultation		2	4	
Funding		1	3	
Governance		6	34	
Impacts		2	13	
Public dialogue		3	10	
Skills		1	1	
Fusion Doctrine		0	0	
Commercialisation		3	9	
International aspects		6	14	
Security and investment		6	12	
Security and prosperity		3	4	
Golden quotes		6	14	
Infrastructure of policymaking		0	0	
Boundaries		5	14	
Challenges		5	8	
Policy and politics, tradeoffs		9	11	
Countering silo-ing		5	12	
Jurisdiction		1	3	
Novel policies		2	2	
Overarching strategy		5	18	
Planning		5	12	
Prioritisation		4	4	
Processes		11	53	
Institutional knowledge and personal aspects		4	10	
Stakeholders		9	44	
Procurement		3	8	
Underlying values		0	0	
Democratisation		3	4	
Trust		1	8	

Figure 11: Coding structure within NVivo for policy case study



## Appendix C: Sample transcript (industry study)

TRANSCRIPT WITH INTERVIEWEE I53

**If we could start with... perhaps you can tell me a little bit more about [your organisation] and the kinds of things you do there, and your perspective and insight on the quantum computing spinouts that you guys have looked at and we can go from there?**

Yeah, definitely. Shall I - on [our organisation], shall I kind of take you from the very top and then get into the quantum bit?

**Yes, all of that sort of background information is really valuable.**

Cool. Great. Yeah. So [our organisation] started in 2015 and we're a venture capital company that invests - almost always - in science spinouts from [the] University that we think can be world-changing, and have a really positive impact.

We've raised about 600 million pounds to date, which we have available to deploy into companies, and we structure ourselves across two or three different sector teams. So on the one hand, we have life sciences team, which tags to everything wet-lab and medical related at [the university]. Then within the other side which we call the technology side of the fund we've got a deep tech team, which covers a lot of hardware based spin-outs and that can be quantum, that can be clean tech, that can be materials and manufacturing. And then we have a software and a health tech team. So they look at Information Engineering, Computer Science, etc.

I kind of split my time, mainly towards software and AI, but also a bit of deep tech and health tech as well just for context. So that's how we structure ourselves.

We like to invest in businesses from the very beginning - we like to be quite active in helping to build them at the start so you'll very, very typically find us starting the journey at the same time the company does with the seed round, and then following our money into later rounds as well. I think one of the nice things about the way that we're structured is we're structured as a company rather than a traditional fund. And what that means is that we have the ability to invest in quite long term technologies and hold our positions in those companies pretty much indefinitely until it makes sense to to exit them. So whereas a conventional fund in London would have to realise gains from positions within a five to seven year time horizon and then pay it all back to their LPs. We have patient capital, we have shareholders rather than LPs. That basically means we can take in money anytime we like and we can hold on to companies for as long as we like and then exit them whenever we like.

And that's particularly relevant to quantum, as an area that requires quite a long, long term commitment and a lot of capital as well. So it helps to be of a decent size, in terms of assets under management, because we can support any company but particularly a quantum one through quite large rounds in a way that a £20 million seed fund couldn't. So I guess that's a bit of background about about us overall.

So within quantum we do a bunch of different things. So hardware we have, I think four or five different companies, split across a few of the different approaches, a few of the different hardware paradigms. The main one that I look after is one called [redacted]. So they do silicon based qubits. And we also have a trapped-ion - at least one trapped-ion company – and a couple of others. And then,

interestingly, I think, for this conversation, in addition to that one I look after another [name] company called [name], and this is a software company and what they do is they are developing quantum secure encryption algorithms and software. So it's actually using classical computing. And the premise is: quantum computers will basically break most classical cryptography. But you can't use quantum computing to protect against that because it'll already be too late. So how can you try and use classical computing to protect against the quantum threat in advance? So that's another of the companies that I look after. Naturally, there's a whole bunch that we see coming through the pipeline, year on year as well, that I won't go through. Does that kind of help as an intro?

**That's all brilliant. I've scribbled down about a dozen extra questions. So can we start with – didn't you used to be [called something else], by the way?**

For a while, and then we kind of rebranded and our name change coincided with {name} our new CEO, coming on board, as well as a whole bunch of other people which is great.

**So in terms of the finance that you raise in order to invest, does that come from a variety of sources?**

Yes, so we have a pretty long list of shareholders, which I think most of it is available either on our website or somewhere else. And it's a bunch of different types of people. So I think we've got [name] in there, some of the banks. I think [name] is on there. So we have a variety of different... and obviously the university's a shareholder. So we have a mix of different people. But I think we've got at least 30 pretty serious institutional ones.

**Particularly when it comes to investing in quantum and AI - I'm sure you're aware of the new NS&I act that prohibits those companies from receiving direct inward investment from people that we might not necessarily think of as our friends in national terms. Do those rules also apply to funding that you would raise to invest in those sorts of startups?**

So I'll caveat this by saying I'm not the expert and the expert would be our head of legal who I'm very happy to connect you with after this call if it's worthwhile. I haven't come across a huge number of cases or any really where we've been restricted from being able to do what we want. I do know that we do data collection and we keep information about the share the level of shareholding that we have in certain companies, which I believe is important for NSI. So, as well as the source of funds as I understand it, it's also important that certain shareholders have to disclose before they go above a certain threshold of ownership. And so that is definitely a live issue for a lot of our companies because - by nature of the fact that we're seed investors and we then follow our money we generally have a - the ones that we've backed hard we have a pretty considerable stake in. So that's a live issue for us, but I can't remember a time where we wanted to do something and weren't able to. There's been a workaround.

**So you mentioned that you have a relatively unique structure. Obviously I'm familiar with standard venture capital, they're always looking for their exit ... they want to know ... we're going to take you through the process and we will see you through the process, but we want to know how we're getting out at the other end. Your structure is different and allows you to take a longer term view. Do you think that's unique to [this institution]? Or would you expect to see that in other [places] - the similar sort of ... presumably there's a [version of your organisation] at [other institutions] as well, their version - would they have a similar sort of structure? Or is your structure unique?**

It's a good question. So it's definitely not unique. Evergreen - our structure is called evergreen. What you would call it an evergreen fund, except that we're technically not a fund we're a company but the evergreen structure is definitely not unique to us. Evergreen funds, as I understand it, are on the rise although still in the minority. They're not ... they're also not unique to science and IP-heavy investments. So the 'patient' or evergreen capital structure makes sense in the context of science-heavy investments, where you might have five years of technical milestones and R&D, and then five years of commercial development and growth. And maybe you get an exit in 10 or 12 years time. But as I understand it, there are some private equity firms and pretty big growth funds that have evergreen structures, even though they might not be investing in particularly technical heavy things. So evergreen structure - not unique, on the rise, but still in the minority of cases.

The ... I think it's rare in the context of university spin out funds at the moment. So [our] equivalent in [the other university you mentioned] is called [name]. Not 100% sure, but I think that they have an evergreen structure - or at the very least they didn't [?] start with an evergreen structure. And I can't speak for the other ones in the UK. Obviously, there are a lot more university funds in the United States and in China, and I can't speak for those either. All I can say really, is that it makes a huge amount of sense to me why you would, why you'd have patient capital if you're, if you're trying to invest in things that can be massively world changing, but realistically, you're going to need 10 or 12 years from the inception of the idea to the exit. And I think it's important to know, in the way that you asked the question - we still need to, we are still a venture capital firm, we need to see that exit route eventually. It is more of a timing ... if it's like a 12 year timeline, that's okay. You know, if it's big and it has the chance to be big, then that's fine for us. And it's not fine for other people.

**As long as there is a timeline, it can be in your case it can be extended.**

Yeah. As long as there's you know, as long as there's an exit route, I think ... I'm pretty sure you would universally get all funds telling you that because at the end of the day is it's what we need. Let's not kid ourselves!

**You've got to go on and invest in other things.**

Exactly.

**So this Evergreen structure, from what you were saying that's not unique to the UK?**

It's not, no. There are examples in the States

**Okay. I was just wondering whether that was something that our government had had the foresight to bring in in terms of you know, allowing for these longer term kind of investments, but...**

Not as far as I'm aware, I mean, in a way it's more similar to a conventional company that would have shares. So it's not necessarily something new full stop, although it is relatively new to the venture capital industry. You know, so the shareholders and the way that money is put in probably resembles a bit more one of the starters we'd invest in than a typical fund but then obviously the way that payouts and dividends work is I'm sure very complex and yeah, there there are things that need to be worked out but it's worthwhile I think.

**From what I understand you get involved at a very early seed stage, and then you go all the way up to the series A, where do you - how far do you go through the funding rounds?**

Yeah, it's a really good question. So we've made this proactively part of our new strategy going forwards and the new strategy is, the answer is there is no endpoint where we absolutely jump off. So we are happy to follow our money and/or companies through all of the rounds. We also have a couple of companies that have IPO-ed and we have a couple of public market positions, which I don't think is a secret. Probably the most famous example is [name], which does the [pharma product], and we have a public market position in that, that we manage very carefully. We also, maybe six to 12 months ago, we hired a head of 'late stage' portfolio and that was a very conscious move to build the team and the expertise internally to be able to manage those positions that are coming down the pipeline. So we are very, very keen for there not to be a point beyond which we can't go.

**And how involved do you get with the companies – because academics who are starting up companies don't necessarily have experience running companies. Do you get involved with the management side? Do you help out on business development issues, compliance – how hands-on?**

Yeah, we make – and I'm particularly passionate about this – but we make a very big deal out of being as helpful as possible to the academic founders, and crucially, that extends even to before incorporation in the seed deal. So what's very typical, in fact, what would be an archetypal process for us is, you know, you have you might meet up with with a couple of academics, you know, six months before a company is even incorporated. And you know, in that six months, the conversation might start with you know, is the spinout journey the right thing for this particular piece of research, or even for this particular academic/prospective founder – because it's quite a personal decision to to found a company and it's a big decision and it's a hard and long road. So we like to think that that from what we've seen, we can be really helpful and kind of empathetic in speaking to academic founders from you know, Day-60. And, yeah, so the conversation can often start, you know, at that super early stage and then eventually it coalesces into 'well, here's a piece of research and here's how a company built around that might look', you know, again, we're very, very happy and we get very excited to be involved in that. So you know, that can look like a bunch of different meetings with the founders, you know, again before any money is in the conversation. And then you get to Okay, the seed deal, and the seed deal, as I'm sure you know, is to fund the company for a certain amount of time, to achieve a certain set of milestones – defining what those milestones are, and what investment the company needs to achieve those is something that that we like to be super proactive in, recognising that, that a lot of the time, academics want that that support to kind of help craft that at least at the start. Equally, if they don't, then that's fine as well. You know, some of our companies, the founding team includes some external founders, and some of those are serial entrepreneurs who've done this lots of times before so you have to kind of flex it depending on the context, but I think the headline is, we like to be super-proactive with shaping the story, particularly at the start where it's really important. And then once the money's in and it's a portfolio company and the seed deal's done, you know, it starts to look a bit more conventional, you know, you'll have a board seat, we always take a board seat, and in that we are as helpful as we can be with – maybe it's investor intros, maybe it's helping to bring new talent in, maybe it's refining pitch decks. You know, small, small things that can make a big difference.

**And do you find that academics make good business people? Would you generalise?**

Yeah, I think in general, the best, or the thing that I've been most impressed by is the ability of academics to identify where they're strong and where they need other people. And even just having that ability and awareness is really it's 90% of the battle. You know, the great advantage of doing what we do, and you know, we're [location]-specific, so, so quite focused geographically. The great

advantage of that is you're just immersed into this pool of researchers where, basically everyone's a world leader in their field. Literally 100% of people that come to us pitching companies are world leaders in their research, which is a pretty privileged position to be in. You know, the flip side of that, which I think everyone including us and academics, would have no problem saying is that, you know, the commercial awareness that you have [here] is probably less than you have in Silicon Valley or London or Berlin. And that's fine when you kind of make yourself aware of that and call it out. And you say, okay, therefore – and the solution is different in different cases – so we have a number of academic founders who are fantastic CEOs. And that's either because they had commercial experience when they were academics, or they've picked it all up and have been fantastic. We also have a number, a large number, of academic founders who are fantastic CSOs or CTOs, and know that they're not the CEO archetype. And again, it's very personal. It's very, very personal. So you have to kind of have that conversation each time. But as I say, I think the main dynamic that I pull out to, to answer your question is that in almost all cases I can think of, you know, the, when the academic is maybe not the right person to lead they almost always say that before we say it, and that makes everything manageable.

**I was going to ask whether you ever have to have a difficult conversation with somebody who is great at what they do, but they suck at management.**

I can't remember a conversation where they also, if that was true, they also didn't acknowledge it and happily asked for help to kind of bring in people around them. I think if you talk to, if you talk to a lot of funds, you'll often hear the phrase kind of bringing people around someone to support them. And that's definitely true. With a lot of what we do.

**And so do you see part of your role as as matching your key academic person with a business development person or somebody who has that other expertise – is that something that you would do, or do you just say to them, 'you need to go out and find somebody who can do this stuff'?**

Yeah, very often, that's something that we help with. So at least two major ways that we help with that, depending on what's right in the situation. Quite often we will help with bringing in talent into the company. So there may, there often comes a certain stage where it's time for a serial, experienced CEO to be hired in to run the business. Very often we will help with or even run that search, that talent search. And so you know, whether that's working with a talent agency to find a longlist and shortlist of candidates, whether that's helping and sitting in with the academic founders on the interviews, and then eventually making the selection, we do all of that for our companies. And the other way that we can help on the talent side, particularly on the commercial, sort of matchmaking, which I think is a great phrase because it's something we need to be doing, we have a what we call an Entrepreneur In Residence scheme [here] and basically what that means is, we go out, and find and hire and bring into the [redacted] ecosystem a handful of either serial entrepreneurs or people with huge experience in a certain industry and we bring them in and we literally, almost let them loose on the departments or particular professors who are interested in teaming up. And the goal of that, you know, over a, let's say, six month or 12 month period, is absolutely to end up with a founding team where you've got an EIR in that executive management position. But actually, the magic of the company is the academic founders and the IP. And there you then have a vehicle with all the parts that you need to get going.

**So you don't just have a pool of people in your back pocket and you know, you can say, Oh, I've got this fabulous marketing person or this amazing operations person.**

Yeah, easier said than done, I guess. But, I mean, that's always the aim – and speaking of new hires, I think about three months ago, we brought in a really, really impressive head of talent into our business, called [name]. And she's, I think, worked at a bunch of funds before and you know, 80% of her job is matching portfolio companies with the best possible people for the roles that they need. And you know, the reason for that – you know, that's an investment from us into that kind of role and that person – and the reason for that is quite simply that having the right people in the company, probably the number one driver for the value of that company eventually.

**And between success and failure.**

Yeah, exactly. Exactly. And you know, whether that's – some of the time that's CEO, some of the time that's other C-suites, but a lot of the time that's Board – having the right people around the Board. Both from an experience point of view, but also from a personal skills point of view. So again, back to bringing the right people around the founders, particularly academic founders. You know, a board that works well with academic founders is a really great way to do that. But it takes a lot of time and effort. Hence, you know, we're investing in that function. And I think portfolio companies are happy that we're doing that.

**Yes, I'm sure. When you're getting involved with helping and supporting these companies, particularly in the early stages, do you have any kind of facility for infrastructure support – office space, that kind of thing? Because often if they're spinning out of the university, they've had access to all these lovely labs and things and suddenly they find themselves, potentially without that necessary hardware?**

Yeah, yeah, definitely. So we have a head of property in our team. So we have that as a formal function. So that's a guy called [name] who I think's been with - he was at [the organisation] before me so I think he's been there for for a while. And so we either own or help to operate or have helped to build a whole bunch of property and office space in [our location]. So I'm trying to think of a good example. There's the [name] building is the first one that comes to mind which is like a massive sort of lab space. The other one that comes to mind is directly underneath our offices, which you absolutely must visit some time if you haven't already. We have an area called [name], which is essentially it's an incubator slash office space hub. And there's, I think there's six or seven little offices in there that are specifically designed for very early stage spin outs and companies. You know, teams of five to 10 people who need a nice space to collaborate with the right IT facilities, you know, also a place to host investors and early customers in a professional way. So, and we kind of rent that out to a bunch of our portfolio companies. So, I think we do try to make efforts to make sure that infrastructure is there.

As soon as you said infrastructure, the other – my mind also goes straight away to cybersecurity. And this is a thing that we're actively working on. And again, myself and a couple of other people in the software team are, as you can imagine, particularly passionate about that kind of area. So we're, I guess refreshing at the moment our view of how we can best support a spinout. Number one because you know, maybe it's not the top of mind thing for an academic founder. But number two, because it's also – regardless of whether you're an academic founder or not, when you found a company, there's just so much to do, and you don't want to be stopping doing the other things just so that you can be cybersecurity compliant. Unfortunately, that is a very, very important thing to do as a company. So you know, the ideal solution is that the investors and maybe also the board kind of help to hold them to account and also to make sure they've got the right systems in place. So I'm doing a piece of work at

the moment of like, is there kind of like a repeatable package of support that we can give companies where you know, with minimal distraction or minimal resource swallowing, on the parts of founders, we can just make sure that they've got all the right things in place.

**I imagine that's particularly acute for academic founders, who come from this essentially very cosy cybersecurity environment where they're protected by the central IT, by departmental IT, and everything they do is safe. And suddenly they're in the big bad outside world where people can steal your IP, they can hack you, they can ransomware you – all of this stuff.**

Yeah, exactly. And, you know, I think that's particularly pertinent to a lot of the technology-heavy stuff that we invest in, you know, quantum is a great example. But that's not the only one. You know, we've got a lot of fairly heavy AI companies with software based intellectual property. And I think it's been said before, and a lot of other people will say this as well, you have to operate on the assumption that someone is trying to infiltrate you right now, every day. And I'm stealing those words from a government run cybersecurity workshop that I was on the other week, and they were like, you have to operate as if someone's trying to hack you right now. So.

**I've spoken to some of the guys at [large cybersecurity organisation] before and they are – those are some scary conversations to have. You know – “you have *no* idea what's going on!”**

Yeah, I mean, those guys are really great. They're doing, so they did, they did a project a while ago where it was similar, similar things to what I just mentioned, but that they're from their side, they're also asking the question, like how do we provide startups with guidance that like both ticks the right boxes, but also is realistic in terms of the amount of time that founders have to spare and they reached out to us to get our point of view and work together and basically come up with what the right sort of list of actions should be. So I really like those guys.

**A really good piece of engagement on their part. I can't remember who was telling me about it, there's a company that's been apparently operating without computers since November because they were subjected to a ransomware attack and they just lost everything. It would be horrifying.**

**How do you decide who to invest in? – and kind of relatedly do people usually approach you or do you go and seek out opportunities?**

On the first one, we don't have a super strict list of kind of investment criteria, but there are a few questions that we like to ask of companies. You know, if I had to, I could give you a really long answer, obviously, if I had to distil it down to what's the one question that I would ask if I could only ask one, the one that I keep coming back to is, “if this works, could it be big?”. And I've kind of deliberately phrased that in the optimistic way. And the reason for that is that, particularly when you're doing early stage venture capital, you will always be able to find a risk or something that could go wrong, or a way in which this could fail.

But because of the way that early stage venture capital works, that's actually ... your problem isn't investing in things that fail. Your problem is not investing in the billion pound company that works. And so if you're – I'm sure you're familiar with the idea of kind of outsized returns, and sort of nonlinear distribution of returns and, but all that really means is that if you look at the early stage VC funds that do well, the kind of rule of thumb that I've been taught, is: of the things that you invest in, probably 70% of them fail, 20% of them maybe break even, and one in 10 returns all the rest and

three times over. And so if you know that that's what's going to happen, obviously, it doesn't help you when you're about to invest in something because you don't know if it's the one

### **You don't know which one of the ten it is!**

Exactly! So in that situation, the strategy is to not worry about investing in things that might fail. You've got to say, 'OK, is there a chance that, OK, there's a bunch of things that go wrong, but if it goes, right, can this be a billion pound company?' or whatever your benchmark is? And if the answer to that is yes, then you have to seriously consider it. So that's kind of what I distil it down into. Now obviously, the longer answer is there's a bunch of stuff we do around - the usual things, the market the IP, particularly important for us because we do spinouts is sustainable competitive advantage from that IP. So you know, we always ask questions about replicability, prior art, you know, both patentability and also other forms of protection. You know, trade secrets, know how, proprietary data, but with everything that we do, we look for some kind of IP edge or technical edge and that probably makes us different to a lot of other funds that do you know, whether it's just pure B2B SaaS<sup>54</sup> or even B2C marketplaces where there isn't, there isn't really a moat [unclear?]. That probably makes us a little bit different. Yeah, I could go on for ages, but that's probably the main stuff

### **So when you're looking at whether to invest in something, it sounds like you're more 'pro' than cautious. Is it the person that sells it or the technology that sells it?**

Yeah, that's a really good question. There have been times where both of those have been, or one of those have been true, and then the other has been true. So I guess good examples there are times where, you know, I think always the technology piece needs to be there. And then the 'how do you put the best team on top of that' is where you have multiple different answers. So sometimes a founding team, whether it's just academics or whether it's an academic and an external person, or an entrepreneur in residence, sometimes an opportunity will just come to you with a team that's just knockout and it's there and it's fully formed, and we're good to go. A lot of the time, an opportunity will come and the technology is fantastic – and when I say fantastic I don't just mean that it works and is interesting but also solves a problem that is a real commercial problem, etc. So sometimes you get things coming where the technology is great – And then you think, OK, we have pieces of the team here and now we need to embellish that and make sure that we have everything. And by the way, you know, sometimes that doesn't need to happen right at the start. So another archetype that we often see, which can be really exciting, is, let's say you've got a couple of academic founders. And let's say that it is the case and they are aware that it's the case that you know, maybe they're not the CEO and someone else needs to come in – it doesn't mean you to find that person on day one. Sometimes there is a seed or a pre seed stage that is super valuable, you know, maybe it allows you to validate the technology in a more real world setting. Maybe it's about building out certain other functionality before you properly make a commercial push. And you know, sometimes you can have that stage where it is just an academic team, and we fund the seed deal and it's sort of a 12 month, we sometimes call it exploratory phase. And it can be really valuable to just have a super lean academic team there. So – case-dependent and, but you know if there's a story there around the technology we have to find a way to make the team and the plan work and that's our responsibility as much as the company's

---

<sup>54</sup> Business-to-business software-as-a-service

**Thinking a bit more specifically about some of the quantum companies that you're funding. Do you – how important is it that the government is trying to create this quantum sector? I mean, does that kind of sort of high level strategy have any influence on the kind of decisions that you make in terms of, let's look at what we're funding or let's go and look for something quantum-y for our portfolio?**

I think it does. Yeah. So I think the one that immediately springs to mind is, we see things in our pipeline that have been germinated as a result of either government funding, as part of the quantum strategy, or the government's views on which areas are particularly attractive. So I think it directly influences the pipeline of opportunities that we see and we then choose to invest in some of them. So that's that's definitely the case. I'm sure that it will have an effect on how we treat our portfolio companies. I think there, you know, a lot of that at the same time has to be our own thesis areas about what we think is exciting.

You know, one of the questions that is probably interesting for your research is, should investors essentially spread-bet across different types of quantum hardware? Or should they try and pick one that they think is going to be the winner and double down on that? And I think that's a really interesting question, because people are coming up with different answers at the moment. So I think it – probably at the very beginning of the quantum wave, it used to be the case that almost everyone said 'look, I'm going to try and spread my bets. I'm going to try and go across the four or five different types of hardware'. And you know, when it's seed deals, you can afford to do that. And I think what we're seeing now is as round sizes get bigger<sup>55</sup>, and the companies get more capital intensive, you're starting to see people are forced or incentivized into picking certain approaches and having to form their own views about which parts of either the stack or which parts of the hardware they're most excited by, and then focusing their money there. As I say, we're – at the moment, I think we're a fairly privileged exception to that because we're a £600m fund. If you're a £30m, or £50m, or even £100m pound fund, you can't do that with quantum companies, because by the time they get to Series B, they're raising like £40-£50 million pounds. If you're trying to defend your stake, you're gonna have to – you're having to do maybe £10, £15, £20 million rounds. And you can't do that over five companies if you're a £100m fund. So I think that is a really interesting dynamic, and I'm looking forward to how that plays out in the next five years.

**Do you – does government ever talk to you, or other VCs about the way that the sector is funded, is being encouraged, is being governanced is ... you know, do policymakers ever come to you and say, What do you think, what are you seeing?**

So, I'm trying to think about the different kind of arms that we engage in. So you know, we've got a, we happen to have a really close relationship with Innovate UK, which I guess is kind of a pretty key touchpoint. And we often have conversations about what's exciting, what they're backing and what we're backing – not just with quantum, by the way, but with agtech and other areas. Also one of the guys that does quantum at Innovate UK used to work at [our sister organisation]. We obviously know him. And so that's nice.

As I say, you know, it's not a quantum example but the NCC, you guys, and the people that do ... the CPNI<sup>56</sup> – the people who do guidance around cyber, you know, they often come to us and are like

---

<sup>55</sup> Ie funding rounds

<sup>56</sup> Centre for the Protection of National Infrastructure <https://www.cpni.gov.uk/>

'what do you think?' so, different arms of government, and we obviously welcome that. I think on quantum we probably should do more. I think we should do more engagement with the relevant parts of the government on what they're thinking.

**Is it usually them coming to you or do you ever go to DCMS and say, 'we think you need to do something about this?'**

Honestly, it's probably more people coming to us but there's no reason why it should be. Other than naturally, everyone's busy, which is always the answer, but it's ...

**you wait till they knock on your door.**

Yeah, exactly. I mean, it's kind of it's a cop out, but it's also not a cop out because we are very busy but yeah, we should. I think it will always be – I'll always be in a position of saying this, but I might as well say it anyway. You know, we should do more. We should, we should talk with them more.

**I'm very aware of the time so I'll make this my last question. In terms of the development of the quantum sector, what do you see as the the key challenges and opportunities of the next few years or the enablers and constraints if you like,**

I think the most exciting thing for me, which is always how I conceptualise quantum investing, is to kind of look at the stack from the hardware all the way through to the different software layers. So without being a technical expert, I'll put my hand up straightaway, you could move from the bottom upwards from the hardware manufacturers, the people doing the chips through to the control layer, which is your error correction, calibration of qubits, making sure that things are working properly. One level up from that is operating systems and there are companies that are specifically going after that. And then you've got a layer up from that, which is your developer tools, all of which have analogies in classical computing, which is an interesting thing to look at when you're trying to work out how the space will play out. And then on top of that you've got your actual apps and specific use cases for quantum.

So I kind of – I like looking at that - I like looking at companies who are trying to target each part that there. I like thinking about which areas of that might be consolidated. So you know the hardware guys are going to want to move up the stack and push up. The software guys are going to want to move as far down as possible, so there's a land grab there. So I really like thinking about that.

Naturally, there's the question of end users and end use-cases of quantum and we're doing a piece of work on that at the moment in terms of the thesis areas, you know, the, I think the kind of prevailing view, and we're sort of bouncing around on the hype curve with quantum at the moment. But I think maybe two or three years ago, the answer might have been 'oh, you're going to be able to use quantum for everything'. I think now the prevailing view is see-sawing to actually the range of use cases where you get exponential improvement over classical computing – it's not that big, but they are hugely important. The ones that get mentioned most often – drug discovery, materials, design and simulation, finance. Massive markets. But I think people are coming round to actually the proportion of all computing use cases where we're going to be switching over to quantum soon is not big. And that's fine. Because it's much better to be aware of that, find out where you get that exponential incremental advantage - not try and see it as a silver bullet. So I like thinking about the end use cases.

And of course the other thing that I'm quite excited about in terms of our position is opportunities and threats – as with any new technology. So one of the reasons, I love one of the things that I love is that we have a bunch of quantum hardware companies and ... one of the major threats of quantum computing is breaking classical cryptography. We also have [name of company], which is one of our companies that tries to do post quantum cryptography.

**So you feel secure!**

Exactly. So one of the things that I quite like is as investors if you're investing in a relatively new area of technology, which since the beginning of time has carried - always carried opportunities and risks, why not invest in the opportunities *and* the counteracting of the risks? You ... almost the only analogy I can think of really is kind of carbon neutral – you're almost risk-neutralising the situation and we do that with AI as well. So I look after quite a few of our AI companies. Another great corollary where: massive opportunity, also ethical risks and responsible investing risks, which is why I'm really, I'm really gassed about the fact that we have a bunch of companies that do AI transparency, and explainability, making models accountable. So you're kind of covering both bases. And I think that's really great to be able to do as an investor.

**Fantastic. Thank you so much for your time this afternoon.**

-ends-