

UNIVERSITY OF SURREY

**Expertise as an Object: An
Ontological Study of Cryptology
Research in the United Kingdom
from 1970 to 2000**

by

Richard Fletcher

Submitted for the
degree of Doctor of Philosophy

Faculty of Arts and Human Sciences
Department of Sociology

Supervisors: Dr Christine Hine and Dr Nicola Green

©Richard Fletcher

Words: 95,698

Declaration of Authorship

This thesis and the work to which it refers are the results of my own efforts. Any ideas, data, images or text resulting from the work of others (whether published or unpublished) are fully identified as such within the work and attributed to their originator in the text, bibliography or in footnotes. This thesis has not been submitted in whole or in part for any other academic degree or professional qualification. I agree that the University has the right to submit my work to the plagiarism detection service TurnitinUK for originality checks. Whether or not drafts have been so-assessed, the University reserves the right to require an electronic version of the final document (as submitted) for assessment as above.

Signature:

Date:

“Quis custodiet ipsos custodes?”

Juvenal (Satire VI, lines 347-8)

Abstract

Recent research carried out under the ‘third wave of science studies’ has produced robust categories of expertise, and has developed normative ideas about how it should be used during controversies over technological decision-making. Though separate in the literature, third wave ideas about contributory expertise appear to be compatible with the recent ‘turn to ontology’. Both sets of ideas focus on what it is that practices are able to produce, and consider the results of those practices to be real. It is argued here that contributory expertise can usefully be treated as an ‘object’ under the ontological framework, thus placing additional analytical focus on the practices that are used to enact it.

To explore this idea, documentary analysis and qualitative interviews have been used to produce a description of cryptology research and the crypto wars in the United Kingdom from 1970 to 2000. The cryptology research carried out at four research sites will be described. It is argued that, given divergence amongst the institutional research practices used at each site, the contributory cryptology expertises enacted during this period were ‘multiple’, and can be identified as such using sociological discrimination. A description of how these expertises were then used during the crypto wars - a subsequent controversy over the regulation and use of cryptography in the 1990s - is also provided. It is argued that, as a consequence of this multiplicity, expertises were used during the crypto wars in different ways and for different purposes. In particular, the consequences of basing political decisions on expertise enacted in secret are described. It is argued that acknowledging multiplicity amongst contributory expertise could be used to improve the application of ‘elective modernism’, and to refine its core tenets through the application of a ‘Minimum Transfer Requirement’ and the identification of the ‘problem of expert discrimination’.

Acknowledgements

I would like to extend my gratitude to my supervisors, Christine Hine and Nicola Green, for their excellent supervision over the course of the last four years. I'd like to thank them for their patience in letting me figure where my ideas were leading, for allowing me to draw on their extensive knowledge of the Science and Technology Studies literature, and for their invaluable help in preparing for preparing research papers and conference presentations.

I would also like to extend my gratitude to all of my interviewees for being generous and helpful with their time, and for sharing their memories with me. Similarly, I would like to thank all of the archivists for their help in tracking down the relevant sources.

I would like to thank all of the other members of staff in the Department of Sociology at the University of Surrey for providing a unique and inspiring environment for doctoral research. In particular, I would like to thank Kate Burningham, Rachel Cohen and Rob Meadows for allowing me to teach on their undergraduate Research Methods modules. Special thanks also go to Geoff Cooper for allowing me to have the first pick of his vast collection of academic books upon his retirement. I would also like to thank all of the administrative staff for making everything run so smoothly.

I am also grateful for the stimulating and enlivening presence of my fellow doctoral students in the Department of Sociology. In particular, I would like to thank Peter Johnson for his enthusiasm for discussing social research, his willingness to collaborate with me on a separate project, and most importantly, for driving us both from London to Guildford on a regular basis.

I would also like sincerely thank friends and housemates past and present for their support and motivation over the course of the last four years. Although it probably didn't seem like you were helping me complete my research - you were. Special thanks, of course, go to Clemy Walker.

Finally, I would like to thank the Economic and Social Research council for generously providing the three years' worth of funding that made the research possible.

Contents

Declaration of Authorship	i
Abstract	iii
Acknowledgements	iv
List of Figures	x
List of Tables	xi
Abbreviations	xii
1 Introduction	1
1.1 Controversies and Expertise	1
1.2 Cryptology and the Crypto Wars	6
1.2.1 Why Cryptology and the Crypto Wars?	7
1.2.2 The History of Cryptology	10
1.2.3 Public-Key Cryptography	11
1.2.4 The Literature on the Crypto Wars	13
1.3 Research Questions and Answers	20
1.4 Chapter Guide	23
2 Literature Review	28
2.1 Introduction	28
2.2 Expertise and Controversies over Technological Decision-Making	28
2.2.1 What is a Controversy?	29
2.2.2 Controversies and the Sociology of Scientific Knowledge	31
2.2.3 Controversies and the Sociology of Technology	32
2.2.4 Ways of Understanding Controversies	34
2.2.5 Controversies over Technological Decision-Making	39
2.2.6 The Third Wave of Science Studies	42
2.2.7 The Periodic Table of Expertise	46

2.2.8	Specialist Expertise	49
2.2.9	Elective Modernism	51
2.2.10	Meta-Expertise	57
2.2.11	Criticisms of Elective Modernism	58
2.2.12	Summary	60
2.3	The Ontological Framework	61
2.3.1	The Epistemological Framework	62
2.3.2	Ontological Fluidity	63
2.3.3	Ontological Multiplicity	65
2.3.4	Developing the Ontological Framework	66
2.3.5	Criticisms of the Ontological Framework	70
2.3.6	Summary	73
2.4	Conclusion	73
3	Methodology	77
3.1	Introduction	77
3.2	Methodological Background	77
3.2.1	Case Study	78
3.2.2	Historical Methods	82
3.3	Data Gathering	85
3.3.1	Research Design	85
3.3.2	Documentary Analysis	87
3.3.3	Practicalities of Documentary Analysis	90
3.3.4	Semi-Structured Interviewing	92
3.3.5	Practicalities of Semi-Structured Interviewing	96
3.3.6	A Note on Ethics	100
3.4	Conclusion	100
4	Cryptology Research at the National Physical Laboratory	103
4.1	Introduction	103
4.2	Historical Overview of NPL	104
4.2.1	Early Computing Research	106
4.2.2	New Public Management	108
4.3	The Data Security Group	112
4.3.1	Cryptology Standards	113
4.3.2	Later Governance of NPL	116
4.3.3	Testing and Accreditation	119
4.4	Conclusion	122
5	Cryptology Research at Royal Holloway	123
5.1	Introduction	123
5.2	Historical Overview of the University of London	124
5.3	Historical Overview of Racal	126
5.3.1	Early Years	127
5.3.2	Development of Communications Security Products	128

5.3.3	Enlisting Cryptology Expertise	129
5.4	Historical Overview of Royal Holloway	131
5.5	The Information Security Group	133
5.5.1	Formation	133
5.5.2	Early Cryptology Research and Industrial Collaboration	135
5.5.3	Teaching	137
5.6	Conclusion	138
6	Cryptology Research at the University of Cambridge	140
6.1	Introduction	140
6.2	Historical Overview of the Cambridge Computer Laboratory	141
6.2.1	The Roots of Computing Research	141
6.2.2	Formation	142
6.2.3	Early Computer Development	144
6.3	Computer Security and Cryptology Research	145
6.3.1	Early Cryptology and Computer Security Research	147
6.3.2	The Security Group	148
6.3.3	Cryptology and Security Systems	150
6.4	Conclusion	152
7	Cryptology Research at the Government Communications Head-	
	quarters	154
7.1	Introduction	154
7.2	The UK's Intelligence Organizations	155
7.3	Historical Overview of GCHQ	157
7.4	The Communications-Electronics Security Group	160
7.4.1	Formation	160
7.4.2	Research on Non-Secret Encryption	162
7.4.3	Increased Public Awareness	166
7.4.4	Public Role	168
7.4.5	Continued Secrecy	169
7.5	Conclusion	170
8	The Political Phase of the Crypto Wars	173
8.1	Introduction	173
8.1.1	A Note on Organization	174
8.2	The Debate over Trusted Third Parties	175
8.2.1	Background	175
8.2.2	The Announcement of the TTP Proposals	178
8.2.3	The First Stage of Consultation	188
8.2.4	The Second Stage of Consultation	197
8.2.5	The End of the TTP Debate	203
8.3	The Debate over Export Controls	205
8.3.1	The Background	206
8.3.2	The Announcement of the Export Control Proposals	208

8.3.3	The Consultation Process	210
8.3.4	The Export Control Bill	215
8.4	Conclusion	217
9	Multiplicity and its Consequences	220
9.1	Introduction	220
9.2	Contributory Expertise During the Technical Phase	221
9.2.1	Identifying Multiplicity Using Sociological Discrimination	223
9.3	Contributory Expertise During the Political Phase	234
9.4	The Transfer of Contributory Expertise	247
9.5	Consequences for Elective Modernism	251
9.5.1	Minimum Transfer Requirement	252
9.5.2	Problem of Expert Discrimination	253
9.6	Conclusion	254
10	Conclusion	257
10.1	Introduction	257
10.2	Research Questions and Answers	257
10.3	Reflections, Limitations, and Further Work	264
A	Information Sheet and Consent Form	270
B	Archival Sources	273
B.1	British Library	273
B.1.1	Collection: General Reference	273
B.1.2	Collection: Science, Business, and Technology	273
B.1.3	Collection: Trade Literature	274
B.2	Cambridge University Archives	275
B.2.1	Collection: Archives of the Mathematical Laboratory and its successor, the Computer Laboratory	275
B.3	Imperial College London Archives	276
B.3.1	Collection: Donald Davies' Papers	276
B.4	National Archives	276
B.4.1	Collection: Records of the Department of Scientific and In- dustrial Research, the National Physical Laboratory	276
B.4.2	Collection: Records of the Home Office, Ministry of Home Security, and Related Bodies	278
B.5	Parliamentary Archives	278
B.5.1	Collection: Records of the House of Commons	278
B.6	Royal Holloway Archives	278
B.6.1	Collection: Royal Holloway College Papers	278
B.6.2	Collection: Royal Holloway and Bedford New College Papers	280
B.7	Wayback Machine	281

Bibliography

282

List of Figures

1.1	An Example of a Simple Substitution Cipher	10
1.2	An Example of a Vigenère Cipher	11
2.1	The Periodic Table of Expertise (Collins & Evans 2007)	48
2.2	The Relationship Between the Technical and Political Phases (Evans & Plows 2007)	55
3.1	Methodology Flowchart	78
7.1	The Organization of GCHQ in 1946 (Aldrich 2010)	159
7.2	The Organization of GCHQ in 1970 (Aldrich 2010)	161
8.1	The Passage of a Bill in the UK Parliament (UK Parliament 2013) .	174

List of Tables

1.1	Research Design	25
3.1	Case Study Bounding	81
3.2	Updated Research Design	85
3.3	Types of Document Analysed	91
3.4	Interviewee Demographics	97
8.1	Summary of Responses to the First TTP Consultation (Hosein 1997)	194
9.1	Eight Types of Scientific Laboratory (van Rooij 2011)	225
9.2	Research Practices of the Data Security Group at NPL	230
9.3	Research Practices of the Information Security Group at Royal Hol- loway	231
9.4	Research Practices of the Security Group at the University of Cam- bridge	232
9.5	Research Practices of CESG at GCHQ	233
9.6	Summary of Multiple Research Practices	234
9.7	The Transfer of Contributory Expertise	251

Abbreviations

ACE	Automatic Computing Engine
ANT	Actor-Network Theory
ARPANET	Advanced Research and Projects Agency Network
ATM	Automated Teller Machine
BAN	Burrows, Abadi, Needham
BMA	British Medical Association
CASM	CESG Architecture for Secure Messaging
CBE	Commander of the Most Excellent Order of the British Empire
CCSC	Commercial Computer Security Centre
CESD	Communications-Electronics Security Department
CESG	Communications-Electronics Security Group
CLEFs	Commercial Licensed Evaluation Facilities
CoCom	Coordinating Committee for Multilateral Export Controls
COMSEC	Communications Security
DES	Data Encryption Standard
DSD	Domain-Specific Discrimination
DSG	Data Security Group
DSIR	Department of Scientific and Industrial Research
DTI	Department of Trade and Industry
EC	European Commission
EDSAC	Electronic Delay Storage Automatic Calculator
EFF	Electronic Frontier Foundation
EPOR	Empirical Programme of Relativism
EPSRC	Engineering and Physical Sciences Research Council

EU	European Union
FIPR	Foundation for Information Policy Research
FoI	Freedom of Information
GC&CS	Government Code and Cipher School
GCHQ	Government Communications Headquarters
GoCo	Government-owned Contractor-operated
GPT	GEC Plessey Telecommunications
GRE	Government Research Establishment
HMG	Her Majesty's Government
IBM	International Business Machines
IMG	Information Management Group
ISG	Information Security Group
IT	Information Technology
ITD	Information Technology Division
ITSEC	Information Technology Security Evaluation Centre
ITSEM	Information Technology Security Evaluation Manual
LCSA	London Communications Security Agency
MAA	Message Authenticator Algorithm
MDP	Minimal Default Position
MinTech	Ministry of Technology
MoD	Ministry of Defence
MP	Member of Parliament
MSc	Master of Science
NAMAS	National Measurement Accreditation Service
NASA	National Aeronautics and Space Administration
NBS	National Bureau of Standards
NGO	Non-Governmental Organization
NHS	National Health Service
NPL	National Physical Laboratory
NSA	National Security Agency
OECD	Organization for Economic Cooperation and Development

PGP	Pretty Good Privacy
PhD	Doctor of Philosophy
PIN	Personal Identification Number
PIU	Performance and Innovation Unit
R&D	Research and Development
RIPA	Regulation of Investigatory Powers Act
RAE	Research Assessment Exercise
RSA	Rivest, Shamir, Adleman
SCOT	Social Construction of Technology
SEE	Studies of Expertise and Experience
SERC	Science and Engineering Research Council
SIGINT	Signals Intelligence
SOG-IS	Senior Officials Group - Information Systems Security
SSK	Sociology of Scientific Knowledge
STS	Science and Technology Studies
TTCC	Tokens and Transactions Control Consortium
TTPs	Trusted Third Parties
UK	United Kingdom
UKUSA	United Kingdom-USA
UNIVERSE	UNIV-Expanded Ring and Satellite Experiment
US	United States

Erratum

The submitted version of Table 3.4 contained the wrong number of interviewees, and has been updated. The author apologises for the error.

Chapter 1

Introduction

1.1 Controversies and Expertise

This thesis is about controversies over technological decision-making. A controversy over technological decision-making is a protracted disagreement over what the right political decision might be on an issue that intersects with science and technology. It therefore refers to controversies over, say, the use of nuclear power, the production of genetically modified food, the right response to climate change, or the use of vivisection to further scientific knowledge. Controversies over technological decision-making are important because, given the prominent role of science and technology in Western societies, their outcomes partly determine the nature of the society that we live in. The outcomes of controversies over technological decision-making have the potential to improve or worsen the lives of many within society. In extreme cases, arriving at a particular decision can determine whether human lives are saved or lost (e.g. Weinel 2010).

Studies of controversies over technological decision-making have been a core component of the interdisciplinary field of Science and Technology Studies (STS) since the 1970s (e.g. Nelkin 1971, 1979, 1984, 1993*b*, Wynne 1982, 1992, Epstein 1996, Mulkay 1997, Jasanoff 2005). It is perhaps surprising, then, that although work within STS has been successful in delivering rich, detailed, case study descriptions, some scholars have expressed a reluctance to use these as a platform for developing normative ideas about how actors ought to behave, and what processes should be used, if the ‘best’ outcome is to be achieved (e.g. Jasanoff 2003, Lynch & Cole 2005). Those that have identified a need for a more ‘engaged’ approach have

often called for various forms of increased democratic public participation in controversies over technological decision-making (e.g. Irwin 1995, Burningham 1998, Wilsdon & Willis 2004). However, in the last ten to fifteen years, a body of work under the heading of the ‘third wave of science studies’ (hereafter ‘third wave’) has emerged that offers an alternative approach (Collins & Evans 2002, 2007, Collins et al. 2010, Collins 2014a). As Harry Collins and Robert Evans have explained:

What we want to do is consider how to make good decisions in the right way. But our particular concern is to find a rationale which is not inconsistent with the last three decades of work in science studies. Our initial claim is that though many others working within the science studies tradition have studied the problem, and contributed valuably to the debate about technical decision-making¹, they have not solved it in a way that is completely intellectually satisfying (Collins & Evans 2002, p.236).

The third wave is rooted in the study of expertise. Put simply, expertise is considered important because it is a powerful tool that can be used to inform decisions, justify preferences, and frame questions. According to Collins and Evans (2002), the study of the relationship between controversies over technological decision-making and expertise can be divided into three ‘waves’. They have argued that sociological analysts working within a particular wave have tended to hold a certain view of the relationship between scientific and technological expertise and controversies over technological decision-making. They argued that, during the first wave of science studies, which preceded the 1960s, good scientific credentials were largely seen as synonymous with expertise. Scientific credentials were therefore used as criteria for demarcating the expert from the non-expert, occasionally even on non-scientific or non-technical matters. Though nowadays considered outdated by STS scholars, echoes of this view can be found in contemporary popular

¹Earlier third wave work refers to ‘technical decision-making’ rather than ‘technological decision-making’. However, ‘technological decision-making’ is now used as standard after clarification from Weinel (2010).

discourse.² One consequence of this view is that, during controversies over technological decision-making, given that only credentialed scientists can be thought of as experts, they are the only ones imbued with authority to influence technological decisions. Taken to its extreme, this constitutes a form of technocracy (Collins et al. 2010).

As STS matured, and in line with a broader rejection of technocratic ideals (see Ezrahi 1990), the belief that authority should reside solely in credentialed scientific and technological expertise during controversies over technological decision-making has been shown to be both untenable and undesirable. Sociological analysts working under the second wave of science studies, which started in the 1960s and continues to the present day, have shown that credentialed scientific expertise, when applied by itself to many controversial issues in society, is not necessarily robust or flexible enough to satisfactorily resolve them. In some cases, it is argued, scientific expertise needs to be complemented with other types of knowledge about the problem being tackled (e.g. Irwin 1995, Epstein 1996). However, some working under wave two have gone further in their critique of the wave one understanding. Here, the very use of scientific expertise in controversies over technological decision-making is challenged because it is seen to impose an undemocratic scientific and technological framing that is in conflict with wider moral or political priorities (Wynne 2003). Taken to its extreme, this rejection of technocracy can result in an endorsement of ‘technological populism’, given that it threatens to erode any distinction between experts and non-experts (Collins et al. 2010).

In response, those working under the third wave consider both the technocracy characteristic of wave one, and the technological populism characteristic of wave two, to be unsatisfactory. Therefore, they have attempted to chart a path between the two. They have done this through a reconceptualization of expertise (Collins & Evans 2007). In short, the third wave sees expertise as a product of sustained social experience of a particular domain, rather than resulting solely from possessing formal credentials, studying primary sources, or being an active citizen. Thus,

²This is visible, for example, in debates over whether there should be more UK Members of Parliament (MPs) with scientific expertise. Those that have argued in favour of this have done so not because of the benefits to be gained from the expertise that scientists might have in their own specific field, but rather because of the belief that their training has provided them with the expertise to generally make more rational, evidence-based decisions (e.g. Henderson 2011). However, others have argued that there is no strong evidence to suggest that MPs with scientific expertise vote any differently from those that don’t (Goodwin 2014).

expertise is primarily rooted in the tacit knowledge that is acquired through social experience, rather than the explicit knowledge that can be accrued without it (Collins 2010). Consequently, those who can be said to possess expertise relevant to technological decision-making may include those who don't possess formal scientific credentials but do have practical experience of contributing towards a relevant scientific or non-scientific domain, whilst also allowing for the possibility of excluding those that do possess good scientific credentials if they lack the relevant experience. The third wave, through its periodic table of expertise, states that those who have contributory expertise - acquired through the accumulation of the tacit knowledge that comes from experience of contributing to a field, or interactional expertise - acquired through the accumulation of tacit knowledge based on experience of being immersed in the language of that field, can be said to possess specialist expertise. This is contrasted with forms of non-specialist expertise, which is typically based on knowledge of a domain that can be made explicit, and thus can be acquired without being engaged in certain social practices (Collins & Evans 2007).

This way of conceptualizing expertise has been used as a platform for ideas about how it should be used during controversies over technological decision-making. This aspect of the third wave has been called 'elective modernism' (Collins et al. 2010). Elective modernism, which remains a work in progress, attempts to define a system that avoids both technocracy and technological populism. It does this by embodying the insights from the second and third waves in normative rules about how controversies over technological decision-making should be managed. Elective modernist prescriptions are designed to allow for controversies over technological decision-making to be informed by the relevant specialist expertise, whilst preventing that expertise from dominating democratic processes. There is thus "a preference for democracies which actively promote discussion and debate of technical matters yet which reject populism of all kinds while still rejecting technocracy", and a belief that "scientific values [are] among those which should be at the heart of a good society" (Collins et al. 2010, p.185).

Third wave ideas have not escaped criticism (Wynne 2003, Rip 2003, Jasanoff 2003, Fischer 2011, Forsyth 2011, Epstein 2011). However, critical engagement has often been disappointing because most critiques have not addressed third wave concepts directly. Rather, they have focussed on the premise on which the third wave is based. It is typically argued that the descriptive work of the second wave

should continue, and the normative work of the third wave should be postponed or even abandoned. Those that defend the third wave have (reluctantly) pointed out that these criticisms resemble those typically found in cases of Kuhnian paradigm incommensurability, in that:

- (1) The two sides tend to use disparate sets of exemplary cases to illustrate their points.
- (2) There is a failure to get sufficiently far ‘inside’ the new perspective to know how to give it a fair or even charitable reading before starting on the critique.
- (3) There is a tendency to restate the core beliefs of the ‘old paradigm’ as if they were criticisms in themselves.
- (4) The old paradigm is said to be able to cope with all the things that the new view claims to solve so long as one accepts a few small anomalies and inconsistencies (Collins et al. 2011, pp.340-341).

The style and nature of the debate over the third wave therefore belies the fact that the third and second waves occupy much common ground, and that the third wave attempts to use insights from the second wave to inform its normative arguments. In making this point, one observer has drawn attention to the small pool of case studies used as a basis for understanding expertise during controversies over technological decision-making, and has argued that more are needed:

I am struck on reading the exchanges about the Third Wave by the abstraction of the debate, yet at the same time by the powerful influence of a relatively small number of studies of actual decision-making processes. . . . That empirical investigations have done so much to illuminate relations between knowledge and policy/political processes suggests that we should far more often ‘go and see’, especially in the highly dynamic circumstances of technological innovation and change (Owens 2011, p.331).

It is with this in mind that I introduce a new case study of a controversy over technological decision-making to the third wave debate. In this thesis, I will present a case study of cryptology research and the crypto wars in the United Kingdom from 1970 to 2000. The purpose of this case study is to both contribute towards the third wave debate, and to improve our sociological understanding of cryptology and the crypto wars.

1.2 Cryptology and the Crypto Wars

Cryptology and the crypto wars are not topics that most people are familiar with, so it is useful to introduce the basics. Cryptology is concerned with the writing and breaking of codes and ciphers. ‘Cryptography’ refers to “the science and art of designing ciphers”, ‘cryptanalysis’ to “the science and art of breaking them”, and ‘cryptology’ to “the study of both” (Anderson 2008, p.130).

Cryptology can therefore be used, in its most straightforward implementation, to communicate in secret through the sending and receiving of encrypted messages. Cryptography can be used to encrypt the content of emails such that only the intended recipient is able to read them. This makes cryptographic technologies important to those concerned that individual privacy is being eroded by the increasing prevalence of electronic communication, and the ease with which such communications can be intercepted. Additionally, developments in modern cryptology are significant because, as well as providing the means to encrypt (and decrypt) simple messages, they can also be embedded into electronic technologies to provide: data confidentiality - restrictions on who or what can access data; data integrity - assurances about the accuracy and consistency of data; and authentication - a means of confirming whether someone or something is who or what they claim to be.³ Modern cryptology is therefore used to underpin the security of many electronic systems, including Automated Teller Machines (ATMs), computer passwords, and home security systems. Additionally, cryptology, in the form of the Secure Socket Layer and the Transport Socket Layer protocols - which are built into most web browsers - are used to secure the transmission of financial information between a customer and a vendor during online transactions (Piper & Murphy 2002, pp.130-132). Cryptographic technologies are therefore essential for the prevention of fraud, and other criminal activity, in a world increasingly reliant on electronic communication networks.

However, by the same logic, cryptographic technologies can also be used by criminals, terrorists, businesses, and other actors, to hide their communications and conceal their activity, thus making it harder for law enforcement bodies and intelligence organizations to detect any wrongdoing. This created a tension between the apparent benefits of widespread access to cryptography, conceived of in terms

³‘Modern’ cryptography is used here to refer to developments in the field from the 1970s onwards, particularly those associated with Public-Key Cryptography.

of individual privacy and fraud prevention, and the apparent risks associated with widespread access to cryptography, conceived of in terms of the problems it created for effective law enforcement. When this tension arose, governments in many countries around the world decided that they needed a policy to resolve it. The ‘crypto wars’ - the name most commonly used to refer to the political debates over these issues - emerged out of disagreements over the various governmental attempts to regulate the access to and use of cryptographic technologies during the 1990s.

1.2.1 Why Cryptology and the Crypto Wars?

Before saying anything further about cryptology and the crypto wars, I will explain why I have chosen to use it as a case study to speak about the third wave. There are four main reasons. The first of these has already been alluded to. It has been observed that the third wave debate, and indeed the wider debate about expertise and controversies over technological decision-making, would benefit from being able to draw upon a larger pool of case studies (Owens 2011). Furthermore, given that many of the existing case studies are based on controversies that intersect with the natural sciences, the debate would benefit from the introduction of case studies from the mathematical sciences, such as computing and cryptology.

Secondly, cryptographic technologies, though rarely discussed, are undoubtedly important to modern Western societies. Lawrence Lessig has argued that cryptographic technologies constitute “the most important technological breakthrough in the last one thousand years”, and that “no other technological discovery - from nuclear weapons (I hope) to the Internet - will have a more significant impact on social and political life” (Lessig 1999, pp.35-36). Though Lessig may be guilty of exaggerating the importance of cryptography, as electronic communication networks become ever more prevalent, its significance cannot continue to be ignored.

This links to a third reason. Although this case study will focus on events that occurred in the 1990s, the issues at the heart of the crypto wars are very much alive. Questions over the appropriate level of access to, and appropriate use of cryptology, continue to be raised. Issues surrounding cryptology were central to one of the biggest news stories in second decade of the twenty-first century - the leaking of secret intelligence information by former National Security Agency (NSA) employee Edward Snowden. Although this thesis will not discuss these

events, and it will be some time before the dust has settled and they can be analysed by the academic community, it has been widely reported in the press that, since 2000, intelligence organizations have secretly attempted to build flaws into commercially available cryptographic technologies as part of a coordinated effort to read encrypted information (Perlroth et al. 2013). Now that this information has been made public, it is clear that the issue of whether it is appropriate for intelligence organizations to carry out this kind of cryptology work without public scrutiny is one that needs addressing.

Finally, as will become clear from the overview of cryptology and the crypto wars that will follow, this was a controversy that featured a broad array of expertise. This expertise was produced in a variety of different contexts and for a variety of different purposes. It was produced within academia, the civil service, industry, and intelligence organizations. But more importantly, cryptology expertise has in the past been singled-out as unusual compared to that produced by other scientific disciplines. The reason for this is that it is closely bound up with secrecy. Two lengthy quotes from eminent cryptologists - both of whom will be discussed in later chapters - will help illustrate this. In the first, Ross Anderson argued that:

The practice of cryptology differs from, say, that of aeronautical engineering in a rather striking way: there is almost no public feedback about how cryptographic systems fail. When an aircraft crashes, it is front page news. Teams of investigators rush to the scene, and the subsequent enquiries are conducted by experts from organisations with a wide range of interests - the carrier, the insurer, the manufacturer, the airline pilots' union, and the local aviation authority. Their findings are examined by journalists and politicians, discussed in pilots' messes, and passed on by flying instructors. In short, the flying community has a strong and institutionalised learning mechanism. This is perhaps the main reason why, despite the inherent hazards of flying in large aircraft, which are maintained and piloted by fallible human beings, at hundreds of miles an hour through congested airspace, in bad weather and at night, the risk of being killed on an air journey is only about one in a million. In the crypto community, on the other hand, there is no such learning mechanism. The history of the subject

shows the same mistakes being made over and over again; in particular, poor management of codebooks and cipher machine procedures enabled many communication networks to be broken (Anderson 1994).

In the second, James H. Ellis reflected that:

Cryptography is a most unusual science. Most professional scientists aim to be the first to publish their work, because it is through dissemination that the work realises its value. In contrast, the fullest value of cryptography is realised by minimising the information available to potential adversaries. Thus professional cryptographers normally work in closed communities to provide sufficient professional interaction to ensure quality while maintaining secrecy from outsiders. Revelation of these secrets is normally only sanctioned in the interests of historical accuracy after it has been demonstrated clearly that no further benefit can be obtained from continued secrecy (Ellis 1987).

Although Robert K. Merton (1973) identified ‘communalism’ as one of the norms that underpins science, subsequent work within the second wave of science studies has shown that scientific research can also be guided by secrecy (Mirtoff 1974, Rappert 2009, Balmer 2012). Secrecy can be used to protect research findings in order to establish priority, to maintain a commercial advantage, or to prevent other nations from knowing about defence capabilities. However, judging by the above quotations, it is clear that the role of secrecy in cryptology research - at least in the minds of cryptologists - appears to be much more central than in other fields. Bent Flyvbjerg (2006) has argued that such ‘extreme’ cases are particularly useful in case-study research because they activate a greater number of actors and processes. For example, Diane Vaughan’s (1996) study of the 1986 Challenger space shuttle explosion, though focussed on a very rare event, provided insights into the culture of large organizations like NASA, the role of small technologies in large projects, the unfolding of public inquiries, and the public role played by high-profile scientists. Given that secrecy is intertwined with cryptology to a greater degree than in most other fields, a study of cryptology may be valuable because it can be used to highlight processes that may be present to a lesser degree in other science and technology case studies.

Plaintext:	H	E	L	L	O	W	O	R	L	D
Key:	+3	+3	+3	+3	+3	+3	+3	+3	+3	+3
Ciphertext:	K	H	O	O	R	Z	R	U	O	G

FIGURE 1.1: An Example of a Simple Substitution Cipher

1.2.2 The History of Cryptology

Before discussing modern cryptology and the crypto wars in more detail, it will be useful to say a little more about the history of the field and, in the process, introduce some of the basic technical concepts. The history of cryptology, from its invention in ancient times, to its use during the Second World War, has been well described (e.g. Kahn 1997, 1991, Singh 1999, Sebag-Montefiore 2001). In his definitive 1,200 page history of cryptology, David Kahn (1997) comprehensively described the development and use of cryptology during this long period. Up until the twentieth century, cryptography was primarily used to encipher messages so that individuals could communicate with each other in secret. In one of his numerous historical examples, Kahn described the famous Caesar cipher - named after the cryptographic techniques used by Julius Caesar to communicate with his centurions. The Caesar cipher is an example of what's known as a 'simple substitution cipher'. Here, as in most other examples, it is assumed that someone wants to send a message to someone else.⁴ But, they do not want anyone but the intended recipient to be able to read it. If an interceptor sees the message before it reaches the recipient, it should be unintelligible to them. To create a simple substitution cipher, the ciphertext (the encrypted message) is produced by substituting each letter in the plaintext (the unencrypted message) with the letter, say, three places along in the alphabet. In this example, any 'A' would be substituted with a 'D', and any 'D' would be substituted with a 'G' (any 'X', 'Y' or 'Z' would be replaced with an 'A', 'B' or 'C', as the end of the alphabet joins back to the beginning).

In Figure 1.1, we can see that, using a simple substitution cipher, the message 'HELLO WORLD' would be encrypted to 'KHOOR ZROUG'. The process used to convert the plaintext into the ciphertext, in this case a shift of three places

⁴Most explanations of basic ideas in cryptology use 'Alice' and 'Bob' to refer to the parties wishing to communicate. I don't find this naming convention particularly helpful, as the names used are non-descriptive. Therefore, I will typically use 'sender', 'recipient' and 'interceptor' instead.

Plaintext:	H	E	L	L	O	W	O	R	L	D
	C	R	Y	P	T	O	G	R	A	M
Key:	+3	+17	+25	+16	+19	+15	+7	+17	+1	+13
Ciphertext:	K	V	K	B	I	L	V	I	M	Q

FIGURE 1.2: An Example of a Vigenère Cipher

forward in the alphabet, is known as the ‘key’. The ciphertext can be converted back into plaintext using a process that is the inverse of the process used to encrypt it, in this case, a shift of three places backwards in the alphabet. Due to the symmetry of the encryption and decryption processes, the Caesar Cipher is an example of what’s known as ‘symmetric’ cryptography.

This is a very simple example. More elaborate methods of enciphering messages were developed between Caesar’s lifetime and the twentieth century. A notable example is the Vigenère Cipher - where a message is enciphered based on the position of letters in a keyword in order to further distance the ciphertext from the plaintext (see Figure 1.2). The Caesar cipher is an example of what’s known as a ‘monoalphabetic’ cipher, whereas the Vigenère Cipher is an example of what’s known as a ‘polyalphabetic’ cipher. The famous Enigma Machine - a cryptographic device used to send messages by the German armed forces during the Second World War - though incredibly complex, was essentially based on a polyalphabetic cipher.

1.2.3 Public-Key Cryptography

Modern cryptology can be said to begin with the development of public-key cryptography. In the past, symmetric ciphers had a clear but limited use in certain contexts, such as military and intelligence work, and the private correspondence of interested individuals.⁵ The advent of networked computer systems in the late 1960s and early 1970s, together with rapid developments in computer hardware, prompted a rethink of what cryptology could be used for. This process began when Horst Feistel, a researcher at IBM, proposed a system whereby messages could be represented by a sequence of binary digits, which could in turn be enciphered at a speed and complexity that would have been impossible to achieve manually. Feistel’s algorithm, which he called Lucifer, was identified as a way

⁵Many prominent historical figures, such as Benjamin Franklin and Lewis Carroll, have displayed an interest in using cryptography to communicate in secret (see Kahn 1997).

of securing potentially sensitive financial information transmitted by the newly introduced network of ATMs. Aware of the increasing demand for secure systems of this type, the National Bureau of Standards (NBS) in the US selected Lucifer as the algorithm to underpin their new Data Encryption Standard (DES) in 1975 (Blanchette 2012, pp.34-35).

Importantly, no matter how complicated the encipherment process, for symmetric ciphers, the ciphertext can always be converted back into plaintext if the key is known. Whilst statistical cryptanalysis techniques can also be used to determine the key used given certain conditions, in practice, for symmetric ciphers to be useful, the recipient must know the key in order to be able to read the message. This creates what's known as the key distribution problem - namely, how does the sender make the recipient aware of the key without also making the interceptor aware of it at the same time? Until the 1970s, the insolubility of this problem was believed to be one of the fundamental tenets of cryptology.

Then, in the mid 1970s, cryptology changed radically with the development of 'asymmetric' or 'public-key' cryptography. Although developed independently in a number of contexts, public-key cryptography is usually associated with the work of two American computer scientists - Whitfield Diffie and Martin Hellman. In 1976, they published a paper - with the bold title of 'New Directions in Cryptography' - that proposed techniques for solving the key distribution problem, and in the process, expanded the potential uses of cryptography (Diffie & Hellman 1976). Diffie and Hellman's system was based on both the sender and the recipient having their own key, and it being split in two - with one part made publicly accessible, and the other part kept private. Thus, when the sender wished to send a message to the recipient, they could encrypt it using the recipient's public key. Due to a one-way mathematical relationship between the public and private keys, only the intended recipient would have the correct private key to successfully decrypt the message. Under this system, private keys do not have to be exchanged prior to the communication. Furthermore, the internal logical of the system also allowed Diffie and Hellman to conceive of cryptography being used to provide integrity and authentication, as well as confidentiality.

Although Diffie and Hellman's proposed system lacked the one-way mathematical function required to make it work, a solution was provided two years later by Ron Rivest, Adi Shamir and Leonard Adleman (Rivest et al. 1978). Rivest, Shamir and Adleman proposed a system - now known as RSA based on the initials of

the authors' surnames - built around the intractability of finding the prime factors of large numbers. Multiplying prime numbers requires little computational effort, even if the numbers are very large. However, reversing the process, starting with the product only, requires so much computational effort that for very large numbers, it is practically (though not theoretically) impossible. Though it is unnecessary to state here the exact mathematical processes required for RSA key generation, it is sufficient to appreciate that if two large random prime numbers are multiplied, it is possible to reveal partial information about the result, in the form of a public key, whilst retaining, in the form of a private key, the unique information required to quickly reverse the process. Thus, when integrated with Diffie and Hellman's system, the RSA algorithm enabled the sender to encrypt using the recipients public key, whilst allowing the recipient (and the recipient only) to decrypt it using their private key.

As has already been mentioned, developments in public-key cryptography, and the advent of large-scale computer networks (see Abbate 1999), raised questions over what the appropriate level of access to cryptographic technologies should be. By the 1990s, it was clear that whilst these technologies could be used to uphold individual privacy, and could be key to ensuring trust in electronic commerce, they could also be used to conceal criminal activity. When this tension arose, governments around the world decided that a policy on cryptography was needed. Policies typically attempted to determine whether the potential benefits of widespread public access to cryptographic technologies outweighed potential disadvantages. Debates ensued over whether governmental policies had balanced these concerns correctly, and whether their policy solutions were technologically or politically feasible. These debates took place throughout the 1990s, at a time when the Internet was starting to be used by the general public. Although the crypto wars are usually thought of as particular set of debates that occurred during this period, it is clear that debates over the central issues have continued through to the present day, and it is likely that they will continue.

1.2.4 The Literature on the Crypto Wars

The history of cryptology up to and including the Second World War has been well described, and is a thriving albeit niche academic discipline.⁶ However, the

⁶This history of cryptology has its own dedicated journal - *Cryptologia*. The articles in this journal tend to focus on pre-Second World War cryptology, and are often highly technical. This

current body of non-technical literature, on both cryptology research after 1945 and the crypto wars, is small. Neither cryptology research nor the crypto wars are discussed in any of the major surveys of the history of computing (Ceruzzi 2003, Campbell-Kelly et al. 2013), the history of software (Campbell-Kelly 2003), or the history of the relationship between computers and industry (Ensmenger 2010). This reflects a broader absence in the historical literature of studies of technologies associated with computer security. This situation is largely mirrored in the STS literature. With the notable exception of Jean François Blanchette's *Burdens of Proof* (2012) - which will be discussed later on in this chapter - I am not aware of any published sociological studies of the science of cryptology, or cryptography as a technology.⁷ Interestingly, despite this, some cryptology insiders appear to think of the crypto wars literature as satisfactorily complete. For example, reflecting in 2011, the noted cryptographer Matt Blaze claimed that "the history of the 1990's 'crypto wars' has been well-chronicled" (Blaze 2011, p.238). Statements such as this, in conjunction with the current state of the literature, highlight a disparity between what insiders know about the crypto wars, and what outsiders are aware of or are able to learn.

Though crypto wars debates took place within many national contexts, in the non-technical literature on cryptology and the crypto wars that does exist, most have placed their focus on events that occurred in the United States. Though, in general, studies have not opted for an explicit national framing, most appear to have focussed their attention on the US by default because it was home to key technological developments, its government was the first to attempt to confront the challenges that cryptology posed, and because their policy solutions relied upon, in the opinions of some, rather draconian measures.

In the three best known works on the crypto wars in the US, Steven Levy (2001), Simon Singh (1999), and Whitfield Diffie and Susan Landau (2007) all argued that, as electronic communication became ever more ubiquitous in the second half of the twentieth century, the development of cryptography technologies offered a way for privacy to be upheld in light of the increasing prevalence of surveillance. In

partly stems from a belief that a technical understanding of the history of cryptology can be used to inform modern cryptology research, given that the secrecy that has surrounded the discipline may have inhibited the circulation of some key concepts.

⁷This absence can be demonstrated by searching for "crypto wars" using Google Scholar. In October 2013, this search returned under one hundred academic journal articles, books, and conference papers. Furthermore, most of the sources that were returned refer only to the crypto wars in passing, or as a way to contextualize a piece of current scientific research. Only a handful are primarily concerned with discussing the events themselves.

Crypto, Levy, a journalist for the technology magazine *Wired*, described a David vs. Goliath story in which cryptology scientists and enthusiasts (which Levy referred to as “code rebels”) were successful in defeating the US government’s attempts to ‘control’ cryptography. Here, engaging in cryptology research is synonymous with resisting surveillance, and is bound up with cyber-libertarianism. In setting the tone for his book, Levy argued:

Doesn’t the advent of computer communications means that everyone should have access to the sophisticated tools that allow the exchange of words with lawyers and lovers, coworkers and customers, physicians and priests with the same confidence granted to face-to-face conversations behind closed doors? This book tells the story of the people who asked those questions and created a revolution in the field that is destined to last all our lives. It is the story of those who did their best to make these questions go away. The former were nobodies: computer hackers, academics and policy wonks. The latter were the most powerful people in the world: spies, generals and presidents. Guess who won (Levy 2001, p.2).

This sentiment is partially mirrored in Singh’s *The Code Book* (1999). Essentially an introduction to the science and history of cryptology, Singh, echoing Levy’s tone, enthusiastically walks the reader through the basic concepts and techniques before describing a similar story. In *Privacy on the Line* (2007), Diffie and Landau also described the crypto wars in terms of the US government’s reliance on wiretapping techniques for the purposes of law enforcement - techniques that widespread access to cryptography technologies threatened to render obsolete.⁸

It has been noted that these arguments are often built upon assumptions about the potential for government abuses of surveillance rather than actual cases, and

⁸In addition to these three main works, there are a handful of other less well-known descriptions. Olivia Bosch (2005), in examining the related issue of cryptography export, sought to understand the role of industry in the eventual changes to US export policy. She concluded that the eventual ending of US attempts to regulate cryptography export could be best understood in terms of technological (rather than political) factors, such as its declining role in information security policy and its infrequent use for purposes of confidentiality. Also worthy of a mention is Bruce Schneier and David Banisar’s *The Electronic Privacy Papers* (1997). This book is an edited collection primary source documents related to the crypto wars in the US. Rather than being a neutral collection of material, the documents chosen for inclusion come primarily from US government sources, and appear to have been collected in an attempt to highlight political machinations.

that there is perhaps a sense of “misplaced Big Brotherism” that stems from the authors’ belief in the importance of individual privacy (Staples 1998). Blanchette has noted that cryptology insiders are keen to invoke such descriptions, not only because it casts them as the heroes, but also because they tend to fit with a narrative that allows them to “portray the evolution of the field as that of maturation from a craft to a bona-fide mathematical science, spurred by the embrace of mathematical formalism in the second half of the twentieth century” - a story that, as will be described later in this chapter, Blanchette challenges (Blanchette 2012, pp.17-39).

Though less well-known, the best description of the crypto wars in the US that presently exists has been provided by Gus Hosein (2003). Hosein, in touching on ideas from STS, used a case study of the crypto wars to argue that regulatory discourse could be best understood through consideration of both the technological and the social actor. In describing the crypto wars in the US, Hosein divided the government’s responses to the development of public-key cryptography into three chronological stages: intervening in the algorithms; managing access to keys; and controlling the form and nature of the keys (Hosein 2003, p.137). From 1979 to 1993, shortly after the development of public-key cryptography, the US government attempted to regulate cryptography by ‘intervening in the algorithms’. The NSA argued that strong encryption was fundamentally harmful to national security, and used this to justify attempts to prohibit the publication of developments in cryptology through the screening of research and the acquisition of patents. In the case of cryptology research carried out within US universities, this strategy was partially successful. Despite this approach, in 1991, towards the end of this phase, an amateur programmer named Philip Zimmermann uploaded a piece of software called Pretty Good Privacy (PGP) to the Internet. PGP, which is still widely-used, is a computer program that can be used to encrypt electronic emails using public-key cryptography. Zimmermann had developed PGP at home in his spare time. He was not a scientist, nor was he working at a university or a research environment. Therefore, as well as being one of the first usable implementations of public-key cryptography, the development of PGP, and the manner in which it was produced, signalled the extent to which cryptography had become publicly available. At the time, exporting cryptography technologies from the US still required a licence from the government, as it was seen as having potential military uses. On this basis, attempts were made to prosecute Zimmermann for ‘exporting’

cryptography software using the Internet, but the case was dropped after three years.

In 1993, with attempts to intervene in the development of cryptology partially unsuccessful, Bill Clinton's newly-elected Democrat administration attempted to introduce the Clipper Chip, and in doing so, initiated a phase of the crypto wars marked by attempts to 'manage access to keys'. The now-notorious Clipper Chip was based on the principle of key escrow. Escrow is a legal term that refers to the practice of handing over information or goods to a third party, so that other authorised parties may also gain access to them. Clipper was therefore an electronic chip designed by the NSA to be implemented in a number of devices, such as telephones and modems. It provided strong encryption, but also provided a what's known as a backdoor - essentially a built-in security flaw - which the US government would be able to exploit in order decrypt communications if they wished. The Clipper proposal was controversial and unpopular. It was heavily criticised by organizations such as the Electronic Frontier Foundation, and provoked a strong negative response from many cryptologists. What's more, in August 1994, Matt Blaze (1994) published a paper detailing a serious vulnerability in the scheme that could be used to disable the key escrow function. However, the most serious problem that the Clipper Chip faced was that industry expressed little interest in adopting devices that housed it, and a convincing case could not be made for them to incur the extra costs of doing so. The Clipper Chip proposals were dropped.

Following the failure to promote the adoption of the Clipper Chip, the US government instead set about 'controlling the form and nature of the keys'. Though attempts to integrate Clipper in hardware form were abandoned, the US government attempted to introduce key escrow in software form. These proposals were equally unpopular with the cryptology community. Eventually, in an attempt to drum up worldwide support for key escrow, the US government lobbied the Organization for Economic Cooperation and Development (OECD), and raised the issue at the 1996 G7 summit in Paris. Despite receiving support from some members of the G7, the problems over industry adoption remained. During the protracted debates over Clipper, many industries had adopted standards of their own. The software Clipper proposals were also dropped. This marked the end of the debates on this particular issue in the US. Whilst the US government had failed to establish a key escrow-based standard, they had partially succeeded in

stalling the development of open cryptology research, particularly during the early years.

The literature on the crypto wars, then, is dominated by descriptions of events that occurred in the US - events that took place within a particular social and political context. However, parallel crypto wars debates occurred in a number of other countries around the world. Bert-Jaap Koops' (2013) Crypto Law Survey project provides an overview of the legal and political developments related to cryptography in just under one hundred countries. Of course, it is true that these national debates did not take place in isolation from one another. They overlapped chronologically, in terms of the broad issues discussed, and occasionally featured the same actors. What's more, activity relevant to the debates was often mobilized using the Internet, and ideas about the independence of this activity from that of nation states - exemplified by John Perry Barlow's *Declaration of the Independence of Cyberspace* (1996) - had wider currency than they do now. However, as Jack Goldsmith and Tim Wu (2006) have argued, although there has in the past been a tendency to view the Internet, and the debates surrounding it, as relating to a borderless global community that defies attempts at control from nations, the importance of governmental coercion at a national level has often been seriously underestimated. Thus, arriving at an understanding of other national crypto wars debates, and, indeed, a more general understanding of the crypto wars across the world, based solely on what's known about the crypto wars in the US, is likely to be a flawed strategy.

Though it is known that they took place, only a handful of national debates outside of the US have received any reflective academic attention (e.g. Koops 1998, Parviainen 2000). As well as providing the best description of the crypto wars in the US, Hosein (2003) has also provided the only attempt that I am aware of to describe events in the UK. Using (what I would call) a comparative ethnography of the crypto wars in the US and the UK, Hosein argued that policy discourses and the dynamics of regulatory change are best understood by considering both social and technological actors. To support this claim, Hosein devoted a chapter of his doctoral thesis to providing a description of the political and regulatory developments related to cryptography in the UK. Hosein essentially described the attempts by successive UK governments to gain access to the public's encryption keys for the purposes of law enforcement during the 1990s. He broke this process down into two phases. In the first, John Major's Conservative government

attempted to gain access to keys by attempting to link a key escrow scheme to the benefits of electronic commerce. In the second, after the first initiative had failed, Tony Blair's Labour government successfully passed legislation that allowed direct lawful government access to encryption keys. During the first phase, the Conservative government proposed the establishment of Trusted Third Parties (TTPs) - state-licensed organizations that would provide encryption services to the public. It was hoped that the mandatory government-run licensing scheme would improve public trust in electronic commerce, as it would allow people to feel confident about performing financial transactions over the Internet. However, TTPs would also be required by law to keep keys in escrow, and would be further required to make them available on demand to law enforcement and intelligence bodies. Hosein then moved on to describe how, following Labour's election victory in May 1997, law enforcement priorities were 'de-coupled' from electronic commerce. The proposals for a mandatory TTP licensing scheme were replaced with proposals for a voluntary one. When these proposals also proved unpopular, the government changed tack, and instead pursued a policy based on new powers that allowed law enforcement bodies to demand keys from citizens. This initiative was successful, and resulted in the controversial Regulation of Investigatory Powers Act 2000 (RIPA).

Though Hosein's description covers many of the important political developments that took place during the UK crypto wars, it is incomplete. In particular, the crypto wars are isolated from society, and perhaps more importantly, from cryptology expertise. The political processes, particularly the arguments used in parliament - which are quoted extensively - are disconnected from the relevant cryptology research - which is not described at all. This isn't, by itself, a criticism of Hosein's approach. After all, Hosein was not concerned with understanding the relationship between expertise and technological decision-making, but with regulatory discourse. This understandably placed an emphasis on the political and rhetorical activity associated with the crypto wars. However, this framing served to isolate the crypto wars from the expertise that underpinned the political arguments being made. As it was not relevant to his research, Hosein did not consider how and why the expertise that underpinned the political arguments had been produced, and instead focussed on providing detailed descriptions of the rhetoric used. In short, for Hosein, cryptology expertise arrived ready-made.

If anything unites the literature on the crypto wars, whether focussed on events

in the US or the UK, it is this lack of a proper description of how the relevant cryptology expertise was produced. The implication from the sources that focussed on events in the US is that cryptology expertise emerged primarily out of an academic cyber-libertarian movement. Under this view, cryptology research was undertaken because electronic communication threatened individual privacy. Recently, Blanchette (2012) has partly challenged this view. Though Blanchette specifically stated that he did not wish to return to the crypto wars debates of the 1990s, he does provide some valuable insights on the nature of cryptology research. In attempting to understand why digital signatures - a particular technology based on cryptography - have failed to gain wider acceptance, Blanchette made three overlapping arguments about the complex nature of cryptology research. Firstly, he argued that the characterization of cryptography and digital signatures by researchers as immaterial has made their translation into hardware and software artifacts problematic. Secondly, that attempts to mathematize certain areas of cryptography, with the aim of providing provable security, have marginalized areas of research that, although resistant to mathematization, have the potential to deliver a greater social impact. Thirdly, that the way in which cryptographers have modeled digital signatures has served to obscure the trade-offs inherent in producing technologies that are to function in the real world. Though not directly relevant to my research, through these arguments, Blanchette at least highlights that it is possible to gain an understanding of the nature of cryptography research, and that the associated practices and assumptions are more complex than the impression thus far given in the crypto wars literature.

1.3 Research Questions and Answers

Keeping what's been said thus far about the third wave and the crypto wars in mind, this thesis is structured around formulating and answering the following three sequential research questions. The questions aim to probe both the creation of cryptology expertise and its relationship to the crypto wars, and to use this as a platform to examine some of the concepts that underpin the third wave.

1. How was contributory cryptology expertise produced during the technical phase of the crypto wars in the United Kingdom?

My answer to this question will be provided in the form of four descriptions of the cryptology research that was carried out at four separate UK sites. I will describe the nature of the research carried out at: the Data Security Group at the National Physical Laboratory; the Information Security Group at Royal Holloway, University of London; the Security Group at the University of Cambridge; and the Communications-Electronics Security Group (CESG) at the Government Communications Headquarters (GCHQ). Each individual description highlights the different practices used to produce expertise in cryptology from 1970 - when cryptology research began outside of a military or intelligence context, to the mid-1990s - when the political debates over cryptology policy began. I will describe how the Data Security Group at the National Physical Laboratory, following reforms that commercialized government research establishments, came to produce cryptology expertise geared towards the development of standards and accreditation schemes. The Information Security Group at Royal Holloway, following changes to the structure of the University of London and an emphasis on industrial collaboration, produced expertise in the mathematics of cryptology for technological solutions. The Security Group at the University of Cambridge, in line with a traditional departmental emphasis on systems research, developed expertise on how cryptography technologies function as part of real-world security systems. Finally, CESG at the Government Communications Headquarters, though responsible for providing cryptology expertise to public bodies, thanks to their organizational links to GCHQ and their intelligence priorities, produced cryptology expertise designed to be secret. What emerges from these descriptions is an overview of, in a short period of time, just how diverse contributory expertise in cryptology became prior to the crypto wars.

2. Can the ontological framework and the third wave be used in conjunction to develop a re-conceptualization of the production of this contributory expertise as ‘multiple’?

My answer to this question is a qualified ‘yes’. Although the concept of interactional expertise has received much attention from those working under the third wave, contributory expertise has been thought of as unproblematic. However, given that the third wave recognises that contributory expertise is rooted in collective practices, it follows that a variation in the practices used to produce contributory expertise within a particular discipline has the potential to produce a

variety of different expertises. To probe this possibility, I draw upon recent work within STS on the ‘ontological framework’. Although largely separate from the third wave in the literature, the ontological framework, like the third wave, emphasises the productive power of collective practices, and holds that the results of these practices are real. One particular strand of the ontological framework has stressed that divergent practices will ‘enact’ multiplicity. Though this multiplicity is partially discernable in the descriptions that make up the answer to the first research question, I argue that it can be demonstrated more clearly through a formalised process of sociological discrimination - a particular type of expertise as defined by previous work under the third wave. However, whilst ontological multiplicity can be identified in this fundamental sense, it must also be recognised that some of the more detailed concepts from the ontological framework do not align well with ideas from the third wave. It is for this reason that I have described my answer as a qualified ‘yes’. In particular, concepts from the ontological framework that aim to understand how divergence amongst enactments is resolved, do not map well onto political attempts to deal with and interpret contributory expertise during the crypto wars. Furthermore, given that there appears to be unresolved philosophical contradictions inherent in adopting extreme forms of some ontological ideas, it is sensible to postpone further alignment with the third wave.

3. What were the consequences of this multiplicity of contributory expertises during the political phase of the crypto wars?

In response to this question, I argue that once recognised, multiplicity of contributory expertise should be taken seriously by the third wave if the aim is to “make the best decisions in the right way” during controversies over technological decision-making. This is because, during the crypto wars, the multiplicity amongst contributory expertise, was, to a certain extent, translated across to the political decisions that were eventually made. Each enactment of contributory cryptology expertise was used for different and often distinct purposes. The contributory expertise enacted at the National Physical Laboratory was not drawn upon. The expertise enacted at Royal Holloway was used to provide a technology to underpin the government’s controversial proposals. The expertise enacted at the University of Cambridge, though initially ignored, was used to construct an opposition to these proposals. Finally, the expertise enacted at CESG, given its secret nature, both contributed to the formation of the proposals, and added an extra layer of

uncertainty to political decisions. Importantly, I argue that the nature of each enactment, and its eventual use, were linked. Therefore, from a positive point of view, acknowledging multiplicity offers a tool for political decision-makers to map out sources of available expertise, and given the diverse array of propositional questions that characterize controversies over technological decision-making, to make more informed choices when seeking out expertise. From a negative point of view, acknowledging multiplicity also highlights some of the problems with elective modernist principles. These problems are rooted in the fact that some enactments of contributory expertise, if used, can distort political decisions during a controversy. Therefore, the choices made about what expertises are used to inform decisions become important. In the case of the crypto wars, the use of the secret expertise enacted by CESC meant that it was difficult for other experts to assess its quality, and that more generally, it was difficult for observers to determine CESC's role in the formation of the proposals. Given that debates over technical issues were partially ruled out, speculative arguments over the quality of CESC's expertise, and the motivations behind it, came to be more important. In parallel, the initial exclusion of the expertise enacted at the University of Cambridge prompted them to take steps to have their expertise reflected in political decisions. This 'scientist activism', though it may have exerted a positive influence on political decisions, resulted from the decision to exclude certain expertise during the formation of the government's proposals, and threatened to undermine an essential third wave distinction between the formative intentions of science and politics.

1.4 Chapter Guide

How these arguments were arrived at will be described over the course of nine further chapters. In chapter 2, I survey the literature related to the two main theoretical touchstones of this thesis: the third wave understanding of expertise during controversies over technological decision-making; and the ontological framework. I will use this chapter as a way of laying the theoretical ground for what will follow, and to explain in more detail how my three research questions were arrived at. I will show that case studies of controversies, though in the past used to better understand the creation of scientific facts and technological artefacts, are now often used to understand the relationship between expertise and technological decision-making. The third wave has attempted to partially move beyond

describing controversies over technological decision-making, to formulating normative prescriptions relating to how they should function. The third wave has produced a robust classification scheme for different types of expertise, and has attempted to extend the implications of this scheme to technological-decision making. However, some of the principles that underpin elective modernism are still being developed, and would benefit from refinement. In particular, prescriptions relating to the proper relationship between controversies and contributory expertise are in need of attention. One way for researchers to try and get an intellectual foothold on this particular issue is by drawing on the ontological framework. I will therefore describe how the ontological framework has been developed within STS. I will show that the ontological framework has stressed the importance of multiplicity - the idea that divergent research practices will enact a multiplicity of objects. I argue that, though the full philosophical implications of adopting the ontological framework are best avoided for the moment, it can still be used to probe the production of contributory expertise, and therefore prompts consideration of the implications of multiple contributory expertises for prescriptions on the relationship between science and politics during a controversies over technological decision-making.

In chapter 3, I will describe the methods used to answer my research questions. Given that there is often little reflection on the methods used in STS research, I will situate my methodological choices within the broader literature on research in the social sciences. I will describe the processes used to gather data on the production of cryptology expertise, and on how this expertise was used during the political debates that made up the crypto wars. Put simply, I collected the data for this study using a combination of documentary analysis and retrospective semi-structured interviewing. I argue that the historical nature of the research methods used prompts concerns over what documents and interviews are able to reveal about the past. Given that the historical nature of these methods departs from the ethnographic methods often used when investigating research practices, I will argue that this necessitates a shift in focus from observable, micro-social processes located at the level of the laboratory, to unobservable, macro-social processes, located at an institutional level.

In terms of data collection and analysis, my study can be divided into three stages (see Table 1.1). Documentary analysis and retrospective semi-structured interviews were used to collect data relating to the production of cryptology expertise

Stage	Data Collection	Analysis	Research Activity
First	Documentary analysis; Retrospective semi-structured interviewing	Thematic	Collection of data relating to the production of contributory cryptology expertise at various sites within the UK between 1970 and 2000
Second	Documentary analysis; Retrospective semi-structured interviewing	Thematic	Collection of data relating to the development of UK legislation related to cryptology between 1970 and 2000
Third	None	Concept and theory development	Attempts to understand how expertise from research sites was transferred to, and used during, political processes

TABLE 1.1: Research Design

at various research sites. Then, similar documentary analysis and semi-structured interviewing techniques were used to examine how that expertise was used during the political debates. A third and final stage of theoretical and conceptual development was also undertaken in order to define and then carry out a process of sociological discrimination that can be used to demonstrate multiplicity, to probe how expertise was transferred from research sites to the political arena, and to consider what the ideal form of this process might look like in light of the existence of multiplicity.

In chapters 4 through 7, I will describe in detail what emerged from the data that was collected for the purposes of the first stage of fieldwork (see Table 1.1). More specifically, each chapter will describe the nature of the cryptology expertise produced at one of the four aforementioned cryptology research sites: the Data Security Group at the National Physical Laboratory; the Information Security Group at Royal Holloway; University of London; the Security Group at the University of Cambridge; and CESG at the Government Communications Headquarters. In each case, an historical description will be provided of the institutional context within which cryptology research was initiated, and how this context shaped the nature of that research thereafter. The nature of the cryptology research carried out at each site will be described from its initiation to the start of the crypto wars in the mid-1990s. Taken together, these chapters form the basis of my answer to the first research question.

In chapter 8, I describe the data that was collected for the purposes of the second stage of fieldwork (see Table 1.1). As such, I describe the political phase of the

crypto wars. In contrast to previous descriptions, I describe the politics of the crypto wars in terms of its relationship to cryptology expertise. I divide the crypto wars into two overlapping parliamentary bills. I describe the relationship between cryptology expertise and the Department of Trade and Industry proposals for a nation-wide cryptology network of Trusted Third Parties - organizations licensed to provide cryptology services to the public, but also required by law to provide the means to decrypt the communications they handled to law enforcement bodies. I also describe the relationship between cryptology expertise and proposals to license the export of cryptography from the UK by intangible means. Together, these debates highlight how the government used the available cryptology expertise, the consequences that this had for the nature of the crypto wars debates, and how this shaped the eventual policy decisions.

In the first half of chapter 9, I will bring the content of the preceding chapters together to demonstrate the multiplicity of contributory expertises that were enacted at the four research sites. To do this, I will describe and then carry out a process of sociological discrimination. Moving from the specific descriptions of expertises described in chapters 4 through 7, I will categorise the activity of each research group in terms of the more general categories of: the production of certified knowledge; education and training activities; public research and the innovation process; the participation in public or collective goods and finalities; and public debates about science and technology. This provides a method of operationalizing sociological discrimination that clearly demonstrates that contributory cryptology expertise was multiple. This, together with the descriptions from chapters 4 through 7, will form the basis of my answer to the second research question. In the second half of this chapter, I will describe how this multiplicity of contributory expertises, was, to a certain extent, translated across to political decisions. I will show how each enactment of contributory expertise was used for different and often distinct purposes, and that these uses were linked to the practices used to enact it. To add an extra layer of detail, I also describe a method of differentiating between how expertise, once enacted, can be transferred to those responsible for making political decisions. In the case of the crypto wars, expertise was either: not transferred; commissioned; delivered; or transferred invisibly. This, together with the description provided in chapter 8, is used to arrive at an appreciation of some of the consequences of multiplicity of contributory expertise during the crypto wars. This forms the basis of my answer to the third research question. In light of these consequences, I close this chapter with some suggestions for how the third wave

may be modified in order to accommodate them. I suggest what I have called the Minimum Transfer Requirement - a method of determining whether it can be claimed that political decisions have given adequate consideration to the available contributory expertise. I also identify the problem of expert discrimination - a recognition of the fact that it may not always be possible to publicly assess the quality of the expertise used to underpin decisions during controversies over technological decision-making.

In chapter 10, the concluding chapter, I will reflect upon the answers given to the three research questions. I will use this as a basis for identifying limitations, and for making suggestions for further work. From this, it will be clear that the limitations identified, in both my answers and the third wave more generally, could be addressed through carrying out more empirical case studies viewed through the lens of the third wave.

Chapter 2

Literature Review

2.1 Introduction

In this chapter, I will review the literature relevant to the themes outlined in chapter 1. I have divided this chapter into two sections. The first section provides a review of the literature related to expertise and controversies over technological decision-making. The second section provides a review of the literature related to the ontological framework. I will close the chapter with a conclusion that brings together the ideas from each section, and charts a path for the rest of the thesis through the construction of three research questions.

2.2 Expertise and Controversies over Technological Decision-Making

In this section, I will examine how the study of controversies has been used within STS to understand: the construction of scientific facts; the construction of technological artefacts; and technological decision-making. In particular, I will describe how recent studies of controversies over technological decision-making have focused on their relationship with expertise. Studies have often shown how certain types of expertise can be marginalised during policy-making processes. This has led, in some cases, to calls for increased democratic participation in controversies of this type. More recently, others have argued that this prompts concerns over

how far participation should be widened, given that some possess more relevant expertise than others. I argue that, although the issue of widening participation outside of the expert scientific community is important, we should also examine in more detail the nature of the participation of the expert scientific community. More specifically, I argue that we should embark upon a more careful consideration of how contributory expertise is used in controversies over technological decision-making.

2.2.1 What is a Controversy?

Controversies are the visible result of disagreements between actors. STS is concerned with those disagreements that feature science and technology. Built into the fabric of a controversy is the idea that something becomes controversial when one group (or an individual) disagrees with the ‘actions of’ or the ‘claims made by’ another. Studies of controversies have therefore sought to identify divergent groups, sides, stakeholders, or individuals. This idea is so fundamental to the study of controversies within STS that it is difficult to conceive of one that does not conform to this basic rule. In approaching controversies from a philosophical point of view, H. Tristram Engelhardt Jr and Arthur Caplan observed that:

Scientific disputes resist closure or resolution when the stakeholders in the debate belong to: 1) Different scientific communities with different appreciations of the evidence at stake . . . or 2) Competing social groups with different views of social control (Engelhardt & Caplan 1987, p.11).

The idea that competing social groups may hold different views about social control is, to a certain extent, unsurprising. However, the idea that different scientists can have different appreciations of the evidence at stake challenges the realist account of scientific knowledge production. This challenge has its roots in Thomas Kuhn’s (1962) study of scientific revolutions. Kuhn rejected the idea that scientific knowledge develops through continuous accumulation, and attempted to replace it with the idea that scientific knowledge develops through revolutions. According to Kuhn, during periods of revolution, different groups of scientists come to operate under incommensurate conceptual frameworks, resulting in interpretations of evidence that share little common ground. Kuhn located these incommensurate

frameworks within, amongst other things, social context, thus opening the door to a sociological interpretation of the creation of scientific knowledge.

When describing controversies from the point of view of STS, a common first step is to identify the groups involved. For example, in the opening sentence of Dorothy Nelkin and James M. Jasper's study of an animal rights controversy, divergent groups are used to frame the study:

In the spring of 1976 a **coalition of animal protection organizations** mobilized a protest against research on cats at the **American Museum of Natural History** in New York City [emphasis mine] (Nelkin & Jasper 1993, p.26).

Here, as in many other case studies, the controversy is immediately presented as a disagreement between groups. However, care must be taken when presenting controversies in this way. As Sheila Jasanoff (1996) pointed out, a controversy will often consist of more than just two sides, and as a result, presenting a debate as two-sided can result in an over-simplification. Also, as Bruno Latour (2005, pp.27-43) has argued, when examining controversies sociologically, the analyst should avoid imposing rigid groupings upon the actors under study. Instead, it should be recognised that over the course of a controversy, groups can be formed and dismantled, and their membership can change. Latour advised that, rather than adopting ideas about fixed groupings, the analyst should "follow the actors" and begin by studying "the traces left behind by their activity of forming and dismantling groups" (Latour 2005, p.29).

2.2.2 Controversies and the Sociology of Scientific Knowledge

The literature on scientific and technological controversies within STS can be broadly divided into three strands: controversies over the results of scientific experiments; controversies over the development of technological artefacts; and controversies over technological decision-making.¹ Scientific and technological controversies were first studied in order to generate theories about how scientific knowledge is produced (e.g. Bloor 1976, Collins 1985). As such, they became central to a subfield of STS known as the Sociology of Scientific Knowledge (SSK). Following Kuhn's (1962) study of the structure of scientific revolutions, the intuitive idea of scientific knowledge as cumulative was called into question. Nonetheless, science was also thought of as consistently able to establish facts. This created a dilemma for those wishing to understand how scientific knowledge is created. Harry Collins was among the first to express this problem, when he wrote:

To speak figuratively, it is as though epistemologists are concerned with the characteristics of ships (knowledge) in bottles (validity) while living in a world where all ships are already in bottles with the glue dried and the strings cut. A ship within a bottle is a natural object in this world, and because there is no way to reverse the process, it is not easy to accept that the ship was ever just a bundle of sticks (Collins 1975, p.205).

Collins believed that the study of controversies over scientific knowledge offered a partial solution to this problem, given that they described the processes that occurred during the creation of facts. As Collins went on to explain:

¹Though often clearly noticeable, the separation between these strands is sometimes implicit rather than explicit. This is because, when considering complex case studies, it is almost inevitable that when a case study is framed broadly, it will overlap into more than one of the above categories. For example, although Harry Collins and Trevor Pinch's (1993) study of attempts to achieve cold fusion primarily detailed a controversy over the results of scientific experiments, the way that that these experiments were viewed by policymakers began to influence future research funding strategies. In this example, claims of success led directly to \$5 million being devoted to cold fusion research by the Utah State Legislature. Furthermore, had the experiments ultimately been deemed successful, debates over the appropriate use of cold fusion in society would undoubtedly have followed. Therefore, although attempts have been made to produce more specific controversy typologies (e.g. Giere 1987, Nelkin 1993*a*), they often offer little analytical purchase due to significant category overlap. Furthermore, forcing a controversy into a particular category may unnecessarily exclude theoretical perspectives and explanatory factors.

My contention is that because of the institutionalization of the production of scientific truth, it is possible to make a partial escape from the cultural determinism of current knowledge, in studies of science. It is actually possible to locate this process in scientific laboratories, in letters, conferences and conversations. It is possible then to perform a kind of automatic phenomenological bracketing for ideas and facts, by looking at them while they are being formed, before they have become 'set' as part of anyone's natural (scientific) world (Collins 1975, pp.205-206).

On this understanding, controversy 'closure' is the culmination of the processes used to produce scientific facts. Collins argued that the results of experiments alone are not always sufficient to establish facts. In particular, Collins asked, if the purpose of an experiment is to establish the existence of a particular phenomenon, how can one be sure what the right result for that experiment should be, and how can one be sure that the experiment has been performed correctly? Collins coined the term "experimenters' regress" to refer to the cycle that can result from believing that a good experimental result will emerge from properly functioning experimental apparatus, whilst simultaneously believing that it is possible to tell when the experiment is functioning properly because it has delivered a good result. Collins argued that the experimenters' regress can eventually be (and often is) broken, and controversies over experimental results can be settled, through recourse to social factors, such as the status, credibility, and persuasiveness of an experimenter (Collins 1992). A number of studies of the creation of scientific knowledge have been since been conducted (e.g. Latour & Woolgar 1979, Knorr Cetina 1981, Lynch 1985, Latour 1987), not least Collins' own 35-year project examining the search for gravitational waves (Collins 2004).

2.2.3 Controversies and the Sociology of Technology

Controversies have also been used as a way to understand the development of technological artefacts. In the late 1980s, following the successful use of controversies to illuminate the creation of scientific knowledge, many technologies came to be seen as the result of controversies over the most appropriate solution to a particular problem, rather than as a result of 'technological determinism' - the view

that technology emerges independent of societal influence, but shapes the development of society thereafter.² In the absence of an independent and context-free measure of appropriateness, technological developments came to be seen, in part, as the result of certain groups making more effective use of rhetoric, or furthering their interests in some other way. Taking as their example the development of the bicycle, Trevor Pinch and Wiebe Bijker (1989) argued that the convergence towards pneumatic tyres and equal sized wheels could be understood in terms of the competing interests of certain groups, given the ‘interpretive flexibility’ of the meanings and purposes of cycling.

Interpretive flexibility, an idea that was developed under a branch of SSK in order to understand disagreement over the meaning of the results of scientific experiments, came to be the central idea in a separate strand of STS that aimed to explore the Social Construction of Technology (SCOT). A number of studies of the development of technological artefacts have been conducted using the SCOT approach, including Donald MacKenzie’s (1990) study of nuclear missiles and Boelie Elzen’s (1986) study of ultracentrifuges. However, the notion of interpretive flexibility raised further questions. For example, despite the interpretive flexibility inherent in many systems, controversial elements (such as the air-filled tyre) appeared to have been ‘bypassed’, given that the same technologies were used by groups with varying interpretations of cycling. This phenomenon was accounted for using the concept of ‘boundary objects’. Using an historical study of Berkeley’s Museum of Vertebrate Biology, Susan Leigh Star and James Griesemer (1989) showed how the creation of boundary objects, in conjunction with a process of methods standardization, allowed groups within the museum with divergent viewpoints and agendas - such as collectors, trappers, and administrators - to co-operate in such a way that the museum remained functional in line with each group’s definition. Star and Griesemer defined boundary objects as “objects which are both plastic enough to adapt to local needs and the constraints of the several parties employing them, yet robust enough to maintain a common identity across sites” (Star & Griesemer 1989, p.393). Furthermore:

²The question of whether technology drives society has a long history. As David Edgerton (1993) has pointed out, some scholars had rejected technological determinism outright before the development of an identifiable field focussed on the sociology of technology (e.g. Noble 1977, 1984). ‘Soft’ forms of technological determinism can be found in the work of Langdon Winner (1977, 1980) and Thomas P. Hughes (1994). Furthermore, ‘hard’ technological determinism can still be detected in recent writings on technology outside of STS (e.g. Kelly 2011).

[Boundary objects] have different meanings in different social worlds but their structure is common enough to more than one world to make them recognizable, a means of translation. The creation and management of boundary objects is a key process in developing and maintaining coherence across intersecting social worlds (Star & Griesemer 1989, p.393).

In reviewing the literature on the subject, Joan Fujimura (1992) found that the boundary objects concept has since been widely used throughout STS. For example, Kathryn Henderson (1991) has shown how sketches and technical drawings acted as boundary objects that allowed scientists and engineers involved in different stages of the innovation process to work together on the development of turbine engines, and Jenny Marie (2008) has shown how certain breeds of rabbit and poultry acted as boundary objects that connected fanciers, and scientists requiring animals for genetic research.

2.2.4 Ways of Understanding Controversies

Whether using controversies to understand the creation of scientific facts, or the development of technological artefacts, STS scholars have used social factors to account for the way they unfolded. Numerous strategies have been developed for understanding the creation of scientific knowledge and technological artefacts, including: the Strong Programme; the Empirical Program of Relativism (EPOR); discourse analysis; and Actor-Network Theory (ANT). Although not all of these approaches are compatible with one another, most are based on the idea that the study of controversies can reveal something about the role of social factors in the processes used to create scientific knowledge and technological artefacts.

Studies guided by the Strong Programme and EPOR have tended to emphasise the role played by 'interests'. The basic idea is that disagreements over scientific knowledge, or the development of a technology, can be understood in terms of the different social interests of the actors involved. Interests may be classed as external or internal to science and technology, but the explanatory framework remains broadly the same. The Strong Programme, which was originally stated by David Bloor (1976), aimed to uphold four tenets in the course of offering sociological explanations: causation; impartiality; symmetry; and reflexivity. The causation

tenet stressed the study of the conditions that bring about knowledge. The impartiality tenet stressed that studies would be impartial with respect to the truth or falsity of the creation. The symmetry tenet stressed that the same kinds of explanation would be applied to true and false claims. The reflexivity tenet stressed that tenets would also apply to the claims made under the Strong Programme itself. Studies carried out using the Strong Programme have tended to emphasise the role of external interests, in that they have looked to the macro-social or the historical interests of the actors involved to explain narrower debates within science. For example, Steven Shapin's (1975) study of the early nineteenth century debates that occurred in the field of phrenology located scientific disagreements within the wider social debates taking place in Edinburgh at the time, rather than in the incommensurability of the various intellectual positions. Shapin argued that, given that the Edinburgh Phrenological Society drew its members from the lower and middle classes, and that the Royal Society of Edinburgh drew its members from the upper classes, disagreements over the science could be understood in terms of a class struggle.

EPOR has tended to emphasise the role of more localized, internal interests. EPOR, which was most clearly stated by Collins (1981*b*), consists of three sequential stages: demonstrating the interpretive flexibility of experimental data; showing, despite this, the mechanisms by which closure is reached; and finally, linking these mechanisms to wider social and political structures. Though broadly similar to the Strong Programme, satisfactorily completing the third stage of EPOR has sometimes proved difficult, and descriptions produced using EPOR have not typically drawn on macro-social or historical factors in order to explain disagreement, but have instead located the disagreement within the social aspects of the scientific community. Studies have shown how the participants in a controversy may disagree because they are using, for example, different conceptual frameworks, or different methods of knowledge production. For example, Andrew Pickering's (1981) study of a controversy over the observation of a magnetic monopole - a hypothetical elementary particle - showed that interpretive flexibility was dealt with through recourse to an uncontroversial prior theoretical consensus, rather than through the results of experiments alone.

A number of studies of scientific and technological controversies have been based on an examination of 'rhetoric'. In line with the so-called 'linguistic turn', scholars

have primarily analysed controversies as discourses or texts. On this understanding, the study of changes in rhetoric can not only be used to mark out different stages of a controversy, but the style of the contrasting arguments of divergent groups can be seen as actually constituting, shaping, and closing the controversy. Whilst some EPOR explanations have accounted for controversy closure in this way (Collins 1981a), rhetoric was typically the main focus of studies conducted using discourse analysis and some early studies based on ANT.

The study of rhetoric began with the Ancient Greeks, and contemporary ideas about rhetoric are the intellectual descendants of these concepts (Conley 1990). Following on from Kuhn (1962), historians, sociologists, and philosophers of science and technology have been able to examine the content of argument and form a link between rhetoric and controversy. On this understanding, whilst it may be tempting to conceive of scientific argument as antithetical to rhetorical argument, scientific argument is ultimately designed to persuade, and as a result, the choices made by the proponent of a scientific or technical argument can be analysed. For example, Charles Bazerman (1988) has shown how Isaac Newton took a collection of experiments and observations and recast them as one single experiment when attempting to communicate his ideas in his *New Theory of Light and Colours*. Bazerman argued that, in recasting multiple experiments as one experiment, Newton used a rhetorical device in an attempt to convince his readers of his argument. G. Nigel Gilbert and Michael Mulkey's (1984) study of discourses within biochemistry showed, amongst other things, that scientists are able to draw on both an 'empiricist repertoire' - when justifying the formal experimental procedures associated with their own work - and a 'contingent repertoire' - when describing the social or psychological factors that explain why other scientists disagree with them. Similarly, in their famous historical study of the debates between Robert Boyle and Thomas Hobbes, Steven Shapin and Simon Schaffer (1985) showed that the rhetorical strategies employed by both individuals mirrored their respective views on the scientific method and the impact this had on the idea of certainty.

Before delving any further into how rhetoric has been used to understand controversies, it is useful to examine what it means for something to be considered uncontroversial. In carrying out an anthropological study of the Salk Institute, Bruno Latour and Steve Woolgar (1979) defined a completely uncontroversial statement as a 'fact'. They argued that a fact is a piece of information that has shed both its modalities and the history of its creation. Given that science aims to produce

facts, they argued, rhetoric is used to remove both the modalities and history from arguments, so that the acceptance of a statement is universal. These ideas were developed further in Latour's *Science in Action* (1987). Underpinning this work, and many other studies of controversies, is the idea that in order to understand how scientific knowledge is produced, it must be examined whilst facts are in the process of being created, instead of when they arrive fully-formed - a state Latour referred to as 'black boxed'. Latour attempted to trace the development of a set of controversies by studying the published claims and counterclaims of scientists. In one of his examples, Latour took a debate on the structure of hormones between two Nobel Prize-winning endocrinologists - Roger Guillemin and Andrew Schally. Latour showed how both Guillemin and Schally attempted to add authority to their claims, and simultaneously subtract authority from their opponent's claims, by drawing on the status of the investigator and the context of the citation.

The Latourian approach to understanding controversies through rhetoric, and the work on ANT that this fed into, has not been universally praised or accepted (e.g. Amsterdamska 1990). One objection that the focus on rhetoric raises is that it appears to downplay the importance of the tangible or the material - whether in the form of objects or evidence. Furthermore, if a controversy is conceived of in terms of discourses, texts, claims and counter-claims, this may distort our understanding of other important factors. The emphasis on closure that's noticeable in the literature on controversies may be linked to the fact that many studies of controversies have taken rhetoric as the focus. If rhetoric is used as the lens through which controversies are examined, aside from potentially marking out the start of a controversy, it is difficult to see how it would be possible to use the study of claims and counter-claims to analyse events prior to the disagreement. This serves to isolate the controversy from the uncontroversial scientific research that preceded it, and goes some way to prohibiting the study of how something can become controversial in the first place.

It is also questionable whether the analysis of rhetoric can capture how actors use certain aspects of non-communication, such as secrecy, silences, and absences, strategically (or otherwise) during controversies. Such features have clear relevance for the study of controversies. However, they also carry with them equally clear barriers to their descriptive representation. This may explain why they have rarely been studied in the past. However, as Brian Rappert (2010) has argued, there is now an emerging belief that such barriers are not necessarily insurmountable,

but may instead require novel approaches to research. Brian Balmer (2012) has used a number of case studies of research into biological and chemical weapons to highlight how secrecy can alter the dynamics of knowledge production, and thus the way scientists communicate. Balmer has shown that science produced in secret can exclude scientists from their traditional reward system, and can insulate scientists from the moral objections to their work, resulting in a lack of rhetorical engagement. Even when claims and counter-claims relating to secret science are exchanged, and it is therefore possible to analyse the rhetoric used, acknowledging the role of secrecy should surely be central to their interpretation. For instance, Balmer argued that one way in which outcomes of secret science can be observed is when secrets are revealed, either intentionally - through publication or press release, or unintentionally - through leaks. From the point of view of those carrying out secret research, when secrets are revealed intentionally, the process can be carefully managed, and decisions about what to release and how to release it are paramount. When secrets are revealed unintentionally, they can initiate a process of information management that aims to deal with the new circumstances. Furthermore, secrecy can also actively construct uncertainty, gossip and rumour that can then be used either positively or negatively by actors on all sides of the secrecy divide.

Latour's early work on the use of rhetoric during scientific controversies fed into the development of ANT. ANT - which is principally associated with Latour, Michel Callon and John Law - was developed during the 1980s, and came to dominate STS throughout the 1990s and beyond. Though Latour's work on the use of rhetoric in controversies shares similar goals with the Strong Programme, EPOR and discourse analysis, ANT departs from them in that it aims to describe a much broader set of relationships. ANT can be described as a material-semiotic method, because it aims to describe how material objects and concepts can come together to form networks. Nodes within ANT are referred to as actors, and the actor label can be applied to almost anything. Furthermore, things that we may more readily think of as networked, can themselves be actors within broader networks. Importantly, actors, whether human or non-human, are granted agency during the network building process. Where ANT has been used to understand controversies, success and failure has been accounted for in terms of how well actors were able to build and maintain networks. For example, in Michel Callon's (1986) study of the controversy over the decline of scallops in France, the survival of scallop stocks at St Brieuc Bay was accounted for through the success of actors - including the

scallops themselves - in building a stable network. Conversely, in Latour's (1996) study of the Aramis project - an attempt to implement a rapid transit system in Paris - it was argued that the project's failure was due to the crumbling of the network of which it was a part. In contrast with previous approaches described, using controversies to understand the creation of scientific knowledge and technological artefacts is not the central purpose of ANT. ANT aims to describe a much broader set of social phenomena. Because of this, though it is possible to say much more about the concepts associated with ANT and the work it has spawned, it does not make sense to describe ANT in any further detail during a discussion of controversies.

2.2.5 Controversies over Technological Decision-Making

In parallel to work aiming to understand the development of scientific knowledge and technological artefacts, STS scholars have also studied controversies over technological decision-making. Controversies over technological decision-making have been studied because they are interesting and important in their own right, but also because they can reveal something about prevailing attitudes towards science and technology. Building upon an earlier definition from Collins and Evans (2002), Martin Weinel (2010) described controversies over technological decision-making as "those points where science and technology intersect with the political domain because the issues are of visible relevance to the public". As Weinel (2010, pp.19-20) pointed out, numerous attempts have been made by others to define similar scenarios. For example, Brooks (1964) defined 'science for policy' as "matters that are basically political or administrative but are significantly dependent upon technical factors", and Brian Wynne (2007) has referred to 'public decision-making that involves science'. These are essentially nothing more than different ways of referring to the same phenomena. Given that this thesis will place an emphasis on third wave ideas, I have adopted their use of 'controversies over technological decision-making'.

Dorothy Nelkin (1993*a*) has provided a brief history of studies of controversies over technological decision-making in the second half of the twentieth century. Nelkin argued that, in the years that followed the Second World War, science and technology was seen as the engine for a sustained period of economic growth.

However, this period of growth also spawned an increased awareness of risk. Nelkin described how:

Technological improvements were threatening neighbourhoods and causing environmental problems; drugs to stimulate the growth of beef cattle were causing cancer; efficient industrial processes were threatening worker health (Nelkin 1993*a*, p.x).

By the late 1970s, both scientists and non-scientists were seriously examining the possibility that certain kinds of scientific research should not be done at all (e.g. Morison 1978). This resulted in the perception that science had entered a period of crisis. Reflecting upon this in 1990, Yaron Ezrahi described the recent ‘attacks’ on science and scientific research as a major conceptual shift that marked “its visible decline as a force in the rhetoric of liberal democratic politics” (Ezrahi 1990, p.13). This observation chimed with Ulrich Beck’s (1992) claim that contemporary Western societies during this period came to be thought of as ‘risk societies’. As such, Beck argued that the central problem facing risk societies was not the production of social ‘goods’ - such as wealth and employment, but the minimization of the effect of risks - which Beck called social ‘bads’.

As with debates over the results of scientific experiments and technological artefacts, the examination of interests and rhetoric has been a key part of understanding controversies over technological decision-making. To take just one example, Michael Mulkey’s (1997) study of the embryo research debate in the UK focused on how the rhetorical strategies of scientists, politicians, and others, shaped the debate over the extent to which embryo experimentation should be allowed. Mulkey devoted much of his case study to the examination of parliamentary debates, and as such, showed that the focus on claims and counter-claims of actors can form the basis of a study of this type.

There is probably a link between the research into the social factors that influence the creation of scientific knowledge and the broader trends that Nelkin, Ezrahi and Beck identified. For some, the belief that scientific and technological controversies are amenable to social factors prompted a rethink of the notion that authority lies solely in scientific and technological expertise during controversies over technological decision-making. Brian Wynne’s (1992, 1996) study of Cumbrian sheep farming following the fallout from the Chernobyl disaster described

how scientific expertise came into conflict with non-scientific expertise of the local farmers. In attempting to predict the risks that might result from the fallout, scientists adopted a one-size-fits-all approach based on laboratory evidence that told them that the threat posed by the fallout would clear up relatively quickly. This was at odds with the farmer's own knowledge of the geographical features of the land that told them that the threat was likely to linger. This conflict undermined confidence in the ability of the scientists and policymakers to deal with the problem, and suggested a role for other forms of non-accredited expertise in scientific and technological decision making. Other studies have emphasised that, during controversies over technological decision-making, non-scientific or non-accredited experts can make a positive contribution. For example, Steven Epstein (1995) has described how AIDS victims were able to contribute to the design of clinical trials, and Alan Irwin (1995) has shown how farmers were able to contribute to an assessment of the risk of certain pesticides. Together, these studies suggested that knowledge produced in the laboratory is not necessarily sufficient to meet the challenges of controversies over technological decision-making, and that locating expertise in scientific credentials can exclude people with relevant expertise, to the detriment of eventual outcome.

This can be thought of as a 'conservative critique' of the use of credentialed scientific expertise during controversies over technological decision-making (Evans & Collins 2008). The conservative critique, together with the sociological interpretation of the construction of scientific knowledge and technological artefacts, has been used as a platform for a more radical critique. Here, case studies showing that the successful use of rhetoric or the promotion of interests is all that separates true scientific facts from false claims, together with a broader commitment to interpreting the two symmetrically, serves to strip scientific evidence of epistemic authority. Under this view, prioritising scientific expertise over the economic, political, and moral preferences of the public can no longer be justified. Furthermore, it is argued that the inclusion of scientific expertise results in the framing of the controversy as something that primarily hinges on scientific and technological questions, at the expense of other factors (Wynne 2003). For example:

Those opposed to further developments in genetic testing and screening may question their economic, political, and moral consequences by stressing the way in which they reinforce existing inequalities (e.g., allowing the affluent or powerful to enhance their children's genetic

inheritance); create new forms of discrimination (e.g., a return of eugenics via the “deselection” of embryos seen as likely to have a disease or disability); and/or presume the desirability of increased industrialization, commodification, and control (e.g., by implying that it is proper to choose or design humans) (Evans & Collins 2008, p.612).

The implementation of deliberative or participatory processes - where the opinions of stakeholders and the wider public are solicited - have been advocated as ways of counterbalancing overly scientific or technological framings (Irwin 1995, Burningham 1998, Wilsdon & Willis 2004). Whilst the implementation of such processes may improve the acceptance of decisions, it has also helped promote the notion of ‘lay expertise’ - the expertise acquired by virtue of being a citizen, and the expertise required to arrive at economic, political, and moral outcome preferences. This has served to give credence to the idea that, during controversies over technological decision-making, there should be no boundaries to participation because no special weight should be given to those with scientific or technological expertise, because lay expertise is all that’s required to ensure that public preferences are considered. On this understanding, during a controversy over technological decision-making, “the proper participants are in principle every democratic citizen” (Wynne 2003, p.411). This dissolution of the boundaries between experts and citizens during controversies over technological decision-making chimes with other descriptive studies of how scientific and technological expertise is attributed. Sheila Jasanoff (2005) used a study of the controversies over biotechnology in the UK, Germany, the US, and the EU to show how different countries use different mechanisms and modes of reasoning to make decisions about what constitutes expertise. These mechanisms, Jasanoff argued, form part of ‘civic epistemologies’ that are deeply embedded in the institutions that manage, shape and frame political issues. Given that many of these institutions are made up of people without scientific and technological expertise, citizens and lay people become active participants in the construction of expertise, thus undermining the basis for a clear demarcation between the two.

2.2.6 The Third Wave of Science Studies

In response to these arguments, Harry Collins, Robert Evans, and others, have established a Studies of Expertise and Experience (SEE) research program. In a

much-discussed article, Collins and Evans (2002) argued that given work on the ‘problem of legitimacy’ had shown that a case could be made for consulting those outside of the scientific and technical elite during controversies, in light of recent attempts to increase democratic participation in controversies over technological decision-making, the most pressing problem was now the ‘problem of extension’. Collins and Evans asked:

Should the political legitimacy of technical decisions in the public domain be maximized by referring them to the widest democratic processes, or should such decisions be based on the best expert advice? The first choice risks technological paralysis: the second invites popular opposition (Collins & Evans 2002, pp.235-236).

Collins and Evans (2002) believed that the way to solve the problem of extension was to recognise a preference for “those who know what they’re talking about”. To realise this, they claimed that it was necessary to reconceptualize the notion of what it means to be considered an expert. They argued that developing new categories of expertise, and then taking a normative position on the extent to which they should be drawn upon should constitute a third wave of science studies.

Collins and Evans (2002) argued that STS could be divided into three ‘waves’.³ During what they called the first wave of science studies, which occurred during the 1950s and 1960s, scholarly work served to reinforce the successes of science. During controversies over technological decision-making, it was acceptable for scientists to profess on matters outside of their field, and it was virtually unthinkable for those outside of the scientific and political communities to influence decision-making processes. During what they called the second wave of science studies, scholarly work typically established and maintained that science was socially constructed, extra-scientific factors were often drawn upon to close controversies, and following on from the work of Wynne, Jasanoff, and others, the lines between experts and non-experts were blurred. This, as has been described, led some to conclude that experts should have no more influence on controversies over technological decision-making than non-experts. Therefore, with respect to the problem of legitimacy and the problem of extension:

³Collins and Evans used the term ‘waves’ to refer to the dominant ways of thinking within a particular period, rather than to specify clear boundaries beyond which certain ways of thinking could not be found. As such, Wave One attitudes, though dominant in the 1950s and 1960s, did not disappear with the advent of Wave Two, and can still be observed today within certain discourses (e.g. Henderson 2011).

The First Wave of Science Studies had no Problem of Extension, and was unaware of the Problem of Legitimacy. The Second Wave of Science Studies was good for solving the Problem of Legitimacy that it inherited from Wave One, but replaced it with the Problem of Extension . . . We propose that the Third Wave of Science Studies should accept the Second Wave's solution to the Problem of Legitimacy, but still draw a boundary around the body of 'technically-qualified-by-experience' contributors to technical decision-making (Collins & Evans 2002, p.238).

To complement this idea, and to prepare the ground for future normative claims, Collins and Evans initially defined four types of expertise: no expertise; interactional expertise; referred expertise; and contributory expertise (Collins & Evans 2002, p.254). Rather than being based on qualifications or accreditation, these types were based on experience. Taking the example of a sociological fieldworker aiming to study scientific knowledge production, no expertise was defined as the degree of expertise with which the fieldworker sets out, and is therefore insufficient for both sociological analysis and making a scientific contribution to the field under study. Interactional expertise was defined as the level of expertise sufficient to perform a sociological analysis of the field, and to interact interestingly with those performing the scientific activity. Referred expertise was defined as expertise from an adjacent field that could be applied to the field under study. Finally, contributory expertise was defined as having the level of expertise necessary to make a direct contribution to the scientific field (Collins & Evans 2002, pp.254-259).

In the initial critical response to these initial claims, Wynne (2003) argued that Collins and Evans (2002) had misunderstood the problem of legitimacy, in that they based it on the belief that people with authentic but unrecognized expertise were denied access to deliberations. Wynne claimed that, in fact, the real problem was that meaning was imposed on the public in the form of decision-based, propositional questions, such as "is nuclear power safe?" or "is British beef safe?". Jasanoff (2003) argued that Collins and Evans (2002) provided a misleading characterization of the science studies literature, displayed misconceptions about the foundations of expertise in the public domain, and misunderstood the purposes of public participation in contemporary democratic societies. Pertaining to all three criticisms, Jasanoff argued that:

Expertise is not merely something that is in the heads and hands of skilled persons, constituted through their deep familiarity with the problem in question, but rather that it is something acquired, and deployed, within particular historical, political, and cultural contexts. Expertise relevant to public decisions, I have further shown, responds to specific institutional imperatives that vary within and between nation states. Accordingly, who counts as an expert (and what counts as expertise) in UK environmental or public health controversies may not necessarily be who (or what) would count for the same purpose in Germany or India or the USA (Jasanoff 2003, p.393).

Jasanoff, along with Arie Rip (2003), questioned the appropriateness of categories of expertise, given that the central question is often “what counts as relevant knowledge?”, rather than “who possess the scientific knowledge?”. Furthermore, they argued, even if contributory expertise can be clearly identified as such, we may still wonder about the circumstances under which this came about:

If we regard the very formation of expert ‘core-sets’ as a political phenomenon, then attention inevitably has to focus on the processes by which such sets are created, maintained, patrolled, and protected. In many areas of public policy, we may not be interested in re-examining the foundations of settled expertise in this way, but when controversy erupts, it becomes important to ask what sustains the authority of a particular group of experts and their expertise (Jasanoff 2003, p.395).

In general, then, the criticisms levelled at the third wave suggested - in the minds of the critics at least - that it had failed to sufficiently incorporate ideas from the second wave, and could therefore be interpreted as a backwards step towards first wave ideas. In response to Wynne’s criticism, Collins and Evans (2003) replied that, although the overly scientific or technological framing of controversies over technological decision-making should be avoided, this should not lead to the outright exclusion of all scientific or technological questions, and that propositional questions of the type Wynne described require expert debate. In response to Jasanoff and Rip, Collins and Evans (2003) located the root of their disagreement in a failure to appreciate the difference between descriptive aims of the second wave and the prescriptive aims of the third. Therefore, although Collins and

Evans agreed that the current processes involved in expert attribution do serve to blur the lines between expert and non-expert, this alone should not prevent the third wave from attempting to prescribe a system where this is not the case.

2.2.7 The Periodic Table of Expertise

The differences between the descriptive aims of the second wave, and the normative aims of the third wave, served to leave them somewhat separate in the literature that followed. Work continues on both in parallel. Therefore, the remainder of this section will focus on how the third wave has been developed. Since the publication of their original third wave paper, Collins, Evans and others have refined their categories of expertise and developed complementary concepts. Collins and Evans' third wave ideas were updated and extended with their periodic table of expertise. The periodic table of expertise is now used as the starting point for studies of expertise in the third wave, and remains its most important set of concepts and categories.

Before describing the 'elements' of the table, it is important to recognise that it is built upon the concept of 'tacit knowledge'. A term made famous by Michael Polanyi in his *Personal Knowledge* (1958), tacit knowledge refers to that which we know how to do, but not how to explain how to do. For example, we may possess the tacit knowledge required to ride a bicycle, but we may not be able to express this knowledge logically or explicitly. Polanyi, who'd trained as a chemist, argued that the philosophical descriptions of laboratory science that were available at the time were too reliant on explicit knowledge, and therefore neglected the tacit knowledge often required to make experiments work. Since Polanyi popularized the term, different disciplines have developed their own understanding of tacit knowledge, and in some cases, have refined and updated his ideas. Drawing on work from artificial intelligence and automation, Collins (2010) has made a useful distinction between three types of tacit knowledge: 'relational' tacit knowledge; 'somatic-limit' tacit knowledge; and 'collective' tacit knowledge. Relational tacit knowledge, which may be thought of as the 'weakest' form of tacit knowledge, simply refers to knowledge that could be made explicit, but has not, due to the contingencies of society. However, somatic-limit and collective tacit knowledge refer to knowledge that is thought of as tacit due the nature of the knowledge

itself. As Collins explained in an earlier article (written before the development of the idea of relational tacit knowledge):

Why is such a large component of human knowledge known tacitly? Two different reasons are not distinguished in the literature. The first reason is to do with the limited capacities and particular nature of the human brain and body; this gives us what I'll call 'somatic-limit tacit knowledge'. The second reason is to do with the relationship between individual humans and society; this gives us 'collective tacit knowledge'. The two kinds of tacit knowledge are rarely distinguished, because they are experienced and acquired by humans in the same way: through immersion in society and guided practice. Nevertheless, they have not only entirely different causes, but entirely different consequences (Collins 2007*a*, p.258).

This can be illustrated using the bicycle-riding example. Although it may appear that no formal rules can exist for bicycle balancing, this can be seen as the result of the limitations of human beings, rather than the nature of the knowledge itself. As such, it is somatic-limit tacit knowledge. If we were to imagine trying to balance a bicycle on the moon (or any other environment with low gravity), or if the speed of human reactions and comprehension were drastically increased, it is conceivable that the rider could follow a set of logical step-by-step instructions to balance the bicycle. After all, scientists have already built robots that can balance bicycles in controlled environments. However, balancing a bicycle is not the same as riding a bicycle. Bicycle riding, particularly in urban environments, requires collective tacit knowledge. It is collective tacit knowledge that allows for, say, an appreciation of the social conventions that govern traffic, thus allowing the rider to negotiate safe passage. Collective tacit knowledge recognises that changing the context will change the nature of the activity. Riding a bicycle in London is not the same as riding a bicycle in Amsterdam. Crucially:

Collective tacit knowledge is not a matter of the accident of the human constitution, but a matter of the knowledge itself. This knowledge has to be known tacitly, because it is located in human collectivities and, therefore, can never be the property of any one individual. The simplest way to see this is to note that changes in the content of the

UBIQUITOUS EXPERTISES					
DISPOSITIONS				Interactive Ability Reflective Ability	
SPECIALIST	UBIQUITOUS TACIT KNOWLEDGE			SPECIALIST TACIT KNOWLEDGE	
EXPERTISES	Beer-mat Knowledge	Popular Understanding	Primary Source Knowledge	Interational Expertise	Contributory Expertise
				Polimorphic Mimeomorphic	
META-	EXTERNAL (Transmuted expertises)		INTERNAL (Non-transmuted expertises)		
EXPERTISES	Ubiquitous Discrimination	Local Discrimination	Technical Connoisseurship	Downward Discrimination	Referred Expertise
META-CRITERIA	Credentials		Experience		Track-Record

FIGURE 2.1: The Periodic Table of Expertise (Collins & Evans 2007)

knowledge belonging to communities is beyond the control of the individuals within the communities. For example, the content of the ever-changing argot that children speak, to the irritation of their parents, is not under the control of any child or parent; it evolves at the collective level and no one knows the rules of its evolution (Collins 2007a, p.260).

As such, if collective tacit knowledge is seen as being located within the society that produced it, and varies accordingly, it is of particular relevance to sociologists.

The periodic table of expertises (see Figure 2.1) is built upon these ideas about tacit knowledge. The first row of the table - ubiquitous expertises - refers to the expertise, such as natural language speaking, that “every member of society must possess in order to live in it”. As a result, if one possess ubiquitous expertises, one also has a huge body of tacit knowledge. The second row of the table - dispositions - refers to expertises such as “linguistic fluency or analytical flair”. The third row - specialist expertises - deals with different types of knowledge. It is divided into ubiquitous tacit knowledge and specialist tacit knowledge. Expertise that is based on ubiquitous tacit knowledge includes: beer mat knowledge - isolated facts that serve little purpose outside of general knowledge quizzes; popular understanding of science - knowledge obtained from popular books and journalism that typically eschew discussions of doubt and uncertainty; and primary source knowledge - knowledge garnered from journal articles and other first-hand documents

that provide a shallow description of how an activity operates, and say very little about the context in which it was created. Expertise that is based on specialist tacit knowledge includes interactional expertise and contributory expertise, the definitions of which were largely carried over from Collins and Evans' (2002) original third wave article. The fourth row - meta-expertises - is concerned with the expertise required to assess the expertise of others, and is divided into transmuted expertises and non-transmuted expertises. Transmuted expertise refers to the ability of those who don't possess a particular specialist expertise, but do have the expertise required to judge between those that do. Given that these judgements are based on things like demeanour, consistency, and trustworthiness, they use social distinctions to produce technical distinctions. Non-transmuted expertises refers to the judgements made based on a level of expertise related to the expertise being judged. As such, it includes the expertise of, say, art critics, as they may not create art themselves, but use their expertise to pass judgement, and managers of scientific projects, who can use their management skills on a range other projects if required. The fifth row - meta-criteria - refers to the criteria that outsiders use to judge between experts if they have no expertise themselves. As such, it includes judgements based on information about the expert's qualifications or track record (Collins & Evans 2007).

2.2.8 Specialist Expertise

When analysing controversies over technological decision-making, consideration of the expertise contained in all five rows of the periodic table may be required. However, the types of expertise under the heading of specialist tacit knowledge are those that have received the most scholarly attention. Contributory expertise was defined as "what you need to do an activity with competence", and interactional expertise was defined as "the ability to master the language of a specialist domain in the absence of practical competence" (Collins & Evans 2007, p.14). Much of Collins and Evans' (2007) was devoted to a discussion of interactional expertise, as has been much subsequent work within the third wave (see Collins 2007*b*). This is because it was thought to be the least well-understood type. It also specifically relates to the problem of extension, given that it widens the boundary between experts and non-experts during a controversy. In particular, it licenses the STS scholar to participate in a controversy related to a scientific field that they do not

have any practical experience of contributing towards. In this sense, the idea of interactional expertise is crucial for linking the second wave to the third wave.

Collins and Evans (2007) placed language at the centre of their understanding of interactional expertise. They wrote that:

Typically, sociologists who want to study areas of scientific knowledge that are new to them have to try to grasp something of the science itself. The sociologist begins with no specialist expertise - which is a level insufficient to do sociological analysis of scientific knowledge. The sociologist is likely to move rapidly through public understanding and primary source knowledge, which are also inadequate to allow for competent social analysis of scientific knowledge. With luck, however, interactional expertise, which does allow for social analysis of scientific knowledge, will eventually be attained ... The transition to interactional expertise is accomplished, crucially, by engaging in conversation with experts. Interactional expertise is slowly gained with more and more discussion of the science (or other technical skill) (Collins & Evans 2007, pp.32-33).

Furthermore:

Where interactional expertise is being acquired, there will be a progression from "interview" to "discussion" to "conversation" as more and more of the science is understood. There is no sudden "ah hah" moment that marks the switch to mastery of interactional expertise, but its steady acquisition can nevertheless be recognized (Collins & Evans 2007, p.33).

Experiments have been carried out that confirm the existence of interactional expertise. Here, an imitation game similar to the famous Turing Test was used to show that a judge could not distinguish between colour blind and colour sighted individuals because the colour blind are constantly immersed in the language of colour. In contrast, those without perfect pitch cannot pretend to have perfect pitch because they are not typically immersed in language that refers to it (Collins et al. 2006).

In contrast, less has been said about the nature of contributory expertise within the context of the third wave. Collins and Evans (2007) placed practices at the heart of their description of contributory expertise. They, in part, understood contributory expertise in terms of Stuart and Hubert Dreyfus' (1986) five-stage model of skill acquisition. Under this model, as an individual learns more about an activity, they pass through five stages: novice; advanced beginner; competence; proficiency; and expertise. As an individual progresses towards expertise, increased practical understanding results in a reduced reliance on formal rules, thus more of their actions are based on context and instinct. This understanding of contributory expertise also relates to the problem of extension, given that it can be used to attribute specialist expertise to those without scientific qualifications or accreditation. Instead, what's required is practical experience of contributing towards a particular domain. Therefore, under the third wave, in Wynne's (1996) aforementioned case study, the Cumbrian sheep farmers would be considered experts because of their considerable experience of sheep farming in that particular locale (Collins & Evans 2002).

2.2.9 Elective Modernism

More recently, Collins, Weinel, and Evans (2010) have outlined what the third wave categories of specialist expertise imply for controversies over technological decision-making. They have labelled this 'elective modernism'. Elective modernism consists of normative principles that are designed to, if used during a controversy over technological decision-making, avoid the tendency towards 'technological populism' they argued could be seen during the second wave, whilst also avoiding the 'technocracy' of the first wave that the second wave had shown to be both untenable and undesirable. To be clear, technocracy was used here to refer to the complete exclusion of non-credentialed scientific and technological experts from a controversy, and technological populism to the complete opening up of a controversy to anyone, and with no special preference given to scientific expertise. The purpose of elective modernism, then, was to outline a system that realised a "preference for democracies which actively promote discussion and debate of technical matters yet which reject populism of all kinds while still rejecting technocracy" (Collins et al. 2010, p.185).

In order to do this, Collins, Weinel and Evans (2010) returned to an earlier distinction between what they called the ‘technical phase’ and the ‘political phase’ of a controversy. Put simply, the technical phase is the period during which the science is done, and the political phase is period during which the politics is done. Collins and Evans (2002) originally argued that the two phases were different from one another in terms of: the type of questions they addressed; the actors involved; the role of politics; and the type of values involved. Thus, it was argued that it is possible to distinguish between the technical phase and the political phase of a controversy using Ludwig Wittgenstein’s (1953) ‘family resemblance’ concept:

Science is a distinct ‘form-of-life’ distinguished by the key ‘formative intentions’ of the actors. Philosophical demarcation criteria might have failed but sociological demarcation criteria such as the difference between the values of science and the values of politics can still be robustly applied so long as they are meant to mark out activities that have a family resemblance. Family resemblances stand up even though not every single activity carried out under the description of science matches all the characteristics of the family. . . It is possible to distinguish between the unavoidable ‘intrinsic’ politics of science and the ‘extrinsic’ politics that are an explicit part of the political process. . . Given [the above] it is possible to maintain the distinction between the ‘technical phase’ of a technological decision in the public domain and the ‘political phase’. The technical phase is informed by the formative intentions associated with the scientific form-of-life, whereas the political phase is concerned with the formative intentions associated with the politics of the wider society (Collins et al. 2010, pp.187-188).⁴

Collins, Weinel, and Evans (2010) argued that in order to achieve elective modernism’s goal of avoiding technocracy, the political phase should always have priority over the technical phase. In other words, the outcome of a controversy over technological decision-making does not necessarily have to reflect the work of the technical phase. However, the political phase must at least consider as much of the

⁴Although the technical phase and the political phase can be demarcated in this way, it does not imply that those with specialist expertise necessarily operate within a technical phase based on the formative intentions of science. For example, although the sheep farmers in Wynne’s (1996) case study can be said to possess contributory expertise, and can be said to be operating within a technical phase, this technical phase was based on formative intentions that are different from those of both science and politics.

work of the technical phase as possible. Furthermore, the political phase should make no attempt to subvert or misrepresent the findings of the technical phase. As such, although the speed of politics is faster than the speed of scientific consensus formation, technological populism can be avoided by valuing the judgement of experts. In linking back to the second wave, Collins, Weinel, and Evans (2010) reiterated that, although it had been shown that experts cannot deliver completely neutral findings, and that there is no clear fact-value distinction, experts should at least try to insulate their work from the cultural or political environment.

Collins, Weinel, and Evans (2010) illustrated this using the example of the Brent Spar oil platform - a controversy over the Shell Oil Company's decision to dispose of the platform by sinking it into the North Sea. This decision was opposed by environmental groups - including Greenpeace - who argued that sinking the platform would pollute the sea, and would set a precedent for the future disposal of hazardous material. Ultimately, Shell responded to these arguments, and the platform was disposed of on land. Collins, Weinel, and Evans (2010) argued that, in this situation, it was possible to make two types of argument: the utilitarian; and the quasi-religious/populist:

The argument that the Brent Spar was primarily a symbol of a willingness to pollute, or mix the clean with the dirty, also has these two possible interpretations. First, there is the 'utilitarian symbolic argument' which states that sinking the Brent Spar would be the 'thin end of the wedge' - that is, it would set a precedent that would license many more similar actions. Thus sinking one rig would justify sinking any number of rigs and, perhaps, other items of industrial waste and this would cause long term pollution damage whether or not the Brent Spar was a potential pollutant on its own. This kind of argument was made by some of the actors involved in the Brent Spar debate. Second, there is the quasi-religious/populist symbolic argument which runs along the lines that the North Sea should not be mixed with unnatural things like oil rigs. This sentiment, though it is not always thought of as quasi-religious, is nonetheless what characterizes arguments for the preservation of the 'natural' environment in this absolute sense (Collins et al. 2010, p.190).

Under elective modernism, quasi-religious/populist arguments can play no part in the technical phase, and cannot be advanced by experts in the political phase. However, non-experts may advance them during the political phase, and if these arguments are successful over utilitarian arguments during the political phase, the elective modernist would have to accept this outcome as legitimate in order to avoid technocracy. In order to avoid technological populism, decisions made during the political phase should be done so with as much of the information from the technical phase as possible:

Political decisions should not be made without considering as much as possible of the technical knowledge which bears upon the decision. The democratic process, in leading up to the decision about whether oil rigs should be dumped, should make visible all that needs to be known about the effects of dumping and that the question of what needs to be known should be given as wide an answer as possible (Collins et al. 2010, p.191).

Elective modernism therefore places a focus on the relationship between the technical phase and the political phase. In order for relevant expertise from the technical phase to have the best chance of being transferred to the political phase, Collins, Weinel, and Evans (2010) highlighted the importance of having institutions in place that can allow for the political phase to be framed as imaginatively as possible. Though this would help ensure that no relevant expertise is ignored, it does not say anything about the nature of the transfer process, and what this might look like in its ideal form.

Robert Evans and Alexandra Plows (2007) have used a study of controversies over medical genomics to expand upon definitions of the technical and political phases, and to describe the relationships that can exist between them. In short, these relationships are circular. The activity of the technical phase shapes the activity of the political phase, and vice versa (see Figure 2.2). The political phase frames the questions that it becomes appropriate for the technical phase to address, and the technical phase informs the political phase of what's known and what's possible. However, Evans and Plows (2007) did not describe what an idealised form of the transfer between phases might look like.

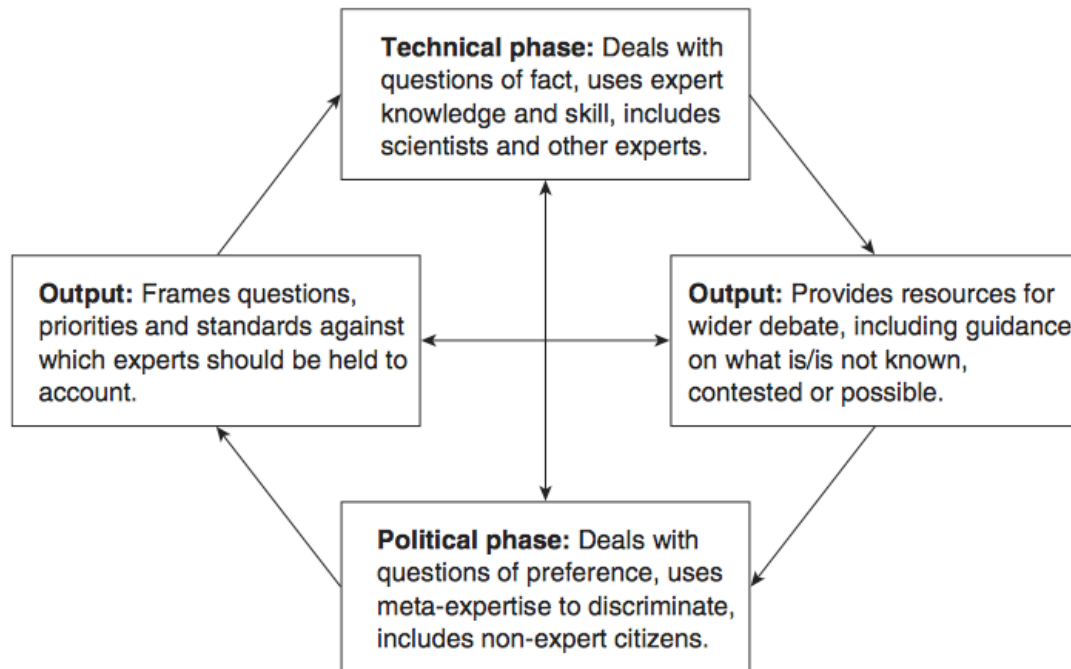


FIGURE 2.2: The Relationship Between the Technical and Political Phases (Evans & Plows 2007)

Martin Weinel (2010) has referred to the issues surrounding the appropriate relationship between the technical and the political phase as the ‘problem of integration’. According to Weinel, the problem of integration:

Arises as a direct consequence of the [third wave] model’s ability to solve the problem of extension. Implicit in demarcating between experts and non-experts is the construction of a niche or ‘phase’, as it has been called by Collins and Evans (2002), in which technical experts can make policy-relevant technical judgements. Separating such a specific ‘technical phase’ from a larger ‘political phase’ makes it logically necessary to establish rules according to which the technical advice formulated by experts feeds into or informs the political phase (Weinel 2010, p.8).

Weinel (2010) offered a partial solution to the problem of integration with the development of what he called the ‘Minimal Default Position’ (MDP). MDP is a rule that specifies a minimum requirement of the relationship between the two phases - namely, that “if experts provide a consensual answer to a propositional

question with a sufficient degree of certainty, this answer should, as a minimum requirement, have a constraining effect on the decision-making in the political phase” (Weinel 2010, pp.71-72). Importantly, the constraining effect of the technical phase is limited to the public justifications that can be used to legitimise choices in the political phase. As Weinel illustrated:

For example, imagine that experts (in the wide sense of SEE, which might include farm workers) are fairly certain that a particular pesticide is safe for users as well as for consumers of agricultural products. According to the MDP, this expert assessment itself must not necessarily impact on the policy-choices of decision-makers. It might be decided that farmers can legitimately use the substance, but it is also feasible that legislation is approved, which forbids the use of the substance. Whatever policy-choice is elected, policy-makers must not misrepresent the expert assessment. If policy-makers legislate against the use of the pesticide, they must not justify their policy-choice by claiming that experts are uncertain about the issue of safety or that experts have even judged that the substance is too unsafe. Instead, they can justify such a decision by stating that despite the fact that experts have found the pesticide to be safe, they have decided not to allow the use of the substance because they aspire to establish a purely organic agriculture (Weinel 2010, pp.72-73).

Weinel (2010) later used the MDP rule to demonstrate how the former South African president Thabo Mbeki had misused the expertise from the technical phase of a controversy over AIDS vaccination.⁵

To sum up, the principles that underpin elective modernism can be expressed in the following tenets:

1. Recognize but do not endorse religious or populist reasons in the making of technological decisions in the public domain.
2. Frame technological issues imaginatively so as to bring as many propositional questions and answers to the table as might bear on the technological decision.

⁵As will be described in the next section, Weinel (2010) also argued that although the MDP was useful for improving the style of debate, it needed to be combined with a consideration of meta-expertise in order to arrive at an outcome consistent with the goals of elective modernism.

3. Never suppress or distort the opinions of experts even if they must always be treated as subservient to politics but, on the contrary, make sure that all relevant answers to all relevant propositional questions are as visible as possible.
4. A good society will be informed by, among other things, scientific values for these are democratic values.
5. A good society will facilitate maximum scope for discussion of political matters and maximum exposure of technical matters.
6. Always aspire to keep the technical and the political phase separate even where they are combined in institutions or individuals (Collins et al. 2010, p.195).

2.2.10 Meta-Expertise

The development of elective modernism has prompted further consideration of the meta-expertise row of the periodic table. As part of the same study of AIDS vaccination in South Africa, Weinel (2010) defined two further categories of meta-expertise that could be used to better understand how expertise can be judged during controversies over technological decision-making. The judgement of expertise is an important element of the relationship between the technical and political phases, given that those in the political phase may be required to make a decision based on apparently conflicting expertise, or based on expertise of varying quality. Weinel therefore developed the additional meta-expertise categories of Domain-Specific Discrimination and Sociological Discrimination.

Domain-Specific Discrimination (DSD) refers to the non-technical expertise used by technical experts to arrive at judgements about other experts within their field. In contrast to transmuted meta-expertises, such as ubiquitous discrimination and local discrimination - where an assessment of the expertise of an individual is arrived at through consideration of social attributes to which everyone has access, or through consideration of social attributes accessible by virtue of a closer social proximity - DSD assesses expertise through social factors intrinsic to the particular field of science, such as experimental skill, reputation, background, and so on. As such, it links back to the factors identified under EPOR that are often

used by scientists to break the experimenters regress and close controversies over experimental results (Collins & Weinel 2011).

Sociological Discrimination refers to the ability of those with expertise in the sociology of science to be able to arrive at a similar judgement about a particular aspect of a controversy. Weinel (2010) argued that those with expertise in the sociology of science would be better equipped than those relying on ubiquitous discrimination to judge, for example, whether a controversy was actually present within a particular scientific field. Weinel argued that those equipped with sociological discrimination would be able to distinguish between authentic and inauthentic scientific controversies through an assessment of four criteria: explicit argument - whether a disagreement within the discipline is visible; expertise of claim maker - whether those making claims that challenges a consensus possess the relevant technical expertise; constitutive work - whether the claims that challenge a consensus were based on scientific work; and conceptual continuity with science - whether the claim that challenges a consensus has any intention of being a part of science (Weinel 2010, pp.170-171).

2.2.11 Criticisms of Elective Modernism

Critical responses to the elective modernism position have argued variously that: it rests on an untenable fact-value distinction; the technical and political phases cannot be disentangled from one another; and there's little evidence that technological populism is a problem in need of addressing (Fischer 2011, Forsyth 2011, Epstein 2011). Such criticisms are reminiscent of those levelled at the original invocation of the third wave, in that they imply that ideas from the second wave have not been sufficiently incorporated into the third. It is important to stress again that, throughout the third wave debate, critics and proponents have, to a certain extent, been arguing at crossed purposes. Proponents of the third wave, in attempting to address the problem of extension, have embarked upon a normative exercise, whereas their critics, in reiterating the arguments of the second wave, continue to be engaged in a descriptive exercise. Collins, Weinel, and Evans (2011), in noticing this tendency, have commented that there may be a certain degree of paradigm incommensurability between the two approaches. However, the similarities between wave two and wave three should not be forgotten, and

nor should the importance of wave two ideas in underpinning the third wave concepts. Therefore, much of the critical engagement with elective modernism has been disappointing due to a lack of direct engagement with its principles.

Consequently, there is plenty of scope for refining third wave ideas. Those that have done this have tended to focus on how the relationship between the technical and political phases should function (Evans & Plows 2007, Weinel 2010). Under elective modernism, it is stated that specialist expertise produced during the technical phase should inform the political phase of a controversy. However, it is worth reiterating that, in practice, this does not always happen. In their recent study of the controversy over the impact of insecticides on honeybee colonies, Sainath Suryanarayanan and Daniel Lee Kleinman (2013) demonstrated that the claims of scientists achieved ‘epistemic dominance’ over the claims of beekeepers. Echoing Jasanoff (2003), they argued that the third wave framework “[does] not consider the factors that legitimize certain claims about methods, data, and truth while delegitimizing others, factors that thus define certain actors as experts and others as nonexperts”. They concluded their case study by arguing that “beekeeper knowledge is constructed via practices that take an informal epistemic form, which makes them conducive to the highly dynamic, local, variable, and complex aspects of their operations”, and more generally that:

Understanding why certain knowledge claims are recognized and others are not demands an analysis that takes seriously the historical and structural bases for the influence of different actors’ claims in technoscientific controversies. Comprehending context and history is crucial to explaining epistemological dominance (Suryanarayanan & Kleinman 2013, p.233).

On this understanding, the history of how contributory expertise was created during the technical phase appears to partly inform its chances of influencing the political. To a certain extent, the periodic table of expertise already acknowledges that this is a possibility. It locates expertise in collective tacit knowledge, and defines collective tacit knowledge as dependent on social context. Although there may be differences between accepted and marginalised contributory expertise, this is not necessarily the only outcome of context dependency. It seems likely that different types of contributory expertise can be produced in different contexts, and that the way in which this expertise is used during the political phase is significant.

For elective modernism to be workable, understanding and acknowledging how contributory expertise can vary therefore becomes important if the best decisions are to be made about the processes used to transfer expertise from the technical phase to the political phase.

2.2.12 Summary

Studies of controversies have been central to the sociological understanding of the creation of scientific facts, the creation of technological artefacts, and technological decision-making. These studies have often focussed on the rhetoric used during controversies. However, if controversies are thought of in this way, certain features - such as materiality, secrecy and silences - may be overlooked. Recent studies of controversies have focussed on expertise. Whereas some have advocated increased public participation and a role for 'lay expertise', Collins, Evans, and others - through their invocation of the third wave, and in particular, their statement of elective modernism - have attempted to consider expertise and its political consequences in normative terms. They have distinguished between expertise based on specialist tacit knowledge from that based on ubiquitous tacit knowledge, and transmuted meta-expertise from non-transmuted meta-expertise. They have also outlined a system that could go some way to solving the problem of extension, whilst avoiding both technological populism and technocracy. Some have criticised the proponents of the third wave for overlooking aspects of the work from the second wave, and in doing so, taking a backwards step towards the first wave. However, much of this disagreement can be resolved through the acknowledgement that second and third wave work both have different goals. Work from the second wave is primarily descriptive, whereas work from the third wave is primarily prescriptive. That being said, links between the two waves must remain strong if the third wave is to develop, given that its normative ideas are informed by descriptive case studies.

Since the original paper, the third wave has done much to integrate work from the second wave - particularly through work on interactional expertise. However, consideration of contributory expertise has been, to some extent, put to one side. This is perhaps because few people working under the third wave would disagree that contributory expertise generated during the technical phase should inform the political phase. However, work has shown that some types of contributory expertise

are nonetheless excluded from the political phase. This should prompt a more detailed examination of how contributory expertise can be produced, acquired and maintained - particularly given that proponents of the third wave acknowledge that the nature of contributory expertise is dependent on the practices used to produce it. Furthermore, it should also prompt an examination of how different types of contributory expertise are transferred from the technical phase to the political phase, and once there, how they are used. Doing so may also require a deeper understanding of how expertise is assessed. Once this has been achieved, it may even be possible to suggest mechanisms that make better use of contributory expertise under elective modernism. One way that we may be able to explore the differences in the nature of contributory expertise further is through a review of the STS literature related to practices, in particular, the more recent work on the concept of multiplicity that has emerged from the so-called ontological turn. This will be the subject of the next section.

2.3 The Ontological Framework

As was described in the previous section, recent attempts to study controversies over technological decision-making within STS have focused on expertise, and in the case of contributory expertise, the practices that are used to produce it. Though the study of practices within STS has a long tradition, much recent work in this area has been carried out within an ontological framework. After briefly defining the ontological framework through a contrast with the epistemological framework, I will describe how, through a process of criticism and refinement, a strand of the ontological framework has come to foreground 'multiplicity'. I will then turn my attention to how ideas about multiplicity have been applied and adapted. I will describe how: case studies drawing upon multiplicity have expanded beyond medicine into other areas; multiplicity has been refined to develop the 'different worlds' argument; acknowledging multiplicity has foregrounded political concerns; and how multiplicity has already been used to understand aspects of controversies over technological decision-making. I will argue that, following some of the concerns that have been expressed about whether the epistemological and ontological frameworks are truly distinct from one another, together with concerns over the philosophical implications of adopting its strict form, the best

approach is to adopt ‘ontography’ - the use of the vocabulary of the ontological framework to build empirical descriptions (Lynch 2013). When thought of in this way, the ontological framework complements the study of expertise under the third wave, and there exists the possibility that the two can be used together to enhance our understanding of controversies over technological decision-making and contributory expertise.

2.3.1 The Epistemological Framework

In order to introduce the ontological framework, it is necessary to take a step back from the study of controversies, and to consider ways of understanding ‘difference’. John Law and Vicky Singleton (2005) identified two broad theoretical frameworks for understanding difference within STS: the ontological; and the epistemological. The epistemological framework takes its name from the branch of philosophy known as epistemology. Epistemology, which is derived from the Greek word for knowledge, is concerned with theories of knowledge production (Blackburn 2005, p.118). In simple terms, work done under the epistemological framework is based around the idea that difference can be accounted for through the recognition that actors view the world from different perspectives (Law & Singleton 2005). As can be seen in the examples from the previous section, this way of accounting for difference has proved to be a fruitful way of understanding controversies. Research under the epistemological framework has informed a number of key concepts and research programs within STS, including SSK and SCOT. For example, the SCOT notion of interpretive flexibility hinges on the idea that groups can have different perspectives on how a particular technological artefact should be.

Though work under the epistemological framework continues, it is often claimed that there has been an ‘ontological turn’ within STS. Some have even argued that:

The turn to ontology in STS is part of a much wider intellectual and political movement. In Western thinking this can be traced back at least as far as Nietzsche, and it now expresses itself both in post-structuralism, and in a range of empirical disciplines, including cultural studies, human geography, parts of feminism, anthropology and post-colonialism (Law & Lien 2013, pp.363-364).

Whether or not this bold claim is true, work done under the ontological framework has certain identifiable characteristics. The ontological framework takes its name from ‘ontology’, which in turn is derived from the Greek word for ‘being’. In the classic philosophical sense, ontology - a term coined in the seventeenth century - is a “branch of metaphysics that concerns itself with what exists” (Blackburn 2005, p.261). The ontological turn is usually used to refer to a series of insights that have emerged primarily from anthropological theory during the last two decades (Henare et al. 2007). How anthropologists have understood the purpose and nature of the ontological framework has varied, and given that their focus is often on the challenges associated with understanding the non-Western world, some interpretations have little direct relevance to STS (e.g. Course 2010). As a result, I’ll focus on how STS has used the ontological framework. As will become clear, within STS, the ontological framework has been concerned with the practices that are used to produce objects - where ‘object’ is used to refer to anything that can be said to exist, rather than something that necessarily occupies three-dimensional Euclidian space. Crucially, whereas the epistemological framework accounts for difference through the existence of divergent perspectives on a single reality, the ontological framework accounts for difference through the belief that divergent practices produce multiple objects, and in some of the more radical interpretations, multiple realities.

2.3.2 Ontological Fluidity

If we are to trace the development of the ontological framework within STS, as this section will attempt to do, the most sensible place to start is with some of the most prominent ways that STS has dealt with objects in the past. The ontological framework is often referred to as a post-ANT framework (van Heur et al. 2013), so it is worth considering how ANT understands objects. One of the initial aims of ANT was to understand how complex systems are able to maintain long-distance control (Law 1992). This led to the concept of ‘immutable mobiles’. Using an example from the Copernican revolution, Latour (1987) defined immutable mobiles as objects that are able to ‘move around’, whilst ‘holding their shape’. Latour argued that the Copernican revolution was only able to take place because astronomers from all over Europe began to record their observations using the same pre-printed charts. This allowed facts to move around, remain universal, and form part of a stable network. The strength of the immutable mobiles concept

lay in that it was able to describe how abstract objects - such as facts about the solar system - behaved, as well as tangible objects. Take the example of a Portuguese imperial ship. The ship must be able to move around, retain its shape (else it will sink), and form part of a network (the empire) (Law 1986). By regarding both abstract and tangible entities as objects, complex systems could be understood using the same framework.

There were, however, anxieties relating to immutable mobiles and the ANT approach. One such concern was the ‘rigidity’ that the immutable mobiles concept appeared to require. Responding to the criticism that ANT overemphasises the immutability of objects, Marianne de Laet and Annemarie Mol (2000) used an examination of the Zimbabwe Bush Pump to show how objects exhibit, and can be designed to exhibit, fluidity. Building upon an earlier argument about the fluidity of anaemia (Mol & Law 1994), they argued that the pump can be thought of as ‘fluid’ because it has the ability to change its shape yet remain functional. In short, it is a ‘mutable mobile’. For example, when the pump breaks, villagers are able to mend it using materials that were not part of the original design. Additionally, the pump can also be thought of as fluid in terms of what it is able to provide. The pump provides water, but it also provides ‘health’ if the water is uncontaminated. The fluidity of the pump is also linked to its boundaries. The pump supports a community at a local level, but at a wider level it supports a nation, as it is produced entirely within Zimbabwe’s borders. To reiterate once more, though this may appear to echo SCOT and interpretive flexibility, de Laet and Mol emphasised that “the fluidity of the pump’s working order is not a matter of interpretation” because “it is built into the technology itself” (de Laet & Mol 2000, p.225). The pump, they argued, was deliberately designed to be fluid, and later redesigns aimed to increase its fluidity.

Whilst recognising fluidity furthers our understanding of the nature of objects, it leaves some important questions unanswered. For example, there are questions that relate to the ‘problem of difference’ - namely, how can the simultaneous differences in objects be accounted for? In particular, sticking with the bush pump example, how can we account for the fact that it provides clean water at a local level, but also political stability in the form of a reliable water network at another (de Laet & Mol 2000, p.235)? The answer to this question came to form the basis of ontological ‘multiplicity’.

2.3.3 Ontological Multiplicity

Ideas about ontological multiplicity can be found in the work of Annemarie Mol. Partial elements of this concept can be found both in Mol's early work on anaemia (Mol & Berg 1994) and in related studies of atherosclerosis - a disease where deposits of fat cause the walls of arteries to thicken (Mol & Elsmann 1996, Mol 1998). In the definitive work on the subject, *The Body Multiple*, Mol (2002) provided a study of atherosclerosis treatment in a Dutch hospital. Mol used this study to argue that objects - including diseases - are constructed through practices. Importantly, Mol also argued that different groups will construct multiple versions of the 'same' object if different practices are used, thus causing reality to multiply. Clearly departing from the epistemological framework, Mol argued that:

It is possible to refrain from understanding objects as central points of focus of different people's perspectives. It is possible instead to understand them as things manipulated in practises. If we do this - instead of bracketing the practises in which objects are handled we foreground them - this has far-reaching effects. Reality multiplies (Mol 2002, pp.4-5).

Mol (2002) chose to use the term 'enact' to refer to the different processes that bring objects into existence. For example, the patient suffering from atherosclerosis may enact the disease in terms of decreased mobility, whereas a healthcare professional may enact the disease in terms of cells observed under a microscope. This is not meant to imply that all health care professionals will enact atherosclerosis in the same way. Different health care practices will also enact different versions of atherosclerosis. To illustrate this, Mol contrasted the atherosclerosis enacted in the clinic with the atherosclerosis enacted in the pathology laboratory. Thus, different instances of atherosclerosis were assigned a locality.

The recognition that different atherosclerosis objects are constructed through different practices in different locations is not meant to imply that they are kept completely separate at all times. Clearly, patients often need to embark on a single course of treatment. This can be achieved through a process of 'coordination'. Mol (2002) identified two methods by which coordination may be achieved. Firstly, a hierarchy of evidence may be constructed. This can then be used to explain away contradictory test results. For example, if a patient complains of pain

consistent with atherosclerosis, but this is not corroborated with test results, the next course of action will often be based on the fact that test results carry more weight. Alternatively, information about different enactments of atherosclerosis can be calibrated where the above incommensurability is negotiable. A process of translation can occur if common measures are established among different practices and techniques. Mol described how coordination was achieved when medical professionals were faced with conflicting results from two different practices: duplex (a non-invasive ultrasound technique); and angiography. In the end, angiography won out because one translation technique judged it to be the ‘gold standard’ (Mol & Elsmann 1996, Mol 2002).

2.3.4 Developing the Ontological Framework

The Body Multiple provided a starting point for thinking about multiplicity within STS. The work that has followed has developed these ideas. In particular: case studies have expanded beyond descriptions of health and medicine; multiplicity has been refined through the ‘different worlds’ argument; acknowledging multiplicity has foregrounded political concerns; and multiplicity has been used to understand aspects of scientific and technological controversies.

The first area I have identified is concerned with the scope of multiplicity case studies. In particular, it refers to the extent to which multiplicity has been applied across the STS landscape. As Steve Woolgar and Javier Lezaun (2013) have observed, of the case studies that have directly drawn on Mol’s ideas about practices and multiplicity, the majority have been based on case studies of biology or medicine. In addition to Mol’s (2002) study of atherosclerosis, there have been ontological case studies of hypoglycaemia (Mol & Law 2004), alcoholic liver disease (Law & Singleton 2005), biological reproduction (Thompson 2007), phantomatic organisms (Schrader 2010), and foot and mouth disease (Law & Mol 2011). This is despite the fact that “in principle, there seems to be no limit to the kinds of entity that might be treated as susceptible to enactment. Objects, persons, things, facts, theories, instruments and so on can all be enacted” (Woolgar & Lezaun 2013, p.325). It is therefore encouraging that ideas about multiplicity have also been tentatively applied to objects outside of biology and medicine - such as software. In a study of the Connexions network - a free educational content delivery system

at Rice University - Christopher Kelty (2008, p.270) briefly argued that multiplicity could be used to understand the finality of open source projects, as well as how multiple enactments of various groups are coordinated towards delivering a finished product. It would appear, then, that the scope of multiplicity is expanding, and that in the future, it will be applied to a wider and more ambitious range of objects.

The second area that I have identified refers to the depth of the multiplicity concept. In particular, it refers to how concepts associated with multiplicity have been expanded upon. One of the most significant of these additions is the 'different worlds' argument. In a study of the 2001 foot and mouth disease epidemic in the UK, Law and Mol (2011) argued, as with the study of atherosclerosis, that divergent practices resulted in the enacting of different foot and mouth disease objects. For example, at the clinical level, foot and mouth disease was enacted through the identification of deviances within small groups of animals, whereas the laboratory enacted foot and mouth disease by seeking to identify a particular virus using a microscope. A third enactment of foot and mouth disease came from epidemiology specialists who enacted it by attempting to trace the infection through the entire UK livestock population. Again, as with the study of atherosclerosis, the issue of how each of these enactments interacted to form a single foot and mouth object centred around the idea that there exists a process by which a dominant object is selected - in this case, the laboratory version.

Law and Mol (2011) argued that the differences between each enactment of foot and mouth disease were rooted in the fact that they had each been enacted in 'different worlds'. The final part of their study aimed to define more precisely the areas in which this praxiographic divergence occurred. Four factors were identified: the materials used to enact the object; the qualities of the enacted object; the staging of time when the object was enacted; and the spatial relations of where the object was enacted. Taking materials as an example, they argued that it was clear that the clinic, the laboratory, and the epidemiology department, all used different materials to enact foot and mouth disease. At the clinical level, foot and mouth disease was enacted through the examination of animal bodies. At the laboratory level, foot and mouth disease was enacted using biological samples and laboratory equipment. At the epidemiological level, foot and mouth disease was enacted using livestock transport records. Thus, in identifying these areas, Law and Mol provided a means of differentiating between enactments in other contexts.

The third area that I have identified relates to politics. It has been argued that the ontological framework, through multiplicity, should prompt the foregrounding of political concerns. Whereas the epistemological framework prompts a certain detachment, because difference can be accounted for in terms of perspectives, the ontological framework confronts difference and treats it as real. Given that, under this view, reality is actively enacted through practices, the ontological framework prompts more direct questions about what kind of world we want to create. This has been attended to in the literature through discussions of ‘experimental political ontology’ and ‘ontonorms’. Law and Mol (2011) concluded their aforementioned study of foot and mouth disease with a brief discussion of ‘ontological politics’. They stated that:

This is not a politics that works to establish goals, leaving questions of means for subsequent implementation by experts and technicians. Instead, in an ontological politics technical questions are at stake from the beginning (Law & Mol 2011, p.14).

In other words, Law and Mol argued that policy decisions should not be seen as separate from ontology, and that there is a role for the ontological framework to play in their governance. Similarly, Noortje Marres (2013) attempted to understand what a commitment to the ontological framework should imply for investigations into politics. According to Marres, “political ontology can here be taken to refer to the set of definitions that stipulate the features of specifically political entities (the state, power, citizenship, interest, democracy and so on)” (Marres 2013, p.422). Marres defined three types of investigation into political ontology: theoretical; empirical; and experimental. Theoretical ontology is simply concerned with what exists. Here, “ontology involves the stipulation of a general set of entities and relations on the level of theory or discourse, as a general blueprint of the world” (Marres 2013, p.422). This view is inadequate for understanding political entities, because it does not concern itself with how things come to exist. However, “empirical ontology differs from theoretical ontology by proposing that the question of ‘what the world is made up of’ cannot be answered wholly in theory but is partly settled in practices that must be studied empirically” (Marres 2013, p.422). As a result, much work within STS has been based on empirical investigations into how the world has been made following interventions from science and technology. Experimental ontology - which Marres advocates above the others - goes one step

further. As well as investigating how things come into being through practices, “[experimental ontology] directs attention to efforts to purposefully design politics and morality into material objects, devices and settings” (Marres 2013, p.423). In short, Marres advocates experimental ontology because “it proposes to examine how politics and democracy are accomplished through the deployment of devices, objects and settings, rather than accounting for politics and democracy in an epistemic register, that is, in terms of the deployment of discourses and ideas only” (Marres 2013, p.422).

Mol (2013) advocated a step further still. Although Mol did not precisely define the term ‘ontonorms’, she used it to refer to the behavioural norms that can often be found embedded in practices and the objects they enact. In a study of advice given to patients in consultations with dieticians, Mol found norms embedded in the practices that are that ultimately used to enact patients’ bodies. Mol argued that, the purpose of analysis should therefore not be to merely highlight that this is so, but to “[analyse] the norms embedded in practices while interfering in them through adding a novel, oblique analysis” (Mol 2013, p.481). Unfortunately, Mol did not elaborate sufficiently on what a ‘novel, oblique analysis’ might actually look like. The concept of ontonorms therefore awaits further development.

The fourth area that I have identified is concerned with cohesion. Cohesion refers to the extent to which multiplicity may be used to shed light upon other themes within STS. It is evident from the literature that there may be the potential for a link to be formed with the study of controversies. As we have already seen, there is a link between the ontological framework and politics, and some have even argued that intervening in politics should be the aim of ontological investigations. Scientific and technological controversies could provide an ideal entry point for this, given their centrality in previous studies of the relationship between science, technology and politics. As has already been discussed, ideas relating to multiplicity have already been applied to the 2001 foot and mouth disease epidemic in the UK (Law & Mol 2011). The debates over how the epidemic was governed have certainly allowed this episode to be thought of as a controversy (e.g. Donaldson et al. 2002, Woods 2004). Another example of a link between multiplicity and a controversy over technical decision-making can be found in Michelle Murphy’s (2006) study of ‘sick building syndrome’. Sick building syndrome refers to the symptoms that office workers claimed to suffer due to low-level chemical exposures from the buildings in which they worked. However, at the time of the

controversy, it remained unclear - particularly to scientists - whether the buildings were the actual cause. This lack of a consensus sparked a controversy. Murphy argued that versions of sick building syndrome:

... were brought into existence in multiple, often conflicting circumstances - the result of not just specific environments, but also new arrangements of technologies and practices through which laypeople, scientists, and corporate experts apprehended the health effects of buildings on bodies (Murphy 2006, p.8)

On this understanding, sick building syndrome became 'real' to certain groups, whilst remaining illusive to others. This is an interesting alternative take on what it might mean for something to be thought of as controversial. Rather than seeing a controversy as a result of different perspectives on a single reality, as the epistemological framework might do, ontological multiplicity offers a way of understanding a controversy as a disagreement over what exists.

2.3.5 Criticisms of the Ontological Framework

Though the ontological framework has provided the basis for a number of studies within STS, some have expressed reservations about its adoption. These have included: doubts over claims about the existence of an ontological turn; concerns over the supposed differences between the epistemological framework and the ontological framework within STS; and most seriously, the commitments that adopting a strict form of philosophical ontology would entail.

Some have questioned whether STS has truly turned from the epistemological framework to the ontological framework. This, of course, is not an easy thing to determine one way or the other. Bibliometric analysis from 2013 revealed that although there has been a recent increase in the use of ontological vocabulary, it remains unclear whether the discipline has fully turned to the ontological framework (van Heur et al. 2013). Others have questioned the basis for the demarcation between the epistemological and the ontological within STS. Woolgar and Lezaun (2013) argue that a distinction between the epistemological and the ontological would be simplistic, given the diverse nature of previous STS research:

The history of STS complicates any simplistic distinction (or transition) between ontology and epistemology. Contrary to those who see in ‘constructivism’ a programme focused on the investigation of ideational and discursive forms, the field has long advanced an analytical programme that foregrounds the instrumental, performative and material dimensions implied in the making of facts and artefacts. ‘Representation’ has rarely been treated in STS as the sort of ‘epistemological’ or meta-physical construct that some proponents of the ontological turn seem to want to turn against. When one considers the long tradition of research into the materialization of technoscientific entities, the attention to embodied practices and practices of embodiment or the classic accounts of the co-production of epistemological and political order, it is clear that the field’s interrogation of knowledge-making can hardly be described as a study of conceptual or cognitive ‘perspectives’ (Woolgar & Lezaun 2013, p.322).

In this vein, Malcolm Ashmore (2005) has advocated a both/and approach to case studies, and as such, recognises the possibility that the epistemological framework and the ontological framework can be used in conjunction.

The most serious criticism that the ontological framework has faced concerns its associated philosophical ‘baggage’ - in particular, a commitment to the existence of multiple ‘realities’. Anthropologists working outside of STS have been vocal in expressing concerns over whether those using the ontological framework oscillate between the study of what exists and the view that the “‘radical alterity’ of certain societies . . . consists not in them having different ‘socially constructed’ viewpoints on the same (natural) world, but in them living in actually different worlds” (Laidlaw 2012). Similarly, Paolo Heywood (2010) has argued that advocates of the ontological framework “use the word ‘ontology’ precisely because of the connotations of ‘reality’ and ‘being’ it brings with it; yet they neglect to acknowledge that insisting on the ‘reality’ of multiple worlds commits you to a meta-ontology in which such worlds exist: what Quine would call “a ‘bloated universe’”. These arguments are difficult to ignore, and on this basis, it seems sensible to take a step back from the commitment to the existence of multiple realities that is frequently mentioned in Mol’s work.

Michael Lynch (2013), in perceiving the same tension between a philosophical commitment to ontology and the empirical study of practices, has recommended an approach he referred to as ‘ontography’ - namely, “historical and ethnographic investigations of particular world-making and world-sustaining practices that do not begin by assuming a general picture of the world”. Lynch argued that the value of many of the studies that have been conducted under the ontological framework does not rest on the use of the framework itself, but in the rich empirical descriptions that they provide. Concerned that a philosophical commitment to ontology inevitably results in conclusions about ‘multiple objects existing in multiple realities’, and whilst also recognising that there are analytical benefits of using the ontological vocabulary, ontography “is a descriptive alternative to its grand theoretical counterpart” (Lynch 2013, p.458). Similarly, for Woolgar and Lezaun (2013), the ontological framework, or the increased use of ontological vocabulary, does allow for work within STS to focus on multiplicity:

The interest in ontology within STS points to the fact that, at least in some quarters, the analytical repertoire of the field is seen as insufficiently attuned to the multiplicity and degrees of alterity of the worlds that science and technology bring into being. In this sense, the turn to ontology would be a way of drawing out the full implications of many other turns: the materialist, performative, instrumental or experimental sensibilities developed by the field over the last two decades (Woolgar & Lezaun 2013, p.323).

Based on the above concerns, it is clear that the ontological framework should not be seen as completely distinct from the epistemological. Equally, the ontological framework should not be seen as a hybrid of the epistemological framework and the responses to those who have criticised it. Indeed, scholars who have adopted the ontological framework often acknowledge that their chosen case study could be understood using the epistemological framework, but that they have chosen the ontological framework in order to garner fresh insights (e.g. Law & Singleton 2005). Therefore, despite the reservations described above, there is undoubtedly a sense that the use of the ontological framework - in particular the ontological vocabulary in the form of ontography - can be a productive way to learn more about science, technology and society.

2.3.6 Summary

Previous studies within STS have shown that it is possible to account for difference adequately through consideration of divergent perspectives. Nonetheless, it is also possible to account for difference by shifting the focus onto the practices used to enact objects. Although it appears that prior commitments to the full philosophical implications of this view are best avoided, the vocabulary established under the ontological framework can be sufficient to reconfigure the analyst's view so as to provide alternative empirical descriptions. Since the publication of *The Body Multiple*: case studies have started to expand beyond medicine into other areas; multiplicity has been refined through the 'different worlds argument'; acknowledging multiplicity has foregrounded political concerns; and multiplicity has been tentatively used to understand controversies. With this in mind, the ontological framework appears well suited to examining the enacting of contributory expertise during controversies over technological decision-making.

2.4 Conclusion

In this chapter I have described how scientific and technological controversies have been understood within STS. Many studies have focussed on the role of interests and rhetoric. However, this focus has made it difficult to appreciate certain facets of controversies - such as secrecy, absence and silence - as well as downplaying the analytical significance of the material. More recently, studies of controversies, particularly studies of controversies over technological decision-making, have focussed on expertise. Studies have shown that scientific and technological expertise, when applied in isolation, is not always robust enough to produce satisfactory outcomes. Therefore, a case has been made for adopting a more democratic and inclusive approach to technological decision-making. In response, concerns have been expressed that this could result in a complete dissolution of the boundaries between expert and non-expert, with no special preference given to specialist knowledge. The third wave aims to chart a path between technocracy and technological populism through a reconceptualization of expertise. In assembling the periodic table of expertise, particular emphasis has been placed on understanding the nature of interactional expertise. Less attention has been devoted to exploring contributory expertise. However, given that contributory expertise is rooted in

collective practices dependent on social context, it is reasonable to ask whether, as a category, it is able to meaningfully capture the essence of the various scientific and technological contributions that can emerge from the technical phase of a controversy. One way in which it is possible to probe this is through the use of the ontological framework. In the second half of this chapter, I have shown how work using the ontological framework has emphasised how divergence amongst research practices can result in multiplicity. Although usually separate in the literature, ideas about contributory expertise and multiplicity share key tenets. In particular, both contributory expertise and multiple enactments are seen as real, and both are rooted in collective practices. On this basis, I aim to investigate the relationship between the two. I will investigate whether ideas about multiplicity can be used in a meaningful way to improve our understanding of how contributory expertise is produced. In particular, I will ask whether contributory expertise can be seen as an object under the ontological framework. If this is possible, then we might expect it to reveal a multiplicity of contributory expertises. If we suppose that contributory expertise can be studied using the ontological framework, and that the practices used to enact it are divergent, then it follows that the resulting contributory expertises will be multiple. This prompts - in the first instance - an investigation into the practices that are used to produce contributory expertise, in order to see whether they do differ from one another in a meaningful way. However, although recognising the possible existence of multiplicity is an important first step, in order for this recognition to be worthwhile, it has to be possible to convincingly delineate and characterize different enactments within the technical phase, and then understand their use during the political phase. This, at least in part, requires the ability to judge the expertise of others. Therefore the meta-expertises row of the periodic table of expertise may be able to guide this process.

If the contributory expertises produced during the technical phase of a controversy are multiple - and can be identified as such - then we may also ask what consequences this might have for the political phase. For example, it becomes reasonable to suppose that different types of contributory expertise produced during the technical phase will be available to policymakers during the political phase. However, contributory expertise on, say, the design and construction of a technology may be significantly different from the contributory expertise related to understanding systems within which that same technology will function. On top of this, contributory expertise produced in an environment dominated by certain institutional

imperatives - such as secrecy - may alter the nature of collective practices used, thus altering the nature of the contributory expertise. If contributory expertise is reconceptualized as multiple - and thus it is acknowledged that enactments can be of a different nature - the way in which each enactment is used during the political phase becomes significant. It becomes possible for it to be said that a particular enactment of contributory expertise was favoured, and another was ignored. Additionally, it becomes possible to investigate the consequences of these decisions. If these consequences are significant for the political phase of a controversy, then this should be reflected in the principles that make up elective modernism. For example, elective modernism states that the political phase should consider as much of the work from the technical phase as possible. However, the impact of favouring or ignoring a particular enactment of contributory expertise has not yet been considered. Similarly, elective modernism stipulates that the work of experts should not be distorted during the political phase. However, the use of, for example, expertise that has been enacted under conditions of secrecy may be problematic because the processes normally used to judge expertise under the third wave may not be available. This has the potential to blur the lines between the legitimate questioning of the quality of expertise, and speculation that might otherwise be viewed as an attempt to subvert expert findings.

Responses to these concerns will be sought through an examination of cryptology research and the crypto wars in the UK from 1970 to 2000. In the introductory chapter, I showed how the current literature on the crypto wars, though arguably one of the most important controversies of recent years, is poorly aligned with ideas from STS. In particular, the way in which the crypto wars have been framed has served to efface cryptology research. Though the politics of the crypto wars have been well described, the practices used to create the relevant cryptology expertise have been overlooked. Using the terms from the third wave, the rhetoric of the political phase has been described, but the work of the technical phase, and the relationship between the two, have not. This has created the false impression that contributory cryptology expertise arrived at the political phase of the crypto wars ready-made. Given that the crypto wars was a controversy heavily influenced by experts, and given the importance of secrecy in cryptology research, a re-examination of the crypto wars provides an ideal opportunity for the development of ideas related to multiple contributory expertises, the transfer of expertise from the technical to the political phase, and finally, to better understand an important episode in the recent history of electronic communication.

On this basis, the following three research can be formulated:

1. How was contributory cryptology expertise produced during the technical phase of the crypto wars in the United Kingdom?
2. Can the ontological framework and the third wave be used in conjunction to develop a re-conceptualization of the production of this contributory expertise as ‘multiple’?
3. What were the consequences of this multiplicity of contributory expertises during the political phase of the crypto wars?

The next chapter will be devoted to a consideration of the best way to provide answers to these questions.

Chapter 3

Methodology

3.1 Introduction

In this chapter I will describe the methods used to answer the research questions posed in the concluding section of the previous chapter. In contrast to many studies within STS, I will begin by providing a description of the background assumptions that have underpinned the more specific methodological choices. I will describe the reasoning behind my decisions to use case study and historical approaches, with respect to the area that my research aims to investigate - namely, the relationship between cryptology, expertise, and controversies over technological decision-making. With this in mind, I will then provide a detailed description of the specific documentary analysis and semi-structured interview methods used, and attempt to justify their combined use.

3.2 Methodological Background

My approaches and methods - and what I see as their hierarchical relationship to one another - can be expressed in a simplified flowchart (see Figure 3.1). As Figure 3.1 illustrates, following on from the formulation of my research questions, I decided to use a qualitative approach, and to adopt a social constructivist worldview. I also decided that my chosen case study would be examined using historical research methods, and that most of my data would come from documentary analysis and semi-structured interviewing.

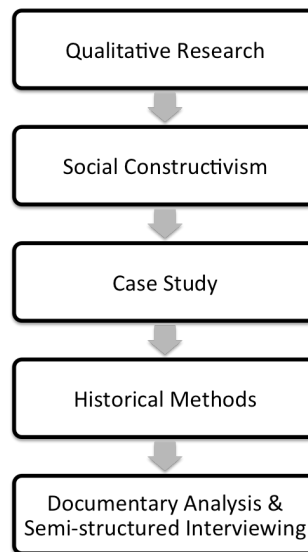


FIGURE 3.1: Methodology Flowchart

Much STS research does not give a detailed description of the methods used to reach its conclusions. Many studies within STS do not discuss research methods at all. Steven Yearley has argued that this is linked to a general lack of reflection on STS' inherent assumptions (Yearley 2005, p.107). However, STS is not completely blind to methodological ideas. Though many do not engage with methodological ideas found in other areas of social research, pointers to the style of investigation used can be found in theoretical approaches. For example, the Strong Programme emphasised reflexivity on the part of the researcher, and provided a basic outline for research into SSK (Bloor 1976). However, it has often been left unclear how the goals of the Strong Programme can be achieved in terms of the methods of investigation used, and how they relate to other ideas from social research more broadly. As a result, it would appear beneficial to link some of the latent methodological ideas within STS to some of the well-developed concepts found in social research. Thus, the first part of this chapter will be devoted to justifying my decision to use case study and historical methods. The second half will be devoted to providing some specific details regarding my data collection and analysis techniques.

3.2.1 Case Study

In common with much research in STS, this study could broadly be described as qualitative and social constructivist. Alan Bryman has described the qualitative

approach as one that “predominantly emphasises an inductive approach to the relationship between theory and research, in which the emphasis is placed on the generation of theories”, “has rejected the practices and norms of the natural scientific model and of positivism in particular in preference for an emphasis on the ways in which individuals interpret their social world”, and “embodies a view of social reality as a constantly shifting emergent property of individual’s creation” (Bryman 2008, p.22). Jan Golinski has observed that most of the major theoretical and practical approaches designed for STS - such as the Strong Programme and EPOR - are essentially social constructivist in nature, given that social constructivism “draws attention to the central notion that scientific knowledge is a human creation, made with available cultural and material resources, rather than simply the revelation of a natural order that is pre-given and independent of human action” (Golinski 2005, p.6). When the aims and objectives of this research - as discussed in chapters 1 and 2 - are considered alongside these descriptions, the use of a qualitative approach and a social constructivist worldview were clearly appropriate.

John W. Creswell (2007) has identified five approaches to qualitative research: narrative; phenomenological; grounded theory; ethnographic; and case study. However, in practice, many research projects will incorporate elements from more than one of the above approaches. My research was no exception. Given that one of the aims was to describe and interpret the practices of cryptology scientists through the “shared and learned patterns of values, behaviours, beliefs and language of a culture sharing group” (Creswell 2007, p.68), it certainly shared some of the basic characteristics of ethnographic research. However, there are some important ways in which my research departed from it. The main reason is that I was aiming to describe and interpret events from the past. As a result, I was unable to collect data through participant observation - often considered the most illuminating method in ethnographic research. Studying the past also prevented me from taking part in some of the activities considered integral to good ethnographic research within STS - such as attending scientific conferences and observing laboratory work (Hess 2001). On top of this, there are also some more general ways in which research thought of as ethnographic within STS differs from much contemporary anthropological ethnography. David J. Hess (2001) has noted that STS ethnographies are typically: focussed on a world of which the researcher is a part; require a more symbiotic relationship between researcher and participant; and are often undertaken in order to subvert an existing ‘standard’ historical or sociological description.

On this basis, my research can be more accurately thought of as a case study. For Creswell, case study research “involves the study of an issue explored through one or more cases within a bounded system” over time, through the use of multiple sources of information (Creswell 2007, p.67). Case study research may examine an issue using one or more case studies, and within those, one or more research sites. The key methodological issue related to the case study approach concerns how the case itself is selected, and following on from that, how that case is bounded. Robert E. Stake (1995) has argued that a case study requires a certain degree of specificity or ‘boundedness’ in order to be useful. However, a precise method for measuring boundedness does not exist. As such, the choice of case study, and the way in which that case is bounded, appears to be in no small part subject to the researcher’s judgement. Therefore, the case study researcher has a responsibility to justify their choices and be aware of how decisions may impact upon the research.

When considering the case study approach, there may also be concerns about the apparent trade off between depth and generalizability. Whilst case studies can offer a richer, more in-depth description of a scenario, this can be seen to be at the expense of descriptions that can be applied to other scenarios. However, as Bryman has pointed out, case study researchers rarely delude themselves into thinking that they are producing results with a high degree of external validity (Bryman 2008, p.55). Bent Flyvbjerg (2006) has gone further in defending the case study rationale. Flyvbjerg rejected the ideas that: general theoretical knowledge is more valuable than case study knowledge; one cannot generalize from a case study; case studies are useful for generating hypotheses but not for testing them; case studies have a tendency towards the confirmation of the researchers pre-conceived biases; and that it is difficult to develop theoretical positions from specific case studies. Flyvbjerg argued that the search for predictive theories and ‘universals’ in the social sciences is in vain, and in this sense case study knowledge can be more valuable. Furthermore, he argued that the idea that advances in the sciences are made on the basis of generalizability is often unfounded, and that case study research can be used as a way of falsifying a theory. He also argued atypical or extreme cases can often reveal more about a particular situation because of their ability to activate more social mechanisms. Finally, Flyvbjerg acknowledged that verification bias is a serious issue, but argued that it’s a feature of all types of research, even in the so-called hard sciences.

There is a clear precedent for the use of case study research within STS. Although

some STS research is thought of as ethnographic, much is based on case studies. Indeed, the prevalence of the case study approach in STS has occasionally generated exasperation from those working in the field (Beaulieu et al. 2007). Nonetheless, it is undoubtedly an approach that is able to produce valuable insights. Although concerns about generalizability may be largely unfounded, the case study researcher must provide a justification of their choice of case study and the way in which that case study was bounded.

Phase	Conceptual Bounding	Geographical Bounding	Temporal Bounding
Technical	Restricted to collective cryptology research	Restricted to research carried out at a site within the United Kingdom	Restricted to research carried out between 1970 and 2000
Political	Restricted to attempts by the government to legislate on matters related to cryptology	Restricted to attempts by the United Kingdom government to legislate	Restricted to attempts to legislate between 1970 and 2000

TABLE 3.1: Case Study Bounding

Though there is general agreement that case studies should be bounded, there is no clear consensus on exactly how this should be done. Therefore, I decided to bound my case study using three of my own criteria: what my chosen theoretical concepts permitted consideration of; the geographical location where events could take place; and the period of time during which these events could have happened (see Table 3.1). The conceptual bounding of my case study was determined by existing third wave concepts. Although many of the types of expertise from the periodic table influenced the crypto wars, I limited my study to the role of contributory expertise. Furthermore, given that the technical phase is currently defined as being informed by the scientific form of life (Collins et al. 2010), I restricted my study of practices from the technical phase to those sites that actually carried out collective cryptology research, rather than, say, sites that carried out relevant research into law or politics. In terms of this geographical bounding, Goldsmith and Wu's (2006) injunction regarding an awareness of how national context can shape debates relating to the Internet chimed with what was known about how the political phase varied with national context (see Koops 2013). As such, it appeared reasonable to restrict my case study to developments that occurred in

the UK, whilst maintaining an awareness of possible contributions from outside. In terms of temporal bounding, the relatively short period of time between the advent of an open cryptology research program in the 1970s and the eventual passing of legislation around 2000 presented the opportunity to study the development of UK-based cryptology research in its entirety.

Given that the aim was to investigate the emergence of multiple contributory expertise, it appeared necessary to carry out a ‘multi-sited’ case study. However, prior to carrying out the research, a decision still needed to be made about what sites to include. Importantly, prior to carrying out the research, I tried not to let my prior knowledge of the crypto wars dictate my choice of research sites too heavily. This was because, given that my literature review had identified that secrecy, silences, and absences might be important, I needed to give these factors the opportunity to emerge, and working backwards from the overt politics of the crypto wars may have effaced them. But perhaps most importantly, I decided that it would be a mistake to use this information to set unmovable boundaries around my case study, given that information gathered during the research may have prompted a rethink of the reasoning used to bound it. As such, I tried to remain open to the possibility that the boundaries may need to evolve with the data.

3.2.2 Historical Methods

After deciding to use the case study approach, and after deciding how to bound it, the next step was to think about the best way to probe the case. Given that I had bounded my case study to examine events from the past, it was clear that I would be carrying out a form of historical research. This is not unusual in STS or the social sciences more generally, particularly if the case study approach is used. However, this does prompt methodological concerns. A traditional view of history and sociology suggests that they attempt to explain things in different ways. According to the terms introduced by the German philosopher Wilhelm Windelband (1848-1915), history can be thought of as ‘idiographic’ - because it seeks to describe singular, particular cases, whereas sociology can be thought of as ‘nomothetic’ - because it seeks to generalize phenomena using theory. However, by the final quarter of the twentieth century, the idiographic-nomothetic distinction had largely fallen out of favour, and the subtle similarities between sociology and

history came to be appreciated. Some even argued that history and sociology were essentially the same. Philip Abrams - in advocating 'historical sociology' - argued that "sociological explanation is necessarily historical" and that "historical sociology is thus not some special kind of sociology; rather, it is the essence of the discipline" (Abrams 1982, p.2). For Abrams, "historical sociology is not a matter of providing historical background, nor is it a matter of imposing grand explanations such as evolution onto the social. It is the attempt to study the social as something that is constructed in time" (Abrams 1982, p.2). Though this defines historical sociology rather broadly, many other definitions see historical sociology as primarily concerned with the development and emergence of modernity - and other large-scale phenomena such as capitalism and the state - over long periods of time (Delanty & Isin 2003, Lachmann 2013). In the case of my research, I did not believe that the historical sociology label was appropriate. Although I argue that understanding the relationship between scientific research and controversies over technological decision-making is important, under the third wave, this is not typically linked to grand sociological themes and does not draw on historical evidence from a period of time fundamentally different to the present.

Others have argued that, although history and sociology share certain similarities, they differ in terms of the available data gathering methods. According to John H. Goldthorpe (1991), historians are concerned with finding evidence from among a stock of relics, whereas sociologists, whilst also able to draw upon relics, have the option of generating their own evidence in the present using ethnographic methods. Therefore:

Because sociologists have the possibility of producing their own evidence - over and above that of exploiting relics - they are in a position of advantage that should not be disregarded or lightly thrown away. In other words, sociologists should not readily and unthinkingly turn to history: they should do so, rather, only with good reasons and in full awareness of the limitations they will thereby face (Goldthorpe 1991, p.214).

Goldthorpe used this as a platform to caution against the "perverse recourse to history" when attempting to describe phenomena that can be studied in the present, and to launch a broader critique of historical sociology. On this basis, there does appear to be a case for being aware of the limitations of using historical methods

such as documentary analysis, as opposed to certain ethnographic methods such as participant observation, to produce sociological insights. It also implies that for the examination of cryptology research practices and cryptology expertise, using historical methods may not be as effective as using sociological methods.

It should be noted that many have been critical of Goldthorpe's claims, arguing variously that he either oversimplified the distinction between historical and sociological methods, or the nature of modern attempts at historical sociology (Bryant 1994, Mann 1994, Hart 1994, Mouzelis 1994). More specifically, though ethnographic methods have been used in the past to study laboratory activity, they have tended to study the processes involved in experimental work. Here, it is thought, ethnographic methods are appropriate because there are lots of observable, tactile, micro-social practices for the sociologist to record, whereas historical methods are inappropriate because these same practices are often absent from written records or considered unimportant by those who had carried them out. However, as Lisa Garforth (2012) has pointed out, it is not clear that non-experimental laboratory activity is as readily observable. For example, much scientific work requires thinking time, and much cryptology research requires time spent sat at a computer. As a result, when attempting to study this type of cryptology research practice, and similar research practices used in other mathematical sciences, ethnographic methods and historical methods may prove equally unrevealing. However, this is not necessarily true of unobservable, institutional, macro-social practices - such as departmental management, seeking funding, communicating with external partners, and so on - that may be equally important in understanding the knowledge-making process. Although largely unobservable in the laboratory, these practices leave traces in documents and memories that can be accessed using documentary analysis and interviews.

Crucially, I decided that the appropriate response was to recognise that, if historical methods are to be used, they are likely to reveal and emphasise different types of practices, and that this will shape any conclusions that may be drawn. Finally, it is worth noting that for any study that aims to empirically investigate the relationship between a technical phase and a political phase, the historical approach may be the only available option. This is because it is questionable whether it would be feasible to justify the study of the technical phase in the present using ethnographic methods, with the aim of understanding how the expertise that resulted was used during a political phase that may or may not occur in the future.

3.3 Data Gathering

Following on from my discussion of the methodological assumptions that underpin my research, I will now discuss some of the more practical considerations associated with my data gathering techniques. An overview of my research plan will be given, followed by a description of the documentary analysis and semi-structured interviewing techniques used.

3.3.1 Research Design

In this subsection, I'll provide an overview of my fieldwork and analysis design. As Robert K. Yin pointed out, "unlike other research methods, a comprehensive catalog of research designs for case studies has yet to be developed" (Yin 2009, p.25). This creates a degree of freedom for researchers when designing their case studies. My fieldwork and analysis can be divided into three stages (see Table 3.1).

Stage	Phase	Methods	Analysis	Research Activity
First	Technical	Documentary analysis; Retrospective semi-structured interviewing	Thematic	Collection of data relating to the production of contributory cryptology expertise at various sites within the UK between 1970 and 2000
Second	Political	Documentary analysis; Retrospective semi-structured interviewing	Thematic	Collection of data relating to the development of UK legislation related to cryptology between 1970 and 2000
Third	Transfer	None	Concept and theory development	Attempts to understand how expertise from research sites was transferred to, and used during, political processes

TABLE 3.2: Updated Research Design

As I mentioned in the previous section, I decided to carry out a multi-sited case study in order to allow for multiple contributory expertises to emerge from the technical phase. Although I knew that I was only interested in sites where scientific research into cryptology was undertaken, I still had to formulate a list of potential sites and then decide upon which of them to study. After I had done this, I reasoned, locating the relevant sources of data would be relatively straightforward.

However, as was explained in the introductory chapter, I also knew that this would be complicated by the fact that an historical account of cryptology research in the UK did not exist. Therefore, in choosing research sites, it would not be possible to make use of a form of probability sampling. Therefore, I decided to carry out some preliminary historical research in order to identify potential sites. Once I had identified a handful of sites, I decided to proceed with the fieldwork, and from then on, use a mixture of purposive and snowball sampling to acquire information about any additional research sites. When examining documents, I made sure to check for any information relating to cryptology research undertaken at other sites, and when interviewing, included questions that probed for information relating to the work of others. This ultimately led to the identification of the following five research sites, all of which produced cryptology expertise during the technical phase:

1. Data Security Group, National Physical Laboratory
2. Information Security Group, Royal Holloway College, University of London
3. Mathematics Division, Racal Electronics plc
4. Security Group, University of Cambridge
5. CESG, Government Communications Headquarters

At quite an early stage in the fieldwork, I felt that I was approaching a saturation point in terms of cryptology research sites, so I decided to focus on these five.

For the first stage of my fieldwork and analysis (see Table 3.2)¹, the broad aim was to gather and analyse data relating to the technical phase of the crypto wars. This resulted in the gathering and analysing of data on the research practices used to produce contributory cryptology expertise at various sites in the UK between 1970 and 2000. For each site, this involved locating and examining documents that shed light on the research practices used. Then, after the documentary analysis stage was complete, potential interviewees were identified and approached. Interviews were then used to ‘triangulate’ information contained in the documents, and to gather data on research practices typically not available from documentary sources. The data from the documents and the interviews was then analysed in order to

¹This table is the same as Table 1.1, but with an added ‘Phase’ column following the discussion of technical and political phases in the literature review.

identify themes. For the second stage, the broad aim was to gather data relating to the political phase of the crypto wars. This resulted in the gathering and analysing of data on the development of UK legislation related to cryptography between 1970 and 2000. Here, the same pattern of documentary analysis and semi-structured interviewing was used. Official documents produced during the parliamentary process were used as a starting point. This would then lead to other documents, and would allow for the creation of a list of potential interviewees. Again, the data from the documentary analysis and the semi-structured interviews was analysed in order to identify themes. For the third stage, the broad aim was to develop concepts related to the ‘transfer’. Here, I use the term ‘transfer’ to refer to the processes used to transfer expertise from the technical phase to the political phase.² This resulted in an examination of how contributory cryptology expertise was transferred and then used during the legislative process.

3.3.2 Documentary Analysis

Turning now to the specific data collection methods used, the first two stages of fieldwork both involved some form of documentary analysis. For Bryman, an analysable document is something that: can be read (including visual material); has not been produced specifically for the purpose of social research; has been preserved; and is relevant to the concerns of the researcher (Bryman 2008, p.515). As a result, this rather broad definition includes paper-based documents, electronic documents, photographs, and websites.

Documentary analysis has much in common with historical enquiry. Therefore, many of the arguments made during discussions about the nature of history are relevant. These discussions are usually thought of - at least initially - in terms of the Carr-Elton debate. E. H. Carr’s *What is History?*, published in 1961, argued against a then-dominant empirical view. Carr argued that “the belief in a hard core of historical facts existing objectively and independently of the interpretation of the historian is a preposterous fallacy”, and thus espoused a form of historical relativism (Carr 1961, p.12). In response, G. R. Elton’s *The Practice of History* (1967) defended the view that emerged from the work of the nineteenth century German historian Leopold von Ranke - namely, that it is possible to arrive

²‘Transfer’ cannot yet be thought of as a phase under elective modernism. It has been placed in the ‘Phase’ column of Table 3.2 to preserve symmetry.

at objective truth about the past through the careful study of primary source documents.

Though some discussions of the nature of history are still thought of in terms of the Carr-Elton debate, it has been largely supplanted by the postmodern critique of history. Although it isn't possible to discuss the postmodern critique in detail, it is useful to observe, as Richard J. Evans (1999) has done, that the critique can take 'radical' and 'moderate' forms. In its moderate form, the postmodern critique of history encourages historians to take problems of interpretation seriously, and to take a reflexive approach to looking at the cultures and belief systems they find themselves in. The radical form attempts to completely undermine any attempt to understand the past by arguing along the lines that "language cannot relate to anything but itself" (Evans 1999, p.3). As Evans argued, "once postmodernist hyperrelativism's principles are applied to itself, many of its arguments begin to collapse under the weight of their own contradictions", and are thus difficult to take seriously, despite their supposed implications for history and sociology as modes of enquiry (Evans 1999, p.190). However, ideas emanating from the moderate view seem to offer a more nuanced way to understand documents, as well as chiming with the reflexivity tenet of the Strong Programme (Bloor 1976).

Taking a more nuanced view of what documents can tell us requires an assessment of their qualities. John Scott (1990) has developed four criteria for aiding researchers with this process, namely: authenticity; credibility; representativeness; and meaning. Authenticity requires an assessment of "whether [the document] is actually what it purports to be" (Scott 1990, p.19). Clearly, in extreme cases, if a document deliberately purports to be something that it is not, this has obvious implications for the interpretation of its content. Authenticity also requires an assessment of whether a document is an original or a copy, and thus an assessment of likely errors in the copying process. Credibility requires an assessment of how distorted the contents of a document is likely to be. Of course, all documents are distorted in the sense that they provide information selected by the author. Thus, "the question of credibility concerns the extent to which an observer is sincere in the choice of a point of view and in the attempt to record an accurate account from that chosen standpoint" (Scott 1990, p.22). Assessing representativeness "involves a judgement as to whether the documents consulted are representative of the totality of relevant documents" (Scott 1990, p.24). Though good research can be conducted without the totality of relevant documents, an honest appraisal of

what is missing can help with the formulation of more realistic analyses. Finally, meaning requires an assessment of whether the document is readable, and if it is, how it can be interpreted.

All four of the above criteria are linked. Of the four, meaning appears to have the largest bearing on the aforementioned concerns related to the moderate post-modern critique. There are many different approaches to the interpretation of documents. Miriam Dobson and Benjamin Ziemann (2009) advocated a checklist approach. They encouraged researchers to consider things like: the context in which the document was produced; the key concepts it draws upon; the metaphors and binary distinctions that are used; and the way in which the document addresses the audience. These are undoubtedly important considerations. However, it is also important to consider the nature of the object of study when deciding on an interpretive framework. Given that I used documents to investigate research practices, Paul Atkinson and Amanda Coffey's (2011) ideas on 'documentary realities' were particularly relevant. Atkinson and Coffey analysed UK Research Assessment Exercise (RAE) documents to show that "documents are not neutral, transparent reflections of organisational and occupational life" but that "they actively construct the very organisations they purport to describe". As such, documents should be interpreted in terms of what the author wanted to convey to their audience and what the document was produced in order to accomplish, rather than as naturally occurring analogues of reality.

On this understanding, the documents produced by an organization should be seen as an integral part of certain working practices, rather than a description of them. Therefore, in terms of my research, documentary analysis was not carried out in the hope of finding a set of documents that described in detail how scientists went about their research into cryptology. Rather, documents were seen as something produced during the course of research activity, and as an integral part of the practices used to produce cryptology expertise. Therefore, in terms of analysis, documents were used to make inferences about the nature of the practices of which they were a part. These inferences could then be checked during interviews with those that had actually been involved in their production.

3.3.3 Practicalities of Documentary Analysis

I will now describe some of the more practical steps taken during my research with respect to documentary analysis. Most of the documents that I examined were sourced from archives.³ A small number were sourced from the Internet, which, it could be argued, shares certain similarities with traditional archives. In both cases, it is important to appreciate that archives are essentially incomplete, and thus the material contained in them paints an imperfect picture of their sources. Michael R. Hill (1993) used the term ‘archival sedimentation’ to refer to the processes by which material finds its way into archives. The eventual content of an archive can be determined by: actors at the primary stage - when the material is held by those who created it; the secondary stage - when material is transferred to the archive; and the tertiary stage - when the material is stored at the archives. All three contain the potential for material to be lost through accidental or deliberate removal. Therefore, Hill advises, it is important to be aware that archives can both challenge and deceive.

Though most of the processes I used to examine archival material could be describe as standard practice, I did make use of two rarely-used techniques in social research. The first of these was the use of Freedom of Information (FoI) Requests. In the UK, the Freedom of Information Act 2000 allows any person to request, in writing, details about whether a UK public body holds information on a particular topic. If the public body possesses that information, then the individual who made the original request can make a further request to be granted access to that information. Though the potential uses of FoI requests are clear, they are rarely used for social research (Brown 2009). Researchers may use FoI requests to access information held by public bodies that is not published or archived. However, FoI requests can also be used to access archival material protected by the so-called thirty-year rule - a principle set out in the Public Records Act 1958 that prevents access to records created by public bodies until thirty years after their creation. Given that my research was concerned, in part, with events between 1970 and 2000, there existed the potential to use FoI requests to access otherwise withheld information. I submitted a total of 22 successful FoI requests.⁴ Once accepted, the requested information could usually be viewed at archives in the normal way.

³A full list of the archival sources used can be found in Appendix B.

⁴Archival sources accessed using FoI requests are marked in Appendix B.

I also made use of an Internet Archive tool known as the Wayback Machine.⁵ The Wayback Machine provides access to an online digital archive of the World Wide Web. It allows a user to view digitally archived copies of web pages from as far back as 1996. This was particularly useful for my research, as it allowed me to view past versions of web pages created by my chosen research sites. It also allowed me to view the complete web archive for the UK Cryptography Policy Discussion Group (ukcrypto) mailing list, which, as will be described in Chapter 8, was a hub of communicative and organizational activity during the political phase of the crypto wars. The web page which hosts the ukcrypto archive only contains the most recent three years worth of conversations, but the Wayback Machine was used to access the most recent three years worth of conversations from any given date, thus effectively allowing for complete access.

Stage	Documents Analysed
First	Research papers, committee minutes, personal documents, published histories, conference records, newsletters, internal publications, product catalogues, product manuals, sales material, prospectuses, archived websites.
Second	Green papers, White papers, Hansard, consultation documents, consultation responses, Select Committee meeting records, Select Committee reports, press releases, research papers, mailing list posts, archived websites, conference records.
Third	None.

TABLE 3.3: Types of Document Analysed

In broad terms, documents were analysed using qualitative content analysis, in that the analysis consisted of the “searching-out of underlying themes” (Bryman 2008, p.529). Once the documents for each research site had been collected for the first stage, they were analysed in order to bring out the themes specific to each. This was done through a process similar to the coding of interview transcripts, and as such, had much in common with David L. Altheide’s (2004) version of ‘Ethnographic Content Analysis’. Documents were analysed iteratively in order to identify themes that could then be used to formulate a description of the practices used to enact expertise. A similar process was used for the second stage. However, in this instance, documents were initially analysed in order to produce a timeline

⁵The Wayback Machine tool can be accessed at: <http://web.archive.org>

of events. This timeline was then coded in order to identify the individual debates that made up the political phase.

3.3.4 Semi-Structured Interviewing

In order to complement the documentary analysis undertaken in the first two stages, I also carried out a series of semi-structured interviews with relevant actors. Interviews were, where possible, carried out after the relevant documents had been analysed. Therefore, the interviews acted a source of new data, as well as a way of validating the data gleaned from documents.

There are various different types of interview used in social research. These types are often thought of as lying on a spectrum, with structured interviewing - where the interviewer is expected to stick rigidly to an interview guide, at one end, and unstructured interviewing - where the interviewer will form questions based upon an interview guide that carries only a list of general topics to be covered, at the other. Semi-structured interviewing sits somewhere in-between the two. Here:

The interviewer has a series of questions that are in the general form of an interview schedule but is able to vary the sequence of questions. The questions are frequently somewhat more general in their frame of reference from that typically found in a structured interview schedule. Also, the interviewer usually has some latitude to ask further questions in response to what are seen as significant replies (Bryman 2008, p.196).

According to Bryman (2008, p.437), semi-structured and unstructured interviews differ from structured interviews - which have much in common with surveys - in that: they place less emphasis on the measurement of reliability and validity; they are more amenable to the generation of theory; there is a greater interest in the interviewee's point of view; rambling or going off on tangents is encouraged as it suggests what is important to the interviewee; greater use is made of prompting and probing; and interviewees can be interviewed more than once. As such, semi-structured interviewing aligned well with the aims of my research, given that I hoped that the interviewees would steer my questions towards themes and pieces of information they felt were significant.

There are numerous types of semi-structured interview. Given that my aim was to use the interviews to learn about events from the past, they could be most accurately described as ‘retrospective interviews’ - a term coined by Andrew Pettigrew (1985) to refer to interviews he conducted at Imperial Chemical Industries (ICI) on their past use of organizational development expertise. Here, as with oral history interviews, the chief methodological issue concerns the impact that memory might have on the quality of the data. The interviewee may have completely forgotten certain details, or perhaps more importantly, the way in which they perceive and understand events from the past may have changed during the intervening period. Linda Shopes (2011) has argued that, following the so-called ‘linguistic turn’ in the second half of the twentieth century, whereas previously there had been a tendency to take the content of interviews about the past at face value, broader intellectual trends prompted a re-examination of what they can tell us. Michael Frisch (1998) argued that, instead of viewing interviews as a way to get closer to the past by bypassing prior historical interpretation, they should be used to examine what happens to personal experience on the way to it becoming memory, and on the way to it becoming history. As Shopes elaborated:

Meaning is conveyed through language, which is in turn shaped by memory, myth, and ideology and through non-verbal expression and gesture, which give both immediacy and emotional depth to the exchange and further command the listeners’ attention. Interviews thus offer clues into narrators’ subjectivities - the intersubjectivity - between narrator and interviewer. Understood in this way, interviews are not documents in the traditional sense, to be mined for facts, but texts, to be interpreted for ways narrators understand - and want others to understand - their lives, their place in history, the way history works (Shopes 2011, p.458).

Therefore, with respect to my research, interviews were used as a way of revealing how actors understood the nature and purpose of their cryptology research, how it related to the research of others, and how it related to the crypto wars.

One way in which researchers using semi-structured interviewing have attempted to achieve this interpretation has been through efforts to establish a rapport with their interviewees. Methodological reflection on research that has drawn on interviews with scientists and other professionals has highlighted two key issues that

may be linked to a lack of rapport: the tendency for interviewee responses to take on the character of official lines; and the tendency for interviewee responses to be delivered in a neutral, teacher-like way. For example, during her interviews with civil servants, Karen Duke realised that “there was indeed an ‘official line’” and that her job was “both to recognise it and probe beyond it” (Duke 2002, p.42). In this particular example, Duke’s interviewees were insistent that they were responsible for ‘implementing’ rather than ‘making’ policy. Similarly, in his study of macroeconomists, Neil Stephens noticed that, on occasion, the interview “would adopt the pattern of a lecturer/student relationship where the interviewee would frequently lean towards teaching the technical issues of macroeconomics as opposed to placing values upon them”, and that there existed “the potential for the teaching voice to depersonalize the account”, forcing the researcher to “resituate the conversation onto the personal position of the interviewee in the debate, rather than recounting the consensual ‘perceived wisdom’” (Stephens 2007, p.208). Responses like these clearly have the potential to create problems during an interview, particularly when the interviewer is asking about an inherently controversial topic. Therefore, I decided that my interviews should more closely resemble an informal encounter in order to avoid the characteristics of a doctoral supervision or a formal data gathering exercise.

There is also the issue of the appropriate level of scientific knowledge demonstrated by the interviewer when communicating with the interviewee. In common with many scientific fields, cryptology and computing are laden with terms and concepts that require specific knowledge and training to unpack. To a certain extent, this is also true of the legislative and policymaking processes surrounding controversies over technological decision-making. The methodological salience of the use of scientific knowledge is proportional to the extent to which the joint construction of meaning between interviewer and interviewee is prioritised. Nonetheless, outside of recent work on interactional expertise, this is an issue that has been infrequently discussed in relation to interviews with scientists. In an exception, Grit Laudel and Jochen Gläser (2007) drew on ideas from laboratory ethnographies in order to formulate their own methodological position. They considered examples of the naive observer, such as Bruno Latour and Steve Woolgar’s *Laboratory Life* (1979) - where the attempt was made to shed any prior knowledge of the science in question and of laboratory practices. They also considered examples of the informed observer, such as Harry Collins’ (1998) study of the search for gravitational waves - where a conscious attempt is made to achieve a level of

comprehension and understanding deemed acceptable for sociological study. Finally, they considered examples of the native observer, such as Andrew Pickering's (1984) study of particle physics - where a former professional scientist attempts to examine their own field from a sociological perspective. Insofar as these positions apply to interviewing, the idea of the informed interviewer is perhaps the most preferable, as it strikes a balance between, on the one hand, the impracticability (and perhaps impossibility) of consciously aiming to shed prior assumptions, and on the other, the problems of 'going native'. It also allows for the possibility of developing a rapport with the interviewee, given that "informed questions signal the interviewees that you have done your homework, made an effort, and have not just come to pick their brain", and that "you have gone as far as you can go with the available material and now you need some help" (Rubin & Rubin 1995, p.195).

In terms of the implications of these arguments for my research, the problems associated with memory appear to be the most salient. As with documentary analysis, they serve to caution against seeing interview data as a mirror of the past, thus placing importance on triangulation with data from other sources. When conducting interviews, I attempted to maintain an awareness of the potential for my interviewees to make factual mistakes, but also that, in some cases, the way in which they looked back on their careers and the crypto wars would be shaped by what had happened in the intervening period. In terms of rapport, given my own undergraduate and employment background in computing science, and the fact that I conducted my documentary analysis before carrying out my interviews, I attempted to adopt a position that was close to the informed observer or the interactional expert. This, it was hoped, would go some way to improving rapport, and would thus reduce the likelihood of being fed official lines or being taught basic cryptology. My concern with improving and maintaining a rapport with my interviewees also extended to my practical approach to interviews. Broadly, as will be described in the next subsection, I decided that my approach would be informal and flexible, given that I did not want to do anything that might alienate my interviewees, and because being overly formal might increase the likelihood of being fed official lines or being treated as a student receiving supervision.

3.3.5 Practicalities of Semi-Structured Interviewing

One of the most important sets of decisions that had to be made during the course of the fieldwork centred on whom to interview. At the most basic level, the population was specified by the conceptual, geographical, and temporal case studying bounding, as defined earlier in this chapter (see Table 3.1). To recap, this meant that I wanted to interview individuals who had: carried out collective cryptology research; carried out this research in the UK; and had carried out this research between 1970 and 2000. However, in practice, I also knew that there were a number of reasons why it would be impossible to interview some of the individuals within this population. For example, given the historical nature of the case study, some of those who had carried out collective cryptology research in the first half of my specified period - particularly those who had held senior positions - had died. Similarly, given the amount of time that had passed between the start of the period under study and the present, it would have been unrealistic to expect to be able to identify every individual within the population. As a result, from the outset, I suspected that the number of individuals within this population was small, and that the number that could realistically be sampled from this population was even smaller.

Due to the lack of literature on the history of cryptology research in the UK during this period, before the fieldwork proper began, there existed no obvious source of a list of individuals within the specified population. However, preliminary research had revealed the names of some of the most prominent individuals and institutions. This was used as a starting point for documentary analysis. As part of this analysis, documents were used to identify the names of individuals who were part of the specified population. In some cases, the available documents included detailed administrative material - such as organizational charts - that contained complete employee information for a particular period. In other cases - where this material was not available - the names of individuals were gathered in a piecemeal fashion from a variety of documents that happened to mention employee information. Furthermore, during the interviews, interviewees were asked for information about others working in the field at the time. As such, the sampling method used shared many of the characteristics of snowballing.

Although it is very difficult to provide exact figures, based on information contained within the available documents, I estimated that the total population (including those who had died or were unidentifiable) numbered around 40 individuals. Of these, I was able to identify 24 individuals. Once a list of potential interviewees had been produced, the next step was to obtain their respective contact details. Given that many of individuals had been identified from documents originating from as far back as the 1970s, this required considerable work, and was often unsuccessful. Typically, up-to-date contact information was sought through the use of Internet search engines, social networking websites, telephone directories, the electoral roll, institutional gatekeepers, and professional and personal contacts. I was able to find the contact details of 14 individuals. As this number was relatively low, I deemed it reasonable to attempt to interview every individual on this list. This resulted in a total of 9 interviews (see Table 3.4).⁶

Interviewee	Institution	Years	Mode	Fieldwork Stage
1	National Physical Laboratory	1987-1997	Face-to-face	First
2	National Physical Laboratory	1978-2000	Skype video	First
3	National Physical Laboratory	1970-1992	Email	First
4	Royal Holloway	1984-2000	Telephone	First and Second
5	Royal Holloway	1989-2000	Email	First and Second
6	Racal	1981-1987	Telephone	First
7	Racal	1987-2000	Telephone	First
8	Cambridge	1992-2000	Telephone	First and Second
9	Various (inc. FIPR)	1998-2000	Face-to-face	Second

TABLE 3.4: Interviewee Demographics

Once contact information had been found, either a letter or an email requesting an interview was sent. The initial contact also contained details about the nature of the research.⁷ Importantly, in the initial letter, potential interviewees were offered a choice of interview mode. It was made clear that if participants agreed to be interviewed, they would have a choice of a: face-to-face; telephone; Skype video; or email interview. Given that many potential interviewees had spent their lives working as computer scientists, offering a choice of Skype video and email was considered appropriate and unproblematic. Although offering a choice of mode goes against a convention in social research that tends to favour face-to-face interviewing, I felt that it was important to be as flexible as possible given the small

⁶The information contained in the ‘Years’ column within Table 3.4 is limited by the case study’s temporal bounding. Therefore, the dates given do not indicate how long each interviewee worked for a particular institution, but rather the years spent working on cryptology between 1970 and 2000.

⁷A copy of the material sent to potential interviewees can be found in Appendix A.

number of potential interviewees, and because being too insistent on a particular mode might have adversely affected rapport. Though it is difficult to know for sure whether this strategy paid off, it is perhaps telling that the majority of those contacted agreed to be interviewed, and that most of the participants elected to conduct their interview over the telephone, with smaller numbers agreeing to meet face-to-face, communicate over Skype video, or via email.

There exists a generally held view that, though telephone interviews are typically cheaper and more convenient than face-to-face interviews, face-to-face should be seen as the ‘gold standard’. However, some recent comparative studies have claimed that there is little difference between telephone and face-to-face data (Irvine 2011, Stephens 2007, Cachia & Millward 2011). Amanda Holt (2010) has even argued that, in some cases, telephone interviews may be preferable because they can: remove misleading ethnographic data; force the interviewee to articulate everything verbally; and remove the basis for a power imbalance between interviewer and interviewee. The last of these points was particularly relevant, given that I wanted to avoid interviews that resembled a teacher-student encounter.

Although the majority of interviews were conducted over the telephone, a small number were carried out using Skype video. Skype is a piece of downloadable computer software that allows the user to conduct video calls with other users. Skype is an emerging research tool that combines some of the facets of face-to-face and telephone interviews (Hanna 2012). For my Skype interviews, visual information was observed but was not recorded. As with my face-to-face interviews, Skype interviews provided an additional layer of data in the form of body language and other visual information. However, it should be noted that Skype video interviews typically provide less visual information than face-to-face due to the constraints imposed by the fixed viewing window (itself determined by the interviewees computer setup) and the frequent stuttering and pixelated video feed caused by insufficient bandwidth.

Whereas telephone, face-to-face and Skype video interviews were conducted ‘synchronously’, email interviews were conducted ‘asynchronously’. Email interviewing is increasingly being recognised as a legitimate interview mode in the social sciences (Meho 2006, Burns 2010), and its use was considered appropriate in this case given the nature of the interviewees. When an email interview was agreed upon, an initial list of questions was sent to the interviewee, and then the interviewee replied with their responses. In one instance, a participant offered to answer

some initial questions by email, but agreed to a telephone interview if they felt that lengthier responses were required. Email interviews were easy to arrange, free from problems of computer literacy due to the nature of the participants, better suited to busier participants as they were able to reply at their convenience, and on the whole, appeared to be a good source of precise information. In terms of drawbacks, though interviewees responded to every question that was asked rather than cherry-picking preferred questions, email interviews made prompting and probing much more difficult, and the collection of visual data impossible.

All interviews - apart from those conducted over email - were recorded using a Dictaphone. After the interviews had been completed, they were transcribed ahead of the analysis. In the case of the face-to-face, telephone and Skype recordings, this was done 'manually', by listening to the recording and typing up what was said. In the case of the email interviews, no transcribing was necessary, as they already existed in written form. Analysing the interview data had to be done in such a way that it complemented the documentary data. Basic historical information, such as names and dates, was obtained by reading through the transcripts, and where possible, checking against other sources due to the potential for error caused by memory.

As with the documentary analysis, for the interview data, salient information was identified using thematic analysis. As Bryman pointed out, although attempts have been made to develop specific thematic analysis techniques, it is not an approach that "has an identifiable heritage or one that has been outlined in terms of a distinctive cluster of techniques" (Bryman 2008, p.554). Therefore, thematic analysis is used here to refer to the process of reading through interview transcripts iteratively to identify dominant trends and ideas. In this way, the analysis process was similar to coding. However, it should be noted that coding is often a much more formal process best suited to multiple interviews on a fixed topic. Data from face-to-face, telephone and Skype video interviews was largely treated in the same way. However, the email interviews were fundamentally different from the other modes used because they offered written data instead of spoken data. As a result, a higher degree of reliability was attached to any short pieces of factual information, such as names and dates, because the interviewee had been afforded the time to check any information of which they may have been uncertain.

3.3.6 A Note on Ethics

As with all forms of social research, ethical considerations were important in deciding upon a methodological approach. All interviews were carried out after receiving informed consent from the interviewee. In the social sciences it is generally considered good ethical practice to anonymize or de-identify interview transcripts before quoting from them, so as to ensure confidentiality. However, this can create a tension in work that has a historical dimension, as it seems to undermine its central purpose of providing detailed information about individuals from the past. In line with Harry Collins' (2014*b*) code of practice for interviews, I have decided to anonymize any quotations used. As Collins argued, though insiders may be able to guess the source of a quote, this rarely matters because they are by definition already aware of the interviewee's position in a debate. Though I made no promises about confidentiality in the consent form issued to all interviewees, I decided later that, given the aims of the research, there was little to be gained from potentially embarrassing or otherwise harming those that had been willing to contribute. Where quotes are not anonymized, they have been taken from referenced documents. Given that all documents analysed were, in theory at least, publicly available, it did not seem necessary to anonymize or de-identify these quotes, as the content was already a matter of public record. Furthermore, given that many of the documents used are available online, it would be straightforward for someone who wished to know the source of a quote to find it using a search engine.

3.4 Conclusion

Much published STS research does not provide the reader with a detailed description of the processes used to reach its stated conclusions. As a result, many of the basic assumptions that underpin STS research are left unquestioned, leading some to draw attention to a lack of reflexivity. Therefore, I sought to ground my methods in established concepts from the social sciences, and have provided justifications of the use of the case study and historical approaches, with respect to the aims of my research.

It was decided early on that qualitative research and social constructivism would be well suited to the investigation of the nature of cryptology research practices,

the nature of the expertise they produced, and how that expertise was used during the crypto wars. Though case studies have traditionally been viewed as consisting of a trade-off between depth and generalizability, this view is often based on idealized conceptions of contrasting methods, and somewhat dated ideas about the potential for social research to provide universal theories. Of more importance is the way in which a case study is bounded. Because 'boundedness' is difficult to measure, it is the responsibility of the researcher to provide a justification of the decisions they made when isolating their study. I have argued that, given the salience of national context and what's currently known about the emergence of open cryptology research, restricting the technical phase of the case study to developments in UK cryptology research from 1970 and 2000, and to the political phase to the first series of debates in the crypto wars, was both sensible and feasible.

Given the temporal bounding of the case, historical methods - such as documentary analysis and retrospective semi-structured interviewing - were required to gather data. Despite some criticism of the use of historical methods to inform sociological conclusions, it should be acknowledged that history and sociology do share some of the same goals and methods, and have been successfully blended by STS researchers in the past. More importantly, though ethnographic methods such as participant observation can be an excellent way of examining certain observable, micro-social laboratory practices - such as those associated with experimental work - documentary analysis and retrospective semi-structured interviewing appear to be better suited to capturing unobservable macro-social institutional and organizational practices. Given that these practices may also have had an important role in the production of expertise during the technical phase, and historical methods seem suited to their examination, it was decided that they would be used to underpin the answers to the research questions.

Turning to the arguments related to the data gathering strategies, recent thinking has led to a more sophisticated appreciation of what documents are able to reveal about the past. Given that documents reveal and sustain documentary realities, documents pertaining to cryptology research sites and controversies were not analysed in the hope that they would offer a clean description of the past. Rather, they were analysed as part of the expertise producing practices that were of interest. Therefore, the thematic analysis of documents allowed for inferences to be made about the nature of the wider processes of which they were a part. Data

from documents was complemented with data from a series of retrospective semi-structured interviews. The primary advantage of using retrospective interviews in conjunction with documentary analysis was that they allowed me to both reveal new information about the past, and to confirm or refute my prior interpretations. However, it should be acknowledged that the experiences of the interviewee in the intervening period can shape the way in which they subsequently make sense of the past. As a result, as with documents, retrospective interviews were not seen as texts able to provide precise descriptions of past practices. Interviews were instead analysed thematically to gain an insight into how interviewees made sense of their past contributions to cryptology and how they related them to the crypto wars.

In terms of the practical steps taken during the research, the range of analysable documents was extended through the use of Freedom of Information requests and the Wayback Machine. In order to maximize the amount of data available during the retrospective interview process, I decided to adopt a flexible and informal approach. This was deemed important because of the small number of potential interviewees, and because it was believed that this would help establish a rapport. Generating a good rapport with interviewees also offered a way of potentially minimizing the number of responses that resembled official lines, as well as the number of interviews based around a teacher-student dynamic.

The successes and failures of this methodological approach will be discussed in more detail in chapter 10. The next four chapters will be devoted to the descriptions of cryptology research practices that these methods were able to produce.

Chapter 4

Cryptology Research at the National Physical Laboratory

4.1 Introduction

The next four chapters will describe what resulted from the first stage of my research design, as outlined in the previous chapter (see Table 3.2). The purpose of this stage was to examine the technical phase of the crypto wars. After carrying out some preliminary investigations into the history of cryptology research in the UK between 1970 and 2000, it became clear that this was not something that could be found in existing historical or sociological accounts of the period. I would have to assemble the history myself. After carrying out some preliminary historical research, it became clear that collective cryptology research was only carried out at a handful of research sites from the 1970s onwards. Though it wasn't unheard of for individuals - usually within computing or mathematics departments - to carry out their own independent research into cryptology, cryptology research groups were rare.¹

Between 1970 and 2000, cryptology research groups could be found within the following five institutions: the National Physical Laboratory; the University of London (in particular at Royal Holloway college); Racal Electronics plc; the University

¹The emphasis on group or collective practices is important because of the distinction made between somatic limit tacit knowledge and collective tacit knowledge, and the fact that the latter underpins specialist expertise within the periodic table. It is only through collective activity that distinct specialist expertise can emerge.

of Cambridge; and the Government Communications Headquarters (GCHQ).² The next four chapters will describe how cryptology research developed at each of these institutions.³

In this chapter, I will examine the cryptology research carried out by the Data Security Group at the National Physical Laboratory (NPL). I will begin by providing a brief history of the laboratory, from its founding at the beginning of the 1900s to the start of its research into cryptology in the late 1970s. I will then provide more detailed information about the computing work that the laboratory undertook, paying particular attention to pioneering contributions from Donald Davies. I will then examine the social and political context in which cryptology work was carried out, focussing attention on the influence of ‘New Public Management’, and the attempts to implement a customer-contractor principle. I will show that, although the aim of the government’s 1971 Rothschild report was to essentially commercialize the work of Government Research Establishments, within the Data Security Group at NPL, it resulted in a move away from basic cryptology research, and towards work designed to produce standards and to provide consulting services for industries that relied on secure electronic transactions. In the late 1980s and early 1990s, as the ‘marketization’ of NPL intensified, the work of the Data Security Group moved away from the production and development of standards, and towards the testing of technologies for their conformance to standards, and the accreditation of testing facilities.

4.2 Historical Overview of NPL

The National Physical Laboratory - based in Teddington on the outskirts of southwest London - is the UK’s measurement standards laboratory. NPL is a Government Research Establishment (GRE) - a government-owned institution that carries out scientific research on behalf of the state. From its official opening in 1902 (it was founded in 1900), to the present day, NPL has received the majority of its funding from government sources. Therefore, historians have argued that ideas

²This is not an exhaustive list. There is evidence to suggest that collective cryptology research was carried out at (at least) one other site. This particular limitation will be discussed in more depth in the chapter 10.

³The reasons why five research sites will be describe across four chapters will become clear later on. Put simply, the cryptology research carried out at Racal will be described in the chapter on cryptology research at Royal Holloway because of the way in which their work overlapped, and the strong links between the two groups.

about the kind of state-funded research, if any, GREs should perform have ultimately shaped the course of its development. Consequently, the history of NPL can be understood within the context of changing political attitudes to the funding of science and technology, and its relationship to industry (Pyatt 1983, Magnello 2000).

NPL was formally opened by the Prince of Wales (later King George V) in March 1902. In his address, the Prince stated that:

In the National Physical Laboratory we have the first instance of the State taking part in scientific research. The object of the scheme is, I understand, to bring scientific knowledge to bear practically upon our everyday industrial and commercial life, to break down the barrier between theory and practice, to effect a union between science and commerce (Pyatt 1983, p.32).

Though today this might sound relatively innocuous, the idea that industry could benefit from a program of state-funded scientific research was a somewhat incongruous, controversial and potentially problematic idea given the *laissez-faire* approach that was the hallmark of Victorian economic policy (Moseley 1978).

When the laboratory first became operational, there were just two divisions: Engineering; and Physics. The Engineering Division was primarily concerned with testing new technologies, whereas the Physics Division carried out research into electrotechnics, metrology, optics, and chemistry (Pyatt 1983, p.43). In the years leading up to the First World War, typical early work included a study of the effect of wind on bridges and roofs, the standardization of shafts and holes in industrial machinery, and the stress testing of iron and steel specimens (Pyatt 1983, pp.40-45). As the war unfolded, it became apparent that NPL could be used to perform research that was geared towards the war effort. In an attempt to develop the UK's scientific capabilities, the government created the Department of Scientific and Industrial Research (DSIR) in 1914, and control of NPL was passed to it in 1918 (Edgerton 2006, p.116). At around this time, science funding was increased dramatically, largely catalysed by the severing of trade links with Germany. Research into areas such as the calibration of gauges for the manufacture of fuses and shells was increased (Pyatt 1983, p.65), as was research into aeronautics and aerodynamics (Magnello 2000, p.61).

Following the First World War, the UK had, in general, acquired a much more scientific outlook. The attitudes of those opposed to the funding of scientific research by the state had softened. In line with a general trend towards rearmament (Edgerton 2011, pp.28-46), NPL maintained much of its military-themed research throughout the interwar period. Although this was part of a general trend, it was particularly prevalent at NPL because many senior staff maintained strong links with the armed forces. With the outbreak Second World War, NPL was once again able to substitute its remaining non-military research with research that was directed towards the war effort. By 1941 - in line with many other GREs (Edgerton 2011, pp.233-271) - all divisions within NPL were once again working on military projects. Perhaps the most well known contribution that NPL made during the Second World War came when it provided the testing ground for an early prototype of the famous bouncing bomb - later immortalised in the film *The Dam Busters* (Pyatt 1983, p.138).

Following the Second World War, with the immediate priority of armed conflict largely absent, questions about the appropriateness of state-funded scientific research came back to the fore. When control of NPL passed from DSIR to the newly created Ministry of Technology (MinTech) in 1965, the emphasis again returned to testing, calibration and metrology. During this period, research at the NPL was mostly funded through a block grant system. A fixed amount of government money would be given to the laboratory, and decisions about how to spend it would be made internally by the laboratory's senior management. This system was therefore aligned with what Michael Polanyi (1962) referred to as the 'republic of science', in that scientists at NPL - in common with scientists working in many other contexts - were able to pursue their own interests whilst still working as part of a collaborative network.

4.2.1 Early Computing Research

It was in the post-war environment that computing research at NPL began. It started with the establishment of the Mathematics Division in 1945, and soon after, the appointment of J. R. Womersley as Superintendent (Yates 1997, p.11). Although originally called the Mathematics Division, and later the Autonomics Division, its role was to examine computational methods and to investigate the possibility of providing computing services to government. The establishment of

the division was hardly surprising given that the testing and calibration work NPL undertook necessitated accurate and fast numerical calculations, and management were keenly aware of the vital role that the wartime computing service had played, given their links with the armed services.

Before work on cryptology and data security began at NPL, two major computing projects were completed. These projects resulted in two technological artefacts well known in the history of computing: the Pilot ACE; and packet switching. Although these were by no means the only computing technologies produced at NPL, they are the ones that have thus far been given the most historical attention, and are also indicative of the nature of the work that the division undertook under the block grant system. They are also noteworthy because of their pioneering nature. This was undoubtedly a confident period for the laboratory. The prevailing political attitudes, together with the post-war economic boom - both reflected in the block grant system - meant that the laboratory felt it could proceed with this kind of innovative, ground-breaking research (Yates 1997, p.200). One of the first projects that the new division undertook was to design and build their own programmable computer - the ACE (Automatic Computing Engine). The project - which was initially led by Alan Turing - foundered slightly in the early stages, leading to both Turing's departure and the scaling down of the project (Davies 1990, 1993). Nonetheless, the Pilot ACE (as it came to be known) ran its first automated calculation in 1950 and was completed in 1951, making it one of only three programmable computers in the UK at the time (Campbell-Kelly 1981).

The second major project resulted in the development of packet switching. When Donald Davies became superintendent of the Division of Computer Science in 1966, he initiated a new programme of computer network research. As Donald Davies is a key individual in the history of computing research at NPL - especially, as we shall see, related to cryptology - it is worth pausing to provide some biographical details. Donald Watts Davies was born in the Rhonnda Valley, Wales in 1924. He entered Imperial College London to study Physics at the age of 19. After graduating with a first in 1943, he spent the remainder of the Second World War working on tube alloys at Birmingham University.⁴ After the war, he returned to Imperial College and took another first in mathematics, winning the Lubbock Memorial Prize in the process. Davies joined NPL in 1947 and worked on the ACE

⁴Though seemingly implausible at first glance, Davies' birth and graduation dates were confirmed during personal correspondence with Martin Campbell-Kelly.

project with Turing, and later on data communications and packet switching. He was head of the computing divisions within the NPL from 1966 to 1979, when he stepped down to head a small data security research group. Davies retired in 1984 aged 60. A year earlier he was awarded a CBE, followed by the von Neumann medal in 1986, before being elected as a fellow of the Royal Society in 1987 (Campbell-Kelly 2008). Though Davies' contributions to the Pilot ACE and packet switching are now well known, his work on cryptology and data security in the final third of his career has been neglected by historians.

One of the new areas of investigation that Davies pioneered was data communications. After an earlier visit to the US in 1965, Davies had experienced first-hand the issues that those working on early computer networks were facing - in particular, the inherent conflict between time-sharing and real-time communication (Campbell-Kelly 1988). In June 1966, Davies produced a report that proposed a solution to these problems and included it in an unofficial proposal for a UK national data communications network. Shortly afterwards, Davies learned of an almost identical proposal produced by the American computer scientist Paul Baran of the RAND corporation. It was this proposal that described the technology that would underpin the proposed ARPANET, which would eventually come to be known as the Internet (Abbate 1999). In contrast, the nationwide network that Davies proposed was never built. It did, however, come to be implemented on a much smaller scale at NPL, and as such, became one of the first local area networks (LANs) in the country (Yates 1997, pp.131-135). In coming to grips with the nature of networked communications, Davies was among the first to glimpse the potential security issues that networked computers raised.

4.2.2 New Public Management

In order to understand the nature of the data security and cryptology research that was undertaken at NPL, we first need to understand something about the broader attitudes to the funding of GREs in the preceding period. The way in which NPL was governed began to change in the late 1960s and early 1970s. These changes can be understood within the wider context of what is referred to as 'New Public Management' (Hood 1991). New Public Management is a term that has been used to refer to the noticeable changes in public administration that occurred in the UK

and other developed countries in the second half of the twentieth century. New Public Management refers loosely to:

A commitment to downsizing the state, cost-cutting, marketization and competition, the devolution of executive functions to quasi-autonomous agencies and a commitment to customer-contractor and other quasi-commercial policy-making and management principles (Boden et al. 1998).

These commitments, it is claimed, had two broad consequences for the public sector. Firstly, the public sector became less distinct from the private sector in terms of personnel, methods and reward structures. Secondly, the discretionary powers of management were replaced with general rules of procedure (Dunleavy & Hood 1994). These consequences can be observed in the recent histories of various GREs (Boden et al. 2004). The changes can be seen as the result of various high-level government policies that attempted to modify the ways in which GREs operated.

The first such policy was detailed in the 1971 Government white paper on research and development, known as the Rothschild Report (1971). The ideas that the report espoused were to impact upon the activities of NPL. Broadly, the report promoted the view that it was unacceptable for state-funded research to overlap with industrial research and development. It attempted to eliminate this overlap by abandoning the block grant funding system and replacing it with the ‘customer-contractor’ principle.

After winning the 1970 general election, the Conservative party - under the leadership of Edward Heath - set about reviewing the functions of government departments (Gummett 1980, Wilkie 1991). MinTech had been established by Harold Wilson’s Labour government in 1964, and although its success is a matter for debate, the Conservative election manifesto of 1970 pledged to reform it. A review process was undertaken by the government, and this ultimately led to a report by Victor Rothschild, 3rd Baron Rothschild, head of the newly created Central Policy Review Staff - now often thought of as the original think tank (see Blackstone & Plowden 1988). The report stated that each government department with a significant link to science and technology should appoint a Chief Scientist to advise them on matters of science policy. The Chief Scientist would be supported

by Requirements Boards who would assist the relevant government department - the Department of Trade and Industry (DTI) in the case of NPL - by providing the technical knowledge required to make far-sighted policy decisions. The most crucial part of the report stated that “applied R&D, that is R&D with a practical application as its objective, must be done on a customer-contractor basis. The Customer says what he wants; the contractor does it (if he can); and the customer pays” (Rothschild 1971). This recommendation became instantly controversial because it not only applied to work done in GREs - such as the work of NPL - but also to the rest of the work done under the research council system. The report further recommended that 25% of the work funded by the research councils should be on applied R&D, and should consequently be subject to the customer-contractor principle. That government ministers, with the Requirements Boards acting as a proxy, would effectively be commissioning R&D was seen by many as a direct attack on the autonomy of science, and a violation of the Haldane Principle (Wilkie 1991).⁵

Although this perceived ‘attack’ on the autonomy of the research councils triggered much of the outcry, the imposition of the customer-contractor principle does not seem to have been welcomed by the senior staff at NPL either.⁶ The internal response to the report from the Superintendents at NPL was almost universally defensive and critical. In summarising the view of NPL for the benefit of the director, one senior employee wrote:

This report can be criticised on many counts. It is dogmatic on arguable points. It ignores large areas of government science. It is devoted to means rather than ends. It is based on a fundamental misconception regarding the different categories of scientific research. Despite its trenchant style it is vague and self-contradictory in many places and ignores the financial problems that would be thrown up by the adoption of its proposals (National Physical Laboratory 1971).

⁵The Haldane Principle, as it is commonly understood, refers to the idea that decisions about what to spend research funds on should be determined by scientists and not by politicians (Gummett 1980, p.25). However, as Gummett and others have noted, the original Haldane Report did not refer directly to a Haldane Principle. It acquired this name at a later date. As a result, its meaning has changed over time and has often been shaped by debates over science funding (Edgerton 2009).

⁶This came in the form of numerous articles in *Nature* and *New Scientist*, 45 published letters, four editorials in *The Times*, and a lengthy debate in the House of Lords featuring contributions from many distinguished scientists (Gummett 1980, p.198).

More specifically, one of the recurring criticisms was that the report failed to grasp the nuanced and diffuse way in which the production of standards can be beneficial, and how the customer-contractor principle would be unable to account for the broad customer-base that standards serve. As another senior figure at NPL stated, “the bulk of the work of the NPL is in the field of standards. It is therefore a matter of concern that the report should make no reference to it, nor provide a category of work into which it fits” (National Physical Laboratory 1971).

It has been argued that - despite the trenchant style of the Rothschild Report, and the negative reaction that it received - “significant changes were not achieved, and for the most part laboratories continued as close adjuncts of their departments, with funds flowing in a way which was not closely monitored, and with the laboratory defining much of the work to be done” (Boden et al. 1998, pp.271-272). Although the block grant system was replaced, staff at NPL were able to come to agreements with the DTI about what work should be contracted to them. As one former NPL scientist explained to me:

Respondent: The customer is not a very informed customer. They weren’t back then, and they’re not now. They don’t really understand the work that’s done. But let’s face it, NPL is working at the forefront in physics, maths and computing. Well, it was then. Not any more. So, basically, you would have to be working in the field in order to understand it. So, yes, it was driven by a Requirements Board, but we used to tell them . . . we used to write their requirements for them.

The report does, however, appear to have prompted a more general change within the Division of Computer Science. In his annual report to the review committee, itself set up to help NPL deal with the Requirements Board system, the director of the NPL - J. V. Dunworth - stressed that the primary role of the laboratory was the production of standards (National Physical Laboratory 1973). However, in the same report, Dunworth also acknowledged that the Division of Computer Science (and also the Maritime Division) had, in the past, been concerned with research of a more pioneering and academic nature, and did not currently perform any standards work at all. From 1973 onwards, judging from the pressure from the Requirements Boards that appears in the minutes of various high-level committees, the nature of the computing work at NPL began to change. Innovative research into computing hardware and data communications was scaled down, and

standards work was initiated. Despite the fact that the Rothschild Report had not covered the production of standards, the Requirements Boards were happy for this work to continue. Indeed, the annual overviews of divisional work seem to indicate a general trend across NPL to consolidate its standards and consulting role. A separate Standards Committee was established and there were (ultimately failed) attempts to turn NPL into the British Bureau of Standards (Department of Trade and Industry 1976).

The work of the Division of Computer Science was eventually brought into line with the rest of the laboratory and this general trend. In 1974, Dr (later Sir) Ieuan Maddock, the Chief Scientist at the DTI, drew attention to the small number of receipts for contractor work related to computing (National Physical Laboratory 1974). This became a consistent refrain through to 1980, and was regularly raised at Review Committee meetings by the Chief Scientist of the day. Although official correspondence between NPL and the DTI was largely courteous and professional, it is clear from the available documents that the Chief Scientist saw the computing divisions as anomalous and potentially problematic. That the laboratory was struggling to satisfy the Requirements Boards in terms of computing-related contractor work no doubt further influenced the turn to standards and consulting work.

4.3 The Data Security Group

By 1978, work on protocol and security standards was seen as part of the future of the computing work of NPL (National Physical Laboratory 1978*a,b*). It was also in 1978 that Davies, who was, according to David Yates (1997, p.123), dissatisfied by the increased level of administrative work that the Requirements Board system generated, stepped down as Superintendent of the Division of Numerical Analysis and Computer Science to head a newly formed Data Security Group. By 1980, the group had begun research on cryptology and data security standards.

4.3.1 Cryptology Standards

Though Donald Davies was appointed head of the Data Security Group in 1978, his interest in data security was much older. In a 1986 interview with the historian of computing Martin Campbell-Kelly, Davies claimed to have had the “usual childhood interest in cryptography” (Campbell-Kelly 1986). His first serious cryptology work came in 1963, when Midland Bank asked him to test a proprietary cipher designed for use in early ATMs. Davies commented that the ciphers the bank produced were very basic, and that he was able to break them almost immediately. Davies first began working on cryptology standards in 1972, when the NBS were soliciting proposals for what was to become the DES. Davies submitted a basic proposal. This, along with all others, was rejected by the NBS. When the time came for the second round, Davies chose not to submit an improved version, and, as was mentioned in the introductory chapter, the NBS eventually selected Feistel’s proposal. In addition to these early forays into cryptology, Davies’ notes - now held at the Imperial College London Archives - suggest that he was keenly aware of the developments that had taken place in the US. Shortly after the development of Public-Key Cryptography, Davies employed a mathematician in the Data Security Group to attempt to ‘break’ the RSA algorithm. Furthermore, Davies and his team at NPL chaired one of the first public seminars in the UK on the protection of data by cryptography in September 1977. The National Computing Centre (NCC) sponsored the seminar, and it was attended by around 100 delegates. The majority of these delegates came from industry, and included representatives from Barclays Bank, General Motors and IBM (National Computing Centre 1977).

From these beginnings, the Data Security Group used research into cryptology to build a body of expertise in the field. As this expertise matured, it was increasingly focussed on the use of cryptography to secure financial transactions. During the 1980s, the group produced a textbook on data security and electronic funds transfer and a series of annotated bibliographies, as well as more practical research into key management and DES (e.g. Bell & Olding 1978, Price 1979, 1980, Davies & Price 1980*a,b*, Price 1982, 1983, Davies & Price 1984). Furthermore, members of the group acted as chairmen for various national and international data security standards committees.

One of the first projects that the Data Security Group undertook was the development of the Message Authenticator Algorithm (MAA). The MAA was designed to be a digital signature standard. Work began on the MAA in 1981, and details of the MAA were published in 1983 (Davies & Clayden 1983). The MAA became part of ISO Banking Standard 8731-2 in 1987. The MAA - an example of Message Authentication Code - is a standard for providing a message with a digital signature. This allows the recipient of an electronic message to be sure that the sender of the message is who they claim to be, and can also be used by a single user to determine if their files have been altered (Schneier 1996, p.455). Davies and Clayden described the workings of the algorithm as follows:

An ‘authenticator’ is a number which is sent with a message so that a check can be made by the receiver of the message that it has not been altered since it left the sender. For authenticators in general the sender and receiver share the knowledge of a key K which is otherwise secret. If M is the message, the authenticator is a function of K and M . It is calculated by the sender and again by the receiver. If the receiver’s calculated value equals the authenticator value received with the message, the message is assumed to be correct. When a well-designed authenticator is used, giving a 32 bit result, the probability that a message alteration will not be detected is 2^{-32} , which is small enough for most purposes (Davies & Clayden 1983).

In other words, if the authenticator number that the recipient receives is a function of the key known to both parties, then it can be confidently assumed that the sender is who they claim to be, and that there have been no changes to the message since it was sent.

The first document to detail the MAA was produced in 1983 (Davies & Clayden 1983). This document does not contain any reference to the potential for the MAA to be used in the field of banking. Nor does it make reference to any other potential practical application. Only in subsequent documents is banking directly referred to. In 1988, a revised document detailing the MAA was published (Davies & Clayden 1988). This document referred to “financial messages” rather than just “messages”, and also stated that “the algorithm attracted the attention of the Committee of the London Clearing Banks and then Technical Committee 68 (Banking) of the International Standards Organisation, which adopted it as one

of the approved algorithms for message authentication” (Davies & Clayden 1988, p.1). Although believed to be in wide use through to the mid 1990s, the MAA was then suspected to contain security vulnerabilities due to its age and relative simplicity (Schneier 1996, p.457). These doubts over the security of the MAA were later confirmed when, in 1997, a team of cryptographers demonstrated a series of potential attacks that could compromise the algorithm (Preneel et al. 1997).

In addition to work on standards like the MAA, the Data Security Group also provided consultancy services. An example of this kind of service was the Tokens and Transactions Control Consortium (TTCC). The consortium was concerned with machine-readable cards and their uses. The TTCC was established in 1982, and was disbanded in 1988 and replaced with the Advanced Tokens Technology Club. The TTCC was essentially a club made up of firms that were willing to pay a subscription fee in return for being kept up-to-date on current developments and advised about the future of the field. At first, the TTCC acted as an information hub for subscribing firms. They would be provided with expertise after explaining their needs and requirements. This did eventually lead to some experimental work and product development, including the development of the prototype NPL Intelligent Token - a piece of technology similar to a ‘smart card’ system that could be used to verify the identity of the token holder. Although similar to a smart card, the NPL Intelligent Token used public-key cryptography to authenticate the user using digital signatures.

The MAA and the TTCC were typical of the early cryptology research of the Data Security Group, and indeed much of the other work of the laboratory, in that expertise was built through the production of commercially useful standards, technologies, and advice. However, there were further considerations specific to cryptology that led the Data Security Group down this particular path. As one former member of the Data Security Group described to me in an interview:

Respondent: [The Data Security Group] was definitely set up to be working in the commercial environment. Donald Davies and Wynn Price had to negotiate with GCHQ at Cheltenham. They had deals. It was not military security at all. That’s why it was interesting. It was definitely set up to do commercial security, across networks in a commercial environment, i.e. business and stuff. He wasn’t interested in

the military side at all. In fact, there would have been severe problems if they had been.

As this quote illustrates, the niche that the Data Security Group came to occupy was fashioned by the commercial priorities expressed in the Rothschild Report, but also the space left vacant by GCHQ's control of military cryptology. As the same interviewee explained further:

Respondent: [Donald Davies and Wynne Price] certainly moved at a different level in terms of having to be very careful in terms of GCHQ all the time. That was the greatest fear, that GCHQ would tell us to stop doing work. Obviously, they can't directly do it but they could make life very difficult for NPL to continue to do the work.

Ensuring that the research of the Data Security Group did not encroach on GCHQ's work, then, appears to have been a pressing concern, and one that also shaped the direction of the group's research.

4.3.2 Later Governance of NPL

From the mid-1980s onwards, the political context within which the work of the Data Security Group was undertaken changed further. Margaret Thatcher's Conservative Government came to power in 1979. Under the Thatcher Government, the speed and intensity with which New Public Management initiatives were pursued was increased. At the same time, the public sector was downsized and public expenditure on science and technology was reduced. The Requirements Boards system was eventually dissolved in 1988. In the same year, the Prime Minister's Efficiency Unit published a report entitled 'Improving Management in Government: The Next Steps'. This report launched what came to be known as the Next Steps Initiative. The initiative aimed to deliver more efficient and cost effective public services through the breaking up of the single Whitehall unit into what it called 'executive agencies'. This initiative was pursued enthusiastically within the DTI, meaning that a number of GREs, including NPL, were seen as strong candidates for agency status (Boden et al. 1998, p.272). NPL subsequently became an agency of the DTI in 1990. This belief in 'agencification' chimed with a more

general feeling about what kind of work GREs should be pursuing. In 1988, the Chief Scientific Advisor for the Cabinet Office, Sir John Fairclough, argued that:

Spending on R&D should be directed to work which was far from the development of a marketable product or process. 'Near market' R&D was to be left to industry, with government expenditure confined to areas where the market would fail to operate to produce maximum benefits to the economy as a whole (Boden et al. 1998, p.272).

As a result, as far as GREs were concerned, scientific work that benefitted the economy, and was simultaneously unlikely to emerge directly from industry, was to be favoured over the production of marketable technologies.

A further series of reports were published in the early 1990s that aimed to deal directly with the role of government laboratories (Boden et al. 2001). The first of these was the 'Review of Allocation and Use of Government Expenditure on Science and Technology' (known as the Levene-Stewart Review after the Prime Minister's Advisor on Efficiency - Sir Peter Levene, and the Chief Scientific Advisor - Sir William Stewart). The review argued that, despite previous efforts, in general, the relationships that existed between GREs and their parent departments did not much resemble those that existed in the private sector. The review proposed that, in the case of government laboratories, procurement should be separated from ownership, and that full privatization should be seriously considered. Further reviews and government exercises followed. Ultimately, direct privatization of government laboratories was favoured over a more gradual introduction of policies to encourage market-liberalization. The end result was that by 1996, a number of government laboratories, including the National Engineering Laboratory and the Laboratory of the Government Chemist, were fully privatized.

Instead of being fully privatized, the NPL became a Government-Owned Contractor Operated (GoCo) agency. Under the GoCo arrangement, the government retained ownership of NPL, but contracted a private company to operate it. This contract was won by Serco plc in 1995. Serco established NPL Ltd as a subsidiary for the purposes of managing NPL. GoCo status undoubtedly entails complex arrangements, but in simple terms:

The contract between the outsourcing firm and the government specifies the work that the government will buy from the business and

arrangements for sharing efficiency gains. This quite substantial organisational reform (in terms of the shape, nature and character of the organisation) is effectively only a privatisation of the management of the business (Boden et al. 2006, p.135).

As such, GoCo was the weakest form of privatisation used during this process of GRE reform. In the case of NPL, some have deemed the reforms successful, given that researchers have claimed that significant savings have been achieved and more commercial practices have been implemented, without adversely affecting the quality of the scientific work produced (Whelan 2000, Wallard 2001).⁷

The reform of GREs also resulted in broad changes to working practices. Boden et al. found an increased emphasis on the customer-contractor principal in this period. Control over the research to be carried out was transferred to the customer, which in the case of NPL, was the DTI:

In the agencies and privatized laboratories alike, the introduction of customer-contractor relationships has almost ubiquitously been described to us as the reform which has had the biggest single change on work and organization. For privatized firms the impact of the introduction of customer-contractor relationships, which preceded any change in ownership, was greater than the change experienced on privatization itself. For many laboratories funding used to come either as a vote directly to the establishment or from the Whitehall parent but following a specification largely drawn up by the laboratory itself. For agencies and privatized laboratories alike funds now rest with the Whitehall customer, who may well now specify the work (or contract a third party to do so) and may use market testing mechanisms too (Boden et al. 1998, p.286).

In general this served to bring a sharper focus to the work of government laboratories, and in the process, phased out “hobby projects” and “Friday afternoon experiments” (Boden et al. 1998, p.286). More specifically, “the ‘business’ approach to the delivery of science and technology from the government research

⁷It should be noted that this view is far from universally shared. It was clear from the interviews I carried out that some - at least within the Data Security Group - viewed these reforms very negatively indeed.

laboratories has involved a shift towards greater emphasis on technology transfer activities, and away from the basic functions which universities are now able to pursue and capitalise upon” (Boden et al. 2001, p.95). As a result, differences between the research of GREs and universities, in general, became more pronounced.

4.3.3 Testing and Accreditation

The consequences of these changes are visible in the research of the Data Security Group. Between 1988 and 1990 - the period between the dissolution of the Requirements Boards system and the switch to agency status - the research of the Data Security Group continued along the same lines as before. Between 1990 and 1995, the research of the group came to be dominated by testing activities. Starting in April 1990, the Data Security Group embarked on a two-year programme of work under the heading of ‘Standards and Conformance Testing in Data Security’ (Data Security Group 1993). By this point, all of NPL’s research was being commissioned on a customer-contractor basis. In the case of this research programme, the customer was the Information Technology Division (ITD) of the DTI. The two-year programme was divided into five sections:

1. Supporting standards in data security;
2. Research and development in conformance testing methods;
3. Technology transfer;
4. Support for the DTI’s CCSC, managed by NPL since April 1990;⁸
5. Technical support for DTI’s ITD (Data Security Group 1991*b*, p.1).

Most of the work carried out under this programme was devoted to research and design in conformance testing methods. The group developed a series of techniques under the heading of ‘Strict Conformance Testing’. The purpose of Strict Conformance Testing was to determine how well a particular piece of technology conformed to a particular standard. Strict Conformance Testing was defined as “The testing and analysis of an implementation of an IT security standard to ensure that:”

⁸‘CCSC’ stands for Commercial Computer Security Centre. This body will be described later on in this section.

1. It implements the mandatory requirements of the standard in a correct manner;
2. It implements only those optional requirements of the standard stated as being supported in the conformance statement, and they are implemented in a correct manner, and;
3. It contains no functionality which would either be prejudicial to the correct operation of the implementation, or would cause a possible breach of security (Data Security Group 1991*b*, p.1).

The Data Security Group also investigated the use of formal methods for testing standards. Formal methods use quasi-mathematical language to logically define processes within software for the purpose of clearly specifying what it is supposed to do. The group used the MAA as a basis for assessing which formal methods would be of most use for the testing of security standards. In this case, the Vienna Development Method was judged the most useful (Data Security Group 1991*a*).

Also in April 1990, the Commercial Computer Security Centre (CCSC) was moved to NPL. The CCSC was established by the DTI in 1987, and was based at the Royal Signals and Radar Establishment at Malvern. The purpose of the CCSC was to act as a technical focus for industry on IT security issues and to stimulate the development of approaches to IT security evaluation (Data Security Group 1991*b*, p.3). Whilst based at Malvern, the CCSC produced the so-called Green Books - a seven-volume collection of criteria and codes of practice concerning the security evaluation of technologies. When the CCSC moved to the Data Security Group, the focus was shifted to harmonising these criteria with those developed separately in other countries. This effort contributed to the formation of the Information Technology Security Evaluation Criteria (ITSEC) - a set of criteria that were used in several European countries, including France, Germany, the Netherlands and the UK. The CCSC played a key role in the formation of ITSEC: "Amongst other things, the CCSC contributed the 'claims language', a system of structured natural language statements used to express the claimed security functionality of a product or system unambiguously" (Data Security Group 1991*b*, p.3). The CCSC also worked on the Information Technology Security Evaluation Manual (ITSEM). ITSEM, which was published in September 1993, specifies the methodology to be used when carrying out ITSEC evaluations.

In May 1991, the UK Information Technology Security Evaluation and Certification Scheme was established. Set up jointly by the DTI and CESG, the purpose of “the Scheme” was to put in place a mechanism for carrying out evaluations of technologies in line with the ITSEC. This led to the creation of a UK Certification Body based at GCHQ in Cheltenham. The Certification Body - a group of around 20 representatives from CESG and the DTI including a small number of Data Security Group employees - managed the running of the Scheme. Actual evaluations of technologies were carried out by Commercial Licensed Evaluation Facilities (CLEFs), which were in turn accredited by the National Measurement Accreditation Service (NAMAS). Those working within the Data Security Group were fully trained NAMAS assessors, whose job it was to take part in the accreditation process for laboratories that wished to become CLEFs (Data Security Group 1991*b*, p.4). The task of the CCSC was deemed complete in 1993, and was closed down. After the establishment of the ITSEC, it was decided that a further harmonisation effort should be undertaken to align the European and North American criteria. This led to the Common Criteria for Information Technology Security Evaluation (‘CC’ or ‘Common Criteria’).

When NPL came to be managed by Serco, research practices within the Data Security Group were focused on producing cryptology expertise relevant to testing and accreditation. However, faced with increased financial pressure, Serco began to look for ways to cut overhead costs. Computing research within NPL was scaled back, and research groups were reorganized. The Data Security Group was formally closed, and cryptology research expertise was transferred to a new ‘Techniques for High Integrity Section’. However, data security and cryptology research, along with other many forms of computing research, were gradually phased out. As one interviewee explained to me:

Respondent: Serco didn’t understand the work we did, and didn’t see the value of it, and didn’t think it should be inside NPL. They knew that the DTI were cutting their funding and making it harder, so I think they wanted to focus on the core work of NPL.

On this understanding, increased efforts to commercialize the work of NPL can be seen as responsible for the shrinking of their cryptology expertise just as the political phase of the crypto wars was getting underway.

4.4 Conclusion

From the 1970s to the start of the crypto wars in the mid-1990s, the research of NPL was increasingly commercialized. This resulted in the replacement of the block-grant system of funding with a system based on the customer-contractor principle. The increasingly strict adoption of the customer-contractor principle shaped the research practices of the Data Security Group. Although the Data Security Group initially carried out basic cryptology standards research, such as the development of the MAA, as the pressure to commercialize intensified, the group increasingly acted as industrial consultants. By the 1990s, it was seen as inappropriate for the group to be developing near-market products, and so instead focussed on work that filled then gaps left by industry. Under a contract from the DTI, the group developed techniques for testing cryptology standards produced by others, and provided the expertise required to establish national and international testing and accreditation frameworks. By the mid-1990s, the practices of the Data Security Group had produced a body of contributory expertise related to the testing of cryptology standards and technologies, that operated within a highly bureaucratized network of small groups and initiatives. It could even be argued that the emphasis on New Public Management served to prohibit the acquisition of contributory expertise related to other aspects of cryptology, and that eventually, it led to a shrinking of the cryptology expertise produced at NPL.

Chapter 5

Cryptology Research at Royal Holloway

5.1 Introduction

Royal Holloway College, University of London, has been an important site for cryptology research in the UK from the mid-1980s to the present. This chapter will describe how research into cryptology at Royal Holloway emerged out of the reorganization of the University of London in the mid 1980s, but also out of collaborations with Racal Electronics plc and other industrial partners. I will begin by examining the circumstances that led to the changes to the organization of the University of London in the 1960s. I will then describe the parallel emergence of Racal and their shift to producing electronic communications devices that utilized cryptography. Then, I will describe how Royal Holloway and Racal came to collaborate with one another. Finally, I argue that the expertise produced by the Information Security Group at Royal Holloway emerged, in part, from practices that were designed for mathematical research, and in part from practices that were designed for industrial collaboration.

5.2 Historical Overview of the University of London

Before describing the cryptology research carried out at Royal Holloway College, it is necessary to situate the college within the context of the University of London. Royal Holloway has been a constituent college of the University of London since 1900. The University of London - which was founded in 1836 - operates under a federal system. Today, it consists of 18 colleges - of which Royal Holloway is one - 10 research institutes, and a number of other centralized bodies. However, these brief details mask a complexity and uniqueness that “arises from its size, its federal structure, its metropolitan role, and above all, the course of its historical development” (Harte 1986, p.10). The structure of the University of London, and the relationships between its various constituent bodies, is complex. As a result, the description that follows will provide only a very brief overview of the issues relevant to the development of cryptology research at Royal Holloway.

The University of London was originally founded in 1836 to provide a federal structure that linked University College London and King’s College London - two new universities that were established in 1826 and 1829 respectively. More educational institutions based in London and the surrounding area were subsequently incorporated into the University at various points. The relationship between the institutions that make up the University and the University’s central organization has varied throughout its history. As F. M. L. Thompson has explained, although for much of the University’s early history it did little more than co-ordinate examinations across the colleges, it still exerted a certain influence:

Since from 1836 to 1900 the University was purely an examining and degree-awarding body the main initiatives in mapping out new branches of knowledge necessarily happened in the separate teaching colleges. Nevertheless, the University controlled the examination system, and the examination syllabus was an important instrument for translating new knowledge into formal qualifications, and these in turn exerted a strong influence on the ways in which the colleges arranged their teaching (Thompson 1990, pp.x-xi).

Unsurprisingly, this relationship has caused disagreements between the colleges and the University over what should be researched and taught. Thompson has argued that the responses to this tension have gone through three phases:

1. 1836-1900: The University, in the shape of a Government-nominated senate, prescribed through its examination syllabuses the content of the college's degrees;
2. 1900-1966: A single University-wide degree for each subject that was managed by Boards of Studies, themselves made up of teachers from the various colleges;
3. 1966-: Centralised University control of degrees has been largely abandoned, with teachers within colleges designed their own degrees.

In the above scheme, the most relevant period for this study is the one that began in 1966 and continues to the present day. As a result, the periods 1836-1900 and 1900-1966 will not be discussed further.

The start of the third phase in the above scheme was, in part, a result of the broader changes to the UK university system that occurred in the 1960s, following the Robbins Report (1963). In 1963, Lord Robbins published his report of the Committee on Higher Education. The Robbins Report recommended the immediate expansion of the university system, and more specifically, that "courses of higher education should be available for all those who are qualified by ability and attainment to pursue them and who wish to do so" (Robbins 1963, p.8). Robbins' conclusions were largely accepted and implemented by Harold Wilson's Labour government.

The Robbins Report had direct implications for the University of London. Although the report recognized that the University's federal structure could offer advantages, it also argued that the system created problems and inconveniences. This led to a process of critical examination, and eventually, reform (Harte 1986, p.262). In 1970, an inquiry chaired by Lord Murray of Newhaven was launched into the organization of the University of London. The Murray Committee's report was in favour of continuing with the federal system, but also asked whether it would be advantageous if the colleges were to merge into half-a-dozen larger

institutions that could then evolve into separate universities (Murray 1972). Although the conclusions of the Murray Report were initially rejected, they set the tone for the subsequent debates about the future of the University. Faced with increased financial pressures - felt by many other universities in the 1970s and 1980s - the decision was eventually taken to merge some of the University's constituent colleges.

Though the federal system remains in place, today, the colleges of the University of London are so distinct that most are considered to be separate universities in their own right. For example, the Research Assessment Exercise (RAE), an official government assessment of the quality of research produced by university departments, separately assesses and ranks each college. Furthermore, those working and studying at each college rarely consider themselves to be a part of a single 'University of London'. As Thompson described:

It is only on grand ceremonial occasions, the graduations days and the conferments of honorary degrees, that the University is made visibly aware of its own existence; and the number actively participating in these rituals cannot be more than a tiny fraction of the whole body of teachers and students. For most of them their world of learning is bounded by their individual College or Institute, and the University remains a remote, unknown, nebulous and vaguely threatening entity, little more than the source of red tape, mountains of largely incomprehensible paper, and unpleasant financial decrees (Thompson 1990, p.ix).

As such, it makes little sense to study the practices that influenced cryptology research at the university level, given the absence of practices that are shared across colleges. Instead, the focus will be placed on individual colleges, and the departments within them, as they were more intimately involved in shaping the research.

5.3 Historical Overview of Racal

In order to understand the nature of cryptology research practices at Royal Holloway, their relationship with the UK electronics company Racal must also be

examined. This subsection will provide a brief historical introduction to Racal, paying particular attention to their work on data communications and communications security. Racal Electronics Ltd was established by Raymond Brown and George Calder Cunningham in 1950. The name ‘Racal’ came from the ‘Ra’ and ‘Cal’ in each of the founders’ names. Brown and Cunningham - who had both previously worked for the British electronics firm Plessey - initially created Racal with the aim of manufacturing and selling high-frequency communications equipment. In 1966, Raymond Brown left Racal to become Head of Defence Sales at the Ministry of Defence, and Ernest Harrison - who had joined the company as an accountant - was appointed as chairman (Wilson 1980). Harrison became synonymous with Racal, and remained as chairman until the company was sold to French electronics and defence contractor Thompson-CSF (which changed its name to the Thales Group shortly afterwards) in 2000. Under Harrison, Racal became something of a success story. At one point it was the third largest electronics firm in the UK. At its peak, it operated in 110 countries, and employed over 30,000 people.

For much of its history, Racal operated under a franchise model. By 1990, the Racal Group was made up of around 150 autonomous companies - such as Racal-Engineering and Racal-Telecom - that each specialised in designing and building certain types of product. However, the vast majority of these companies did not carry out their own cryptology research, and as a result, will not be discussed further. Though not a name that many are now familiar with, Racal is perhaps best remembered for spawning the Vodafone telecommunications company, which at the time of writing, is the second largest mobile telecommunications company in the world. Despite Racal’s size and significance in the second half of the twentieth century, there does not presently exist an historical description of its activities.

5.3.1 Early Years

Racal started out by manufacturing military radio equipment. In the immediate years following its creation, the company struggled to win orders, and would manufacture almost anything in order to remain operational (Jansen 1990). Racal achieved a breakthrough in the mid-1950s when it received a contract from the Royal Navy to build a variant of the American Collins Model 51-J Radio Receiver. The result was the Racal RA17 - a highly successful piece of technology that

set Racal on the path to becoming a major supplier of military communications equipment.

By the late 1960s, Racal had begun to enter the data communications industry. In 1969, Racal formed a partnership with the Milgo Electronic Corporation - an American firm that had recently developed a series of high-speed modems. This partnership created a new company - Racal-Milgo - and helped Racal establish itself in the US. At this time, data communications was a burgeoning field, and as a result, it was the fastest growing area for Racal throughout the 1970s (Wilson 1980). In 1973, Racal established an Advanced Development Division. The purpose of the division - which was initially led by Keith Thrower - was to investigate and probe new areas of research in order to aid product development within other divisions. The Advanced Development Division was in communication with almost all of the companies within the Racal Group, but much of its early work was in the field of communications security (Racal 1975).

5.3.2 Development of Communications Security Products

It was through the research of the Advanced Development Division that research into communications security at Racal began. Early research was deemed successful, and Racal-Datacom was formed in April 1974 to continue it (Racal 1975). As the *Racal Review* - the internal journal for the Racal Group - explained:

Racal-Datacom set up its activity in temporary premises in Salisbury with a small nucleus of Racal personnel transferred from various companies with the group . . . In November 1974 new premises, including laboratory, production area and offices were completed and quickly in operation. First deliveries of production units were made in December 1974 and a new product has been introduced every two months from that date (Racal 1975).

The products that Racal-Datacom produced were originally marketed under the Racal-Mobical brand, as it had been in use since 1966. By 1975, products were being released under the Racal-Datacom name instead. By 1977, Racal-Datacom - which would later change its name to Racal-Comsec (referring to the communications security features that they provided) - had developed a number of products.

These products were designed to offer either secure speech communication, or, secure text communication. The products were typically aimed at militaries and police forces, rather than civilians. The product range included the MA 4014B - a device that used a technique known as time division scrambling. The MA 4014B processed speech into discrete time segments that were then rapidly shuffled in a pseudo-random manner at over 140 times a minute. Over 600,000 code keys were selectable, each of which would be programmed by up to 64,000 codes by means of front panel switches (Racal 1977). Racal also sold devices that offered secure text communication. The MA 4210 was a small device, slightly larger than a pocket calculator, that could convert plaintext messages into cryptograms, and vice versa, using a keyboard and a dot matrix display. It used two interconnected pseudo-random binary generators to produce a complex non-linear key stream of 170,000,000,000 characters (Racal 1977). Importantly, although Racal typically claimed as part of their promotional material that these early products were able to withstand “sophisticated computer-backed cryptanalysis”, the security that they offered did not have any strong mathematical underpinnings. They realised that this would have to change if they were to continue to make similar claims and develop their products further.

5.3.3 Enlisting Cryptology Expertise

In 1978, Racal-Comsec decided to take a more mathematical approach to designing their communications security products. However, they decided that they did not at that time possess the required mathematical expertise. They decided to address this deficiency by drawing on the expertise held within the University of London. Racal-Comsec’s search for cryptology expertise occurred at a time when it was starting to become available outside of military and intelligence gathering organizations. As has already been described in chapter 1, in the US Whitfield Diffie and Martin Hellman (1976) had already published their paper outlining the principles of public-key cryptography. This was followed by Rivest, Shamir and Adleman’s (1978) paper that outlined how public-key cryptography could be achieved in practice. In the UK, research into cryptography at NPL had just formally begun with the establishment of Donald Davies’ Data Security Group. Furthermore, a small number of individuals - such as John Gordon at Hatfield Polytechnic, and R. F. Churchhouse at Cardiff University - had also begun to carry out some early cryptology research. Racal-Comsec, however, decided to

approach Fred Piper - a mathematician based at Westfield College, University of London.

As Fred Piper is an important figure in the history of UK cryptology research, it is worth pausing to consider his career. Fred Piper received an undergraduate degree in mathematics from Imperial College London in 1962. This was followed by a PhD in mathematics from the same college in 1964. Piper then worked briefly as a lecturer at Royal Holloway College before transferring to Westfield College in 1969. There, he was promoted to Reader in 1971, and Professor in 1975. Piper is now informally considered the ‘father’ of the UK’s cryptology research community. This is perhaps reflected in the fact that, together with Royal Holloway colleague Sean Murphy, Piper wrote the cryptography edition of the popular *A Very Short Introduction* series (Piper & Murphy 2002).

When Racal-Comsec approached Piper, he had not yet carried out any cryptology research. Piper had specialized in combinatorics - a branch of mathematics that deals with finite structures such as graphs and sets. Though Piper had not carried out research into cryptology when approached by Racal-Comsec, he was well placed to understand the mathematics associated with the techniques that they were hoping to use. When Racal-Comsec approached Piper, their devices were based on the use of stream ciphers - an example of symmetric cryptography where each plaintext digit is encrypted based on a digit within a random stream of numbers. Racal-Comsec were using physical testing to test the efficacy of their random number generators, which was ultimately unsuited to the task. Placing the random number generators on a mathematical basis offered a cleaner way of understanding the strength of the security that their products were able to provide.

Racal-Comsec wanted to employ their own mathematician. Piper suggested Henry Beker - an ex-PhD student of his who was then working in the Mathematics Department at Swansea University. Beker was appointed by Racal-Comsec to be their first Head of Mathematics. The four individuals that held this post after Beker also completed PhDs that were supervised by Piper, such was the strength of the relationship between him and Racal-Comsec. Beker and Piper worked together on cryptology for the next few years. The results of some this work were described in *Cipher Systems* (Beker & Piper 1982) and *Secure Speech Communications* (Beker & Piper 1985). As is acknowledged in the books’ prefaces, the research that they described was driven by Racal-Comsec’s need for mathematics to underpin the cryptography used in their products.

5.4 Historical Overview of Royal Holloway

After being introduced to cryptology through Racal-Comsec, Piper would go on to be instrumental in establishing a programme of cryptology research at Royal Holloway. In order to understand how this came about, attention must be briefly re-focused on the University of London. The University of London underwent a period of major re-organization during the 1980s. This re-organization process played a major role in the establishment of cryptology expertise at Royal Holloway. This section describes how this came about.

Royal Holloway College was established in 1879 by Thomas Holloway. Holloway was a Victorian entrepreneur who had amassed a fortune from selling patented medicine. After considering the best way to spend his wealth, Holloway decided to establish a women-only college. The result was Royal Holloway College, based in Egham, Surrey. Royal Holloway became a constituent college of the University of London in 1900. The college remained women-only until 1945, when male postgraduate students were accepted. Following the Robbins Report, male undergraduate students were accepted in order to meet the demands of expansion. This was the most significant change in a series of developments designed to move the College away from the Victorian traditions that had previously sustained it (Bingham 1987).

As with many other small universities and colleges, though Royal Holloway made progress towards expansion in the ten years that followed the Robbins Report, the inflation of the early 1970s left the college in dire financial straits. This made further expansion difficult. It prompted the college to consider ways to generate more money. For the departments within the Faculty of Science, this resulted in attempts to form working relationships with industry. In 1978, the Dean of Science recommended:

1. That academic departments should be actively encouraged to organise training courses and to seek liaison with industry where appropriate;
2. That commercially viable rates should be charged for any training courses or services provided by departments;

3. That the larger part of any surplus (i.e. after covering overheads) accruing from such commercial ventures (say 80%) should be credited to the Departments concerned, to provide an independent revenue to assist their teaching and research budgets. The remainder should be credited to the college tuition account (Royal Holloway 1978).

By the late 1970s - following the recommendations of the Murray Report - it was becoming clear that the Royal Holloway would have to seriously consider merging with another college of the University of London. If mergers were to take place, it was clear that Royal Holloway wanted to use them as an opportunity to increase the amount of scientific research based at their current site. It was clear from the documents contained in their archives that it was believed that increasing the amount of scientific research performed by the college would improve its financial situation. In 1981, with college mergers looking increasingly likely, the Principal of Royal Holloway - Lionel Butler - wrote to the Vice-Chancellor of the University of London - Randolph Quirk - to make the case for a 'science consortium' based at Royal Holloway, which would absorb the scientific research currently being undertaken at Bedford College and Westfield College (Butler 1981). The following year, it was agreed that Bedford College would merge with Royal Holloway College to become Royal Holloway and Bedford New College, and would be based at the existing Royal Holloway site in Surrey. This merger took place at the same time as a number of other mergers within the University. The most significant of these came when Westfield College merged with Queen Mary College, to be based at the existing Queen Mary site in East London.

The college mergers of the 1980s also had direct consequences for teaching and research. In order to maximize the resources at the University's disposal, it was decided that certain colleges should focus on particular fields. Rather than it being spread across all colleges, attempts were made to shift scientific research to five colleges: University College; Imperial College (which formally separated from the University of London in 2007); King's College; Queen Mary; and Royal Holloway (Harte 1986, p.284).

Although the science consortium described by Butler did not materialize exactly as he had envisioned, some scientific work was transferred there. Academics based at other colleges of the University were asked to transfer to Royal Holloway - particularly those working on scientific and mathematical research. Despite the

merger, the financial situation faced by the college in the 1980s was still described by the Principal as being “extremely serious” (Royal Holloway 1986*a*). In the mid to late 1980s, the college, particularly within the Faculty of Science, continued to search for new ways to obtain grants from industry (Royal Holloway 1986*c*). Additionally, the college also planned to recruit more overseas students, expand the academic profile to include more vocational subjects, provide short courses and consultancy to industry, and to embark on a zero-based budgeting exercise. It was in this context that the cryptology research at Royal Holloway began.

5.5 The Information Security Group

Fred Piper was asked by the University of London to move back to Royal Holloway College from Westfield College in order to lead research into discrete mathematics and combinatorics. He officially transferred in August 1984. At the same time, Donald Davies - who had previously established a cryptology research program at NPL - joined Royal Holloway’s department of Computer Science and Statistics as a visiting professor (Royal Holloway 1984). Upon arriving at Royal Holloway, the head of the mathematics department there - M. R. C. McDowell - asked Piper to focus on building expertise in cryptology. McDowell had read Beker and Piper’s first book on cipher systems, and saw cryptology as an expanding field, and a potentially useful niche for Royal Holloway to occupy.

5.5.1 Formation

The first steps towards an Information Security Group at Royal Holloway began in November 1986, when a proposal for an academic initiative in ‘Data Security’ was submitted to the Academic Board (Royal Holloway 1986*b*). The documents that accompanied the proposal are worth examining in detail, as they offer a good indication of how the initiative was pitched to the university. They stated that:

The security of stored and transmitted data, and the prevention of unauthorized access to software and other facilities, is the subject of massive international endeavour. The field is rich in technological and

mathematical challenge, and is growing rapidly as an academic specialism, strongly linked with industrial advancement in information technology (Royal Holloway 1986*b*).

It was believed that a gap in the market had been identified related to data security:

British industry has a growing need for recruitment of graduates in data communications and computer systems, with specialist skills in data security. Despite the fact that all the major financial institutions and a number of large corporations now have a senior person responsible for data and computer security, the British Universities do not yet include an adequate centre of excellence satisfying the requirement for data security specialists (Royal Holloway 1986*b*).

Cryptology research was to be central to work done under the heading of data security. Specific objectives of research would be to develop and implement:

1. New, improved key management schemes;
2. Encryption algorithms;
3. Message authentication codes;
4. Public key systems (Royal Holloway 1986*b*).

The importance of collaborating with industry for the the purposes of wealth generation were also emphasised:

The proposed academic initiative would lead to practical implementations. It is expected that the results would be industrially valuable as well as being academically excellent. It is important that this could be a wealth-generating area of university/industry collaboration (Royal Holloway 1986*b*).

It was also stressed that Royal Holloway was “already collaborating with industry, specifically Racal, Hewlett-Packard, ICL and Ferranti, in data security” (Royal Holloway 1986*b*). Specific details were also given of proposed project with Racal:

Racal approached [Royal Holloway and Bedford New College] with a request to use the college's computing facilities as a testbed/showpiece for their network security protocols and equipments. If, as seems likely, this goes ahead, this would constitute further recognition of existing expertise in data security at RHBNC. This collaborative venture would be strongly consistent with an academic initiative proposal to develop a centre of excellence in data security (Royal Holloway 1986*b*).

The academic initiative proposal was approved, and around £40,000 was allocated annually to cover the cost of equipment, and the appointment of a Reader and a Research Associate. This can perhaps be considered the start of the Information Security Group at Royal Holloway. However, it should be noted that the group had a very informal structure through until the early-1990s.

5.5.2 Early Cryptology Research and Industrial Collaboration

Industrial partnerships were central to the early research of the group. When asked to summarize the work of the Information Security Group, one interviewee stated that:

Respondent: [Industrial collaboration was] absolutely crucial. Couldn't be over emphasised enough. Basically, everything I did in cryptography and information security was centred around industry collaboration.

Elaborating further:

Respondent: I set up a consultancy company because if we were really going to understand information security and even cryptography, we had to understand how it was used and why. The last thing industry wanted was academics telling industry what they thought industry should be doing. So, everything was focused around industrial collaboration. That's how I got into cryptography in the first place ... We started by writing to all of the local people saying "Look, we do combinatorial mathematics, if there anything that we can do that's

useful for you?”. That’s how cryptography and coding theory came and jumped up and became something that industry wanted.

The early research of the group was three-fold. It revolved around: block design; theoretical coding; and cryptology. Much of this work was funded by the Science and Engineering Research Council (SERC) (known since 1994 as the Engineering and Physical Sciences Research Council (EPSRC)). From 1988 onwards, the group undertook a series of SERC-funded cryptology projects, on topics such as ‘Encryption Algorithms’ (GR/E64640/01), ‘Stream Ciphers’ (GR/H23719/01) and ‘Digital Signatures and Hash Functions’ (GR/K51259/01) (Engineering and Physical Sciences Research Council 2013). Some SERC-funded projects were undertaken in collaboration with industry, such as UK DTI/EPSRC research project entitled ‘Security Studies for Third Generation Telecommunications Systems’ (GR/J17173/01). This £160,000 project was carried out in conjunction with Vodafone (which had recently emerged from the Racal group) and GEC Plessey Telecommunications (GPT).

Of the early work on cryptology, one of the most notable developments was Sean Murphy’s work on differential cryptanalysis. Murphy - who’d joined Royal Holloway in 1988 shortly after completing a PhD in mathematics at the University of Bath - carried out postdoctoral research on the analysis of block ciphers. This led to work on differential cryptanalysis. As was explained to me in an interview:

Respondent: Differential cryptanalysis is based on encrypting a pair of plaintexts with a specified difference, so is classified as a chosen-plaintext attack. In some circumstances, analysis of the corresponding pair of ciphertexts can give some small information about the key. The technique of differential cryptanalysis is essentially concerned with analyzing many such pairs of ciphertexts to determine the key. When I started working at Royal Holloway as a postdoctoral researcher, I started by looking at FEAL, a block cipher published by the Japanese company NTT. This led to the development of the ideas of differential cryptanalysis and the 1990 paper analyzing FEAL, though others were working with similar ideas at this time.

Though now a familiar concept, Murphy's analysis of the FEAL block cipher came to be recognised as one of the founding contributions to the practice of differential cryptanalysis in the open literature (Coppersmith 1994).

In 1990, Chris Mitchell - another former PhD student of Piper's and former Head of Mathematics at Racal-Comsec - joined the Information Security Group after leaving a post at Hewlett-Packard. The Information Security Group have subsequently claimed that this move was significant because it led to the formation of an annual Hewlett-Packard-funded colloquium on information security. As well as cementing the relationship between Hewlett-Packard and the Information Security Group, it marked a change in the nature of the work of the group from mathematical cryptology to more general computer and information security (Information Security Group 2008).

5.5.3 Teaching

In 1991, the Information Security Group submitted a proposal for the introduction of an MSc course in Information Security (Royal Holloway 1991). The course would cover cryptology, but also computer security, network security, and security management. In 1987, a number of companies had approached the Information Security Group about establishing an MSc in cryptology. However, the group decided that this would be too narrow, so instead they delayed the introduction of a new course in order to recruit those required to offer a course that covered a wider range of fields. The MSc in Information Security that resulted was a vocational course that aimed to produce information security professionals capable of managing the security requirements of large organizations. In its first year in 1992, the MSc had a total of ten students. At the height of the dotcom era in 2000, student numbers peaked at around 250. Though the early collaborations with industry had prompted the group to build mathematical expertise, industry collaboration now increasingly meant training students to be able to go out and provide the information security expertise that many other non-technical industries required.

Though the demand for cryptography expertise had grown, the Information Security Group still collaborated with companies within the specialist cryptology industry. Following the launch of the MSc programme, one of the group's main partners was Zergo. Zergo was founded by Henry Beker - Racal-Comsec's first

Head of Mathematics - and became a leading provider of cryptology solutions during the 1990s. In order to be able to adequately train their staff, Zergo relied on the expertise of the Information Security Group:

[Zergo] introduced a structured Information Security training programme on which members of the ISG lectured. This led in 1994 to the Introduction of the Postgraduate Diploma in Information Security, based on courses offered by Zergo and an MSc level dissertation supervised by Royal Holloway academics (Information Security Group 2008).

The training that the Information Security Group provided also led to them being awarded a Queen's Anniversary Prize for Higher and Further Education in 1998. The group were commended because they had:

... pioneered high-level education in information security and advises both government and industry in one of the most sensitive developing areas of business. The group plays a vital role in training those who work in the field of information security, across industry and commerce as well as in vital security elements of the nation. The courses were the first of their kind in the world and are now seen by many as the benchmark qualification (Royal Anniversary Trust 2013).

Today, the MSc in Information Security is still in operation, and continues to train information security professionals for work in industry. It is recognised as one of the leading postgraduate courses in the field.

5.6 Conclusion

The Information Security Group at Royal Holloway emerged from a combination of the reorganization of the University of London, and a series of industrial partnerships. Following the Robbins Report, constituent colleges of the University of London merged in order to meet the demands for expansion in an unfavourable economic climate. Royal Holloway successfully aimed to fashion itself as a home for scientific research, and was able to capitalize on a niche in the shape of cryptology. In line with what the college expected from departments within the Faculty

of Science, the Information Security Group capitalized on industrial partnerships with companies like Racal. This shaped the research of the group, and led it towards research that would complement the requirements of industry.

The early research of the Information Security Group was abstract and mathematical, because at that time, industry required the mathematical underpinnings for the technologies it was producing. By the early 1990s, through taking steps towards offering postgraduate teaching, the work of the group had expanded beyond cryptology into the broader realm of information security. This brought the group even closer to the needs of a wider range of industries, as they required solutions to specific problems and trained personnel to provide them. Although it would be misleading to argue that the group exclusively engaged in work that followed the contours of the requirements industrial cryptology, the group certainly enacted a body of contributory expertise of this nature - one that, as will become clear in the next chapter, was not necessarily mirrored in other academic contexts.

Chapter 6

Cryptology Research at the University of Cambridge

6.1 Introduction

In this chapter I will consider the cryptology research of the Security Group within the Computing Laboratory at the University of Cambridge. I will begin by briefly describing the work related to computing that was undertaken before the laboratory was founded in the mid-1930s. I will then describe how, following the Second World War, research practices at the laboratory were shaped by the headships of Maurice Wilkes and Roger Needham. Under Wilkes, the laboratory was primarily concerned with large-scale, long-term projects, that typically resulted in the construction of large computing systems. This continued under Needham, but thanks to increases in funding, the laboratory was able to expand in terms of the number of staff it employed, the space it occupied, and the number of research themes it pursued. Security became a key research theme for the laboratory during the 1980s. This led to the formation of an informal Security Group that became particularly concerned with cryptology in the early 1990s. I argue that the focus on systems, and in particular the adoption of interdisciplinary methods, allowed the Security Group to develop expertise on how such systems behaved when put to use in the real world.

6.2 Historical Overview of the Cambridge Computer Laboratory

The University of Cambridge was founded in 1209. It is the third oldest university in the world. Today, it is generally regarded as one of the world's best and most prestigious universities. Furthermore, with an endowment of nearly five billion pounds, it is also one of the wealthiest. Similar to the University of London, the University of Cambridge has a federal structure. It is today comprised of 31 autonomous colleges, and just over 100 departments. One such department is the famous Cambridge Computer Laboratory.

6.2.1 The Roots of Computing Research

The Cambridge Computer Laboratory was established in the 1930s. Before the laboratory was established, individuals associated with the University of Cambridge made some of the most important early advances in the history of computing. Computing at the University of Cambridge can perhaps be thought of as starting with Charles Babbage (1791-1871). Babbage studied mathematics at Trinity College, Cambridge (and later Peterhouse College, Cambridge) as an undergraduate from 1810 to 1814, before nominally returning to Cambridge as the Lucasian Professor of Mathematics from 1828 to 1839 (Hyman 1982). Babbage is now best remembered for his proposals for the construction of two early 'computers' - the Difference Engine and the Analytical Engine - as well as for his work on economics and manufacturing (Schaffer 1994). Babbage also performed some notable work on cryptology. According to one biographer, Babbage was "the outstanding cryptologist of his age", and was "wholly without rival" (Hyman 1982, p.227). David Kahn has described how, though he never published his work on cryptology, Babbage was the first to employ mathematical formulas and notations, and was able to solve both polyalphabetic ciphers and to manage autokeys (Kahn 1997, p.204). Babbage became interested in ciphers as a schoolboy, and maintained this interest throughout his life. Though he remarked in his autobiography that "diciphering is one of the most fascinating of arts, I fear I have wasted upon it more time than it deserves", Ole Immanuel Franksen (1993) has argued that Babbage had intended to write a book entitled *The Philosophy of Deciphering* in 1853. But, like many of Babbage's projects, this book was never completed.

Babbage's work was eventually followed around one hundred years later by the theoretical contributions of Alan Turing (1912-1954). Turing studied mathematics at King's College, Cambridge between 1931 and 1934, before being elected as a fellow of the same college in 1935. In 1936, whilst still at Cambridge, Turing introduced the concept of 'universal machines' in his now famous paper 'On Computable Numbers, with an Application to the Entscheidungsproblem'. This would later become the abstract model on which modern computers are based. Turing also carried out important work related to cryptology. He famously worked on breaking German ciphers at the Government Code and Cypher School (GC&CS) at Bletchley Park during the Second World War. During his time at Bletchley Park, Turing also designed an electromagnetic machine known as the Turing-Welchman bombe - designed to assist with the deciphering of messages encrypted using the German Enigma machine (Hodges 1983).

Though the work of Charles Babbage is sometimes considered to be the first computing work to be carried out at the University of Cambridge, there does not appear to be a link between Babbage and what would eventually become the Cambridge Computer Laboratory. Similarly, though he did come into contact with those involved, Turing was not formally connected to the work that would ultimately lead to the creation of the laboratory. Furthermore, there exists no evidence to suggest that Turing undertook any serious cryptology work whilst at Cambridge. As a result, though any examination of Cambridge computing would, in a sense, be incomplete without consideration of the work of Babbage and Turing, it would be difficult to argue that their work exerted any noticeable influence on the working practices associated with the laboratory or its later work on cryptology.

6.2.2 Formation

The Cambridge Mathematical Laboratory was officially created in 1937. According to Mary Croarken (1992), the 1930s had seen an increase in the amount of computations required for theoretical and applied scientific research. In order to meet this demand, a model differential analyser - based on a similar device constructed at the University of Manchester - was built at Cambridge. The differential analyser came to the attention of Maurice Wilkes - then a researcher with the Cavendish Laboratory Radio Group. Wilkes made use of the machine

for his own computations, and eventually became responsible for the machine's maintenance and for providing guidance to other researchers. The enthusiasm for the differential analyser led to formal approval for the creation of a separate mathematical laboratory. A 1936 General Board Report on the Establishment of a Computer Laboratory stated that the intention, firstly, was "to provide a computing service for general use", and secondly, "to be a centre for the development of computational techniques in the University" (Spärck Jones 1999). The work of the new Mathematical Laboratory would be guided by these two aims up until the late 1960s.

The first director of the laboratory was John Lennard-Jones, who was then also the Plummer Professor of Theoretical Chemistry. Shortly after the laboratory was established, the Second World War broke out, and its activities were temporarily suspended. Following the end of the war in 1945, Maurice Wilkes took over as director of the laboratory. He would hold this position until his retirement in 1980. Wilkes was born in Dudley, Staffordshire, in 1913. He read the Mathematical Tripos at St John's College, Cambridge from 1931-1934, before completing a PhD in the propagation of radio waves at the Cavendish Laboratory. During the Second World War, Wilkes served as a radio operator and contributed towards the development of radar. After the war, Wilkes devoted most of his attention to computing research. As head of the Computer Laboratory, he oversaw a number of pioneering computing projects, most notably, the construction of one of the first stored program computers - the Electronic Delay Storage Automatic Calculator (EDSAC). Wilkes was elected as a Fellow of the Royal Society in 1956, and was a founder member and first president of the British Computer Society. He was the second recipient of the Turing Award in 1967, and was knighted in 2000. He died in 2010 aged 97.

In his autobiography, Wilkes (1985) outlined his approach to university research projects. He believed that projects should be pursued either for the purpose of training graduate students, or to satisfy the intellectual interests of a faculty member. Above all, Wilkes believed that university projects should fall into the mainstream of computer science, and thus contribute towards the field as a whole. In practice, this meant that projects should typically be long-term - requiring around ten years to reach maturity - at which point, the impetus could be passed to industry. Wilkes claimed that the decision to commit to a particular project involves an assessment of what the future will be like, and in particular, "the technological,

economic, and sociological forces that will mould it” (Wilkes 1985, p.222-223). For Wilkes, computers were technologies that had to be useful, and they had to be able to function in the real world. Wilkes and his approach dominated the laboratory during his headship, as can be seen in the projects that were undertaken during this period (Ahmed 2013). As will become clear, the practices that the members of laboratory used under Wilkes allowed them to develop considerable expertise in how to build, maintain and operate usable computer systems.

6.2.3 Early Computer Development

The first major project that the laboratory undertook following the end of the Second World War was the construction of the EDSAC. In line with the aims laid out in the aforementioned 1936 report, Wilkes decided that the immediate objective for the laboratory after the war was to establish a usable and reliable computing service in a short timescale. As historians of computing have noted, the objective “was not to build the best possible machine” (Lavington 1980, pp.31-32). In contrast to the approaches used in most of the other early computer projects, Wilkes “decided that he was interested in having a computer, rather than in trying to advance computer engineering technology” (Campbell-Kelly & Aspray 2004, p.90). After learning about the EDVAC project first-hand during a trip to America, Wilkes and his small team began construction of a similar machine. Though there were short-lived attempts at collaboration with the new computing division at NPL, differing views on how to approach the task meant that the two laboratories would go it alone. EDSAC ran its first program in 1949 and provided a usable computing service until 1958. As such, the EDSAC can be thought of as the first stored-program computer. Consequently, Wilkes’ team was amongst the first to experience some of the now-familiar issues associated with computer programming (Campbell-Kelly 1992). As he explained in his memoirs:

By June 1949 people had begun to realize that it was not so easy to get a program right as had at one time appeared. I well remember when this realization came on me with full force. The EDSAC was on the top floor of the building and the tape-punching and editing equipment one floor below on a gallery that ran round the room in which the differential analyser was installed. I was trying to get working my first non-trivial program, which was one for the numerical integration of

Airy's differential equation. It was on one of my journeys between the EDSAC room and the punching equipment that 'hesitating at the angles of stairs' the realization came over me with full force that a good part of the remainder of my life was going to be spent in finding errors in my own programs (Wilkes 1985, p.145).

Once the EDSAC was functioning, a number of researchers from a variety of scientific disciplines used it to perform calculations. As Joyce Wheeler has noted, "the EDSAC was always regarded as a tool for the solution of problems, rather than just an engineering achievement" (Wheeler 1992, p.27). One of the first uses of EDSAC was to find large prime numbers, but it was also used as a tool by various university departments in fields such as astronomy, genetics, crystallography and economics.

With the EDSAC operational, the 1950s saw much discussion within the laboratory about future research. The stored-program concept on which EDSAC was based was thought of as successful, so it was eventually decided that a more powerful version of EDSAC was the next logical step. The successor to EDSAC - EDSAC 2 - became operational in 1958, and remained in use until 1965 (Wilkes 1992). EDSAC 2 was followed by a machine named Titan. Titan was the name of a time-sharing computer built in collaboration with the British computer manufacturer Ferranti. The advent of the time-sharing model is significant because it marks the point at which concerns about what we now think of as 'computer security' become relevant. Prior to time-sharing, computers did not maintain persistent files linked to a particular user, so in a sense, there was nothing to protect. However, under the time-sharing model, there was a need to prevent users from being able to access files that belonged to others. Those working at the laboratory became aware of this issue whilst using Titan. As such, it is notable in the context of computer security as the first computer to incorporate a one-way mathematical function to protect password files (Needham 1992).

6.3 Computer Security and Cryptology Research

Towards the end of the 1960s, the laboratory began to change direction. With the rise of companies like IBM, it no longer seemed feasible for the main focus to be the development of new computers. The aims expressed in the 1936 General

Board report - concerned as they were with building new computers to provide a computing service to the rest of the University - were now out of date. As a result, Wilkes submitted a report to the General Board in 1969 that argued that the focus of the laboratory should shift from the construction of new computers, and that the user service should be separated from the laboratory in the form of a University Computer Service. The General Board approved these recommendations, and the laboratory set about research into other aspects of computing (Ahmed 2013, pp.80-83).

In common with many other computer laboratories at this time - in particular NPL - the Cambridge Computer Laboratory developed an interest in data communications. In 1974, Wilkes had the opportunity to view a digital communication ring built by the Swiss firm Hasler AG. This inspired the construction of the Cambridge Ring - an early electronic communications network that shared computer peripherals. The expertise gained during the development of the Cambridge Ring, together the expertise in security gained during the development of Titan, would create a platform on which to build expertise on networks and data security.

Research into computer security at the laboratory began in the late 1970s, gathered pace during the 1980s, and by the 1990s, was one of the main avenues of research (Spärck Jones 1999). Cryptology was a key part of the laboratory's computer security research. Research into computer security and cryptology at the laboratory began with the work of Roger Needham. Needham was appointed head of the laboratory when Wilkes retired in 1980. Needham completed an undergraduate degree in Mathematics and Philosophy at Cambridge in 1956, and a PhD on information retrieval and automatic classification in 1961 (Rashid 2004). He was elected as a fellow the Royal Society in 1985. Needham decided to leave academia in 1995, before being involved in the establishment of the UK's Microsoft Research Laboratory. He remained at Microsoft until his death from cancer in 2003. Needham was a classic systems computer scientist. He played a key role in many of the aforementioned Cambridge projects, including Titan, the CAP (standing for 'capability-based') computer, and the Cambridge Ring. Alongside his research, Needham also maintained a public service career. Needham was a member of a number of government committees, including the Science Research Council's Computer Science Committee, the University Grants Committee, and the Alvey Committee. Needham was also a member of the Wass Committee - a group charged with the reform of the structure of the University of Cambridge in the 1980s. The

Committee proposed the creation of the new position of Pro-Vice-Chancellor to assist the Vice-Chancellor. Needham was chosen as the first to occupy this post. Throughout his career Needham maintained ties with industry. He consulted for firms such as Xerox and Hitachi. Whereas Wilkes preferred to carry out projects that existed entirely within the laboratory itself, Needham looked to outside influences.¹ Needham always upheld a positive view of the computer industry. In one interview he claimed that “if there wasn’t an industry concerned with making and using computers the subject wouldn’t exist. It’s not like physics - physics was made by God, but computer science was made by man. It’s there because the industry’s there” (Rashid 2004, p.6).

During the first decade of Needham’s headship, he presided over an expanding laboratory. It expanded in terms of the number of staff it employed, the space it occupied, and the number of research themes it pursued. Much of this was down to generous financial support. Needham took over as head of the laboratory at a time when certain aspects of computing research were well funded. In particular, the Alvey Programme - a UK government sponsored research programme that ran from 1983 to 1987 - provided financial support for a number of high-profile computing research projects (Ahmed 2013, pp.104-105). One of the first projects that the laboratory undertook under Needham’s headship was the development of UNIVERSE (UNIV-Expanded Ring and Satellite Experiment). UNIVERSE was a system designed to connect separate local area networks using satellites. The project was jointly funded by the Science Research Council, the DTI, and a small group of commercial partners. As such, it was one of the first projects that the laboratory undertook that received external input, and was therefore typical of Needham’s collaborative approach to research (Ahmed 2013, p.105-106).

6.3.1 Early Cryptology and Computer Security Research

Needham was also concerned with computer security during the 1980s and 1990s. As a result, research into computer security became a research theme at the laboratory during this period. Needham’s interest in security emerged from his earlier work on networks, time-sharing, and capability-based computing. Needham’s

¹The most notable exception to this general rule came when J. Lyons and Co., a British catering company, assisted with the EDSAC project in order to aid the development of their own computer. This resulted in the LEO (Lyons Electronic Office) Computer, now considered the first computer to be used for commercial or business purposes.

two most well known contributions to computer security - and indeed cryptology - were the co-development of the Needham-Schroeder protocol and BAN logic. The Needham-Schroeder Protocol - developed jointly with Xerox PARC-based researcher Michael Schroeder in 1978 - defined a method for using encryption to provide a decentralized authentication system on an insecure network (Needham & Schroeder 1978). Needham-Schroeder would eventually evolve into the widely-used Kerberos protocol, which is now a key part of Microsoft Windows security (Anderson & Bond 2004). Needham's second key contribution was the co-development of BAN logic. BAN logic - also known as Burrow-Abadi-Needham logic - was first described in 1990 (Burrows et al. 1990). BAN logic is essentially a set of rules for logically defining trust in communication systems. It was designed with the aim of clarifying assumptions about who or what in a system is trustworthy. As a result, it became possible to define more clearly the protection that certain protocols offered, particularly as it highlighted the distinction between trustworthy and untrustworthy insiders.

The Needham-Schroeder protocol and BAN logic are milestones in the development of security and cryptology at the laboratory. As well as being interesting technical achievements, they are notable because of the way they demonstrate how Needham conceptualized security. A consideration of real-world security systems, and the role that cryptology played in them, allowed Needham and others to develop an understanding of cryptology as a component in a system.

6.3.2 The Security Group

The expansion of the laboratory that Needham oversaw during the 1980s fed through to the research carried out during the decade that followed. Building on Needham's research in this area, the laboratory established a Security Group during the 1990s. The group started as an informal collection of computer scientists with an interest in computer security, cryptology, and distributed systems. Though the research group was open to anyone with an interest in security - whether based at the laboratory or not - the group had Cambridge researchers at its core. It was an emphasis on tackling real-world problems that came to define their work throughout the 1990s.

During the 1990s, the key figure in terms of cryptology and security at the laboratory - along with Roger Needham - was Ross Anderson. Anderson's work spanned

multiple fields. He published interdisciplinary work on the economics and psychology of information security, peer-to-peer and social network systems, reliability of security systems, cryptology protocols and algorithms, information hiding, and privacy and freedom issues (Anderson 2012). He has also written an influential textbook on security engineering, the second edition of which was published in 2008 (Anderson 2008). Anderson studied Mathematics and Philosophy at Trinity College, Cambridge, graduating in 1979. His first job was in avionics, before working in security in the 1980s when he acted as a consultant for companies that designed cryptography equipment for banks. Anderson returned to Cambridge in 1992, where he undertook a PhD under the supervision of Roger Needham. He remained at the laboratory after the completion of his doctoral research. He was elected as a fellow of the Royal Society in 2009.

During the early 1990s, the Security Group carried out research into cryptographic protocols, cryptographic algorithms, formal methods and steganography. In 1994, Needham also co-developed - along with Cambridge colleague David Wheeler - the Tiny Encryption Algorithm. The algorithm - which was designed to be both small and secure - can be expressed in as little as nine lines of code. It was designed to be small enough to be used in almost any situation, whilst still delivering an acceptable level of security (Wheeler & Needham 1994). Later, the group developed a number of cryptographic protocols and cryptographic primitives (low-level cryptographic algorithms). For example, Anderson co-developed the Tiger hash function and the BEAR and LION block ciphers in 1996, and the Chameleon stream cipher in 1997 (Anderson & Biham 1996*a,b*, Anderson & Manifavas 1997).

The group also collaborated with industry. As the group were keen to emphasise, “much of our best research has been inspired by tackling real problems, and our funding comes from a wide range of sources; we collaborate with commerce and industry both in the UK and overseas” (Computer Security Group 1998*a*). An example of the group’s work of this kind is the NetCard project, the goal of which was to “design protocols to support emerging services in high speed networks”. Funding for the project was provided by the DTI and the EPSRC, and was carried out in collaboration with Energis Communications Ltd and General Information Systems Ltd (Computer Security Group 1998*b*).

6.3.3 Cryptology and Security Systems

In addition to these research themes, the group also began to study the reliability of security systems, electronic commerce, and medical information security. This work is best exemplified in one of Anderson's (1994) most influential early papers on why cryptosystems fail. The use of interdisciplinary research methods is what clearly differentiates this from other cryptology work. The conclusions of this research were based on a survey of known retail banking fraud cases. Through using this method, Anderson was able to show that "the threat model commonly used by cryptosystem designers was wrong" and that "most frauds were not caused by cryptanalysis or other technical attacks, but by implementation errors and management failures" (Anderson 1994, p.32). The survey highlighted that many instances of bank fraud took place with some kind of insider knowledge or access. For example:

In a recent case, a housewife from Hastings, England, had money stolen from her account by a bank clerk who issued an extra card for it. The bank's systems not only failed to prevent this, but also had the feature that whenever a cardholder got a statement from an ATM, the items on it would not subsequently appear on the full statements sent to the account address (Anderson 1994, p.33).

In another example from a bank in Scotland, "a maintenance engineer fitted an ATM with a handheld computer, which recorded customers' PINs and account numbers. He then made up counterfeit cards and looted their accounts" (Anderson 1994, p.33). Anderson observed that high-tech attacks - such as those that directly attempted to break a cryptographic algorithm - were very rare, and even those that did exist tended to exploit the difficulties associated with integrating cryptology into a security system. Thus, it was concluded that:

Designers of cryptographic systems have suffered from a lack of information about how their products fail in practice, as opposed to how they might fail in theory. This lack of feedback has led to a false threat model being accepted. Designers focussed on what could possibly go wrong, rather than on what was likely to; and many of their products are so complex and tricky to use that they are rarely used properly (Anderson 1994, p.39).

This paradigmatic insight formed the basis for much of the work of the laboratory in the years that followed. In 1995, Anderson and Needham (1995) coined the phrase “programming Satan’s computer” to refer to the problems associated with securing a system under the control of an adversary. Given that “the great majority of actual security failures resulted from the opportunistic exploitation of various design and management blunders” and that “one can always check that a protocol does not commit the old familiar sins, but every so often someone comes up with a new and pernicious twist” it becomes difficult, if not impossible, to prove a cryptographic protocol ‘correct’. Consequently, Anderson and Needham proposed what they called an ‘explicitness principle’ - the idea that:

Robust security is about explicitness. A cryptographic protocol should make any necessary naming, typing and freshness information explicit in its messages; designers must also be explicit about their starting assumptions and goals, as well as any algorithmic properties which could be used in an attack (Anderson & Needham 1995, p.439)

Again, the way in which the laboratory researched real-world uses of cryptography was clearly evident. It could also be argued that this work carries with it explicit policy implications. Anderson set his survey of retail bank fraud against the backdrop of the UK’s legal framework, and contrasted it with the situation in other countries:

In some countries (including the USA), the banks have to carry the risks associated with new technology. Following a legal precedent, in which a bank customer’s word that she had not made a withdrawal was found to outweigh the banks’ experts’ word that she must have done, the US Federal Reserve passed regulations which require banks to refund all disputed transactions unless they can prove fraud by the customer. In Britain, the regulators and courts have not yet been so demanding, and despite a parliamentary commission of enquiry which found that the PIN system was insecure, bankers simply deny that their systems are ever at fault. Customers who complain about debits on their accounts for which they were not responsible - so-called ‘phantom withdrawals’ - are told that they are lying, or mistaken, or that they must have been defrauded by their friends or relatives (Anderson 1994, p.33).

Here, then, it is argued that the results of the survey carried out by the group pointed to a potential problem with the UK's legal framework. By researching cryptology as a part of larger systems, the group were able to produce expertise that could map directly onto policy debates.

Throughout the 1990s, it is clear that the research into the reliability of security systems carried out by the laboratory led to the realization that the level of security provided depended on more than just the strength of a particular technology. It also depended on the way in which that technology was positioned within systems that also included human beings. Needham expressed this view most clearly when he delivered the Clifford Paterson Lecture at the Royal Society in 2002. Needham argued that:

Despite all the theoretical progress that has been made, and the very ingenious papers that have been published, systems remain rather insecure. This is not primarily because of bad algorithms or protocols. It is to a substantial extent because of ignoring the human element. An example is non-repudiation, where the purpose of a protocol is to furnish evidence that will convince an arbitrator that a party attempting to repudiate a transaction did in fact commit to it. The arbitrator is, and has to be, human (Needham 2003, p.1550-1551).

He added that security systems are often complex, and “to compound the effects of complexity, humans involved in managing security are fallible, lazy, and uncomprehending”, and ended the lecture with a plea for “computing researchers [to] climb down from their ivory towers to look at the real-world contexts in which their systems will be deployed” (Needham 2003, p.1554-1555).

6.4 Conclusion

The Cambridge Computer Laboratory was established in the 1930s to develop computing techniques and to deliver a computing service to the rest of the University. As such, from the outset, a significant proportion of the work of the laboratory has relied on practices that were geared towards the development, maintenance, and operation of computer systems. These practices allowed those working in the

laboratory to become sensitive to some of the more practical issues associated with systems. As a consequence, those working at the laboratory were among the first to develop technologies - such as the Needham-Schroeder protocol - that dealt with specific security issues. When Roger Needham was appointed head of the laboratory in 1980, thanks to a climate of generous funding, he was able to increase the number of research themes pursued. Research into computer security was one theme that the laboratory developed. As a result, the laboratory was able to build expertise in cryptology. In addition to the expertise required to build individual cryptology technologies, the laboratory also produced a body of expertise on how security systems behave when they are put into use. During the early 1990s, research practices even came to include the use of surveys in order to develop ideas about how cryptosystems fail. This allowed the group to appreciate the role of human beings in cryptosystems, and in the process, allowed ideas from psychology, economics and criminology to begin to influence their cryptology expertise.

Chapter 7

Cryptology Research at the Government Communications Headquarters

7.1 Introduction

In this chapter I will consider the cryptology research of the Government Communications Headquarters (GCHQ). I will briefly describe the cryptology research carried out there from 1919 to 1970, before describing in more detail what's known about their modern cryptology research. I will pay particular attention to the research carried out by CESG¹ - the body within GCHQ that is now officially responsible for providing information security expertise and cryptology solutions to public bodies in the UK. I will discuss in detail the development of 'non-secret encryption' - a series of theoretical ideas developed within CESG that appear to both mirror and pre-date the independent development of public-key cryptography in the US. Like much of the work of GCHQ, when this research was carried out during the 1970s, it was a closely guarded secret. I will describe how - starting in the late 1980s - in line with broader political trends and changing attitudes towards the role of intelligence organizations, CESG took on a more public role, and attempts were made to fashion a culture of openness around it. Despite this,

¹Prior to 2002, the group was referred to as the 'Communications-Electronics Security Group', and was also abbreviated to 'CESG'. Following a name change in 2002, the group's full name became 'CESG', given that it was felt that 'Communications-Electronics Security' no longer fully described their work (Communications-Electronics Security Group 2012).

I will argue that, though thanks to the declassification of certain materials we now have a fragmentary knowledge of some of the theoretical cryptology research carried out within CESG, and also an understanding of some of the public duties they took on during 1990s, it would still be more accurate to characterize their cryptology expertise as secret given the nature of the practices that were used to produce it.

7.2 The UK's Intelligence Organizations

The UK has three main intelligence organizations: MI5; the Secret Intelligence Service (SIS) (commonly referred to as MI6); and the Government Communications Headquarters (GCHQ). Whilst it is often difficult to clearly delineate their respective remits, MI5 is broadly responsible for delivering national intelligence and security, SIS is responsible for delivering foreign intelligence, and GCHQ is responsible for delivering signals intelligence and communications security (Herman 1996). All three now operate under the Joint Intelligence Committee (JIC), which is responsible for providing the UK Cabinet with the intelligence required for governmental decisions.²

MI5 and SIS were founded as the same organization in 1909. GCHQ was founded ten years later in 1919. Throughout their history - particularly in the case of GCHQ - most of their activities, and even their knowledge of their existence, have been kept secret from the public. This is, of course, unsurprising, given that an intimate knowledge of their activities would undermine their objectives. This is even less surprising in light of the culture of secrecy that is seen to pervade the work of the UK's civil service more generally (Rogers 1997, Vincent 1997, Moran 2013). This secrecy, though perhaps rooted in culture, has been continuously upheld through legislation and institutional practices. The most influential of these has been the Official Secrets Act. The Official Secrets Acts 1911 to 1989 have prevented - and still prevent - members of intelligence organizations from publicly disclosing information relating to their work. Christopher Moran, in surveying the changing ways in which secrecy has been maintained by the civil service, described the broad scope of the original act:

²'Intelligence' refers here to information that is of military or political value.

Section 1, commonly known as the ‘spying clause’, made it a criminal offence for anyone, ‘for a purpose that could be prejudicial to the safety or the interests of the state’, to collect, communicate or publish any plan, drawing or other item of official information to an enemy. The accused had no ‘right to silence’ and a trial could be held in camera. Section 2, which was targeted at civil servants, politicians and journalists, made a felony of both the unauthorised communication and the receipt of official information. It was widely drafted, embracing all types of information without any discrimination (Moran 2013, p.23).

When an individual is employed by an intelligence organization, they are required to sign the Official Secrets Act. However, given that it is a law, individuals are bound by it whether they sign or not. The Official Secrets Act therefore criminalizes the unauthorized dissemination of classified information. This includes information relating to scientific research carried out at GCHQ and much information related to working practices more generally.

The Official Secrets Act has also informed a number of practices common to many UK government bodies. For example, the ‘Government Protective Marking Scheme’ was designed to label documents according to the sensitivity of the material they contained. Under the scheme, documents produced by the state were labeled either: Top Secret; Secret; Confidential; Restricted; Protect; or Unclassified. This system was used in conjunction with a vetting procedure that assigned government employees a clearance level commensurate with one of these labels. The clearance assigned to an employee would then be used to determine what documents, and ultimately what information, they would have access to. Employees were prevented from viewing documents labeled at a level above their clearance. As a result, even employees within governmental organizations are subject to practices designed to uphold secrecy. In the case of intelligence agencies, throughout their history, almost all of the records they have produced have been classified, usually at the level of Top Secret or Secret, meaning that individuals must be vetted and assigned a clearance before they can view them.³

³It was announced in October 2013 that the Government Protective Marking Scheme would be replaced with the similar yet simpler Government Security Classifications Policy. The changes came into effect in April 2014 (Cabinet Office 2013).

The thrust of the Official Secrets Act also influenced the way in which records have, or have not, been made public. In most cases, documents produced by public bodies are sent to the National Archives (formerly known as the Public Record Office) when they cease to be of use. Once at the National Archives - located in Kew, South West London - records can be viewed by any member of the public. The Public Record Office was established through the Public Record Office Act 1838. Though initially established for the preservation of legal documents, records from government departments were accepted from the 1840s onwards. However, government departments were under no obligation to submit documents to the Public Record Office, giving them the option of retaining certain records at their discretion (National Archives 2012). Under the Public Records Acts 1958 and 1967, government departments were subject to a thirty-year rule. Under this rule, all records created by public bodies would be transferred to the Public Records Office thirty years after their original date of creation. However, their creators could retain records that were over thirty years old if they were granted an exception by the Lord Chancellor under Section 3(4) of the acts. Additionally, it was decided that all records created by, or referring to, intelligence organizations would always be retained and would never be released. In 1967, the Lord Chancellor, Lord Gardiner, approved this blanket retention policy for another 25 years, thus carrying it through to 1992.

As will be described in full later, almost all of what follows in this chapter is a result of more recent attempts to reform the UK's intelligence organizations, with a view to making them more 'open' and more publicly accountable. Almost this entire chapter is based upon information that would not have been available to researchers or the general public before 1970. Such was the secrecy surrounding GCHQ in particular that, prior to 1970, it is likely that the majority of the general public would have been scarcely aware of its existence. Therefore, it should be remembered that almost all of what's described in this chapter would not have been public knowledge at the time.

7.3 Historical Overview of GCHQ

This section will review what is now known about the history of GCHQ. GCHQ was founded as the Government Code and Cypher School (GC&CS) in 1919, ten years after the founding of MI5 and MI6. The negative attitude towards spying

and espionage held by Victorian society had meant that by 1904, the UK had been without a code-breaking centre for over fifty years (Porter 2009, p.20). The priorities of the First World War meant that the British army began an effort to break encrypted German radio communications. By 1919 it was decided by foreign secretary Lord Curzon that a unified peacetime code breaking organization should be formed. The stated function of the GC&CS was the defence of the communications used by government departments, but in reality, it began attempting to read the communications of others almost immediately. Prior to the Second World War, the activities of GC&CS were mainly based around the decrypting of Russian diplomatic communications, but also those of France, the US, and Japan (Aldrich 2010, p.16).

By 1939, GC&CS had moved to Bletchley Park in Buckinghamshire, where it set about breaking encrypted enemy communications as part of the war effort. By 1941, this effort came to be known under the codename of Ultra - referring to 'Ultra-Secret' - and consisted of attempts to break various German codes, most famously, those generated by the Enigma machine (Kahn 1991, Sebag-Montefiore 2001). There also existed a parallel project to break the Lorenz cipher, which resulted in the construction and use of Colossus - now recognised as one of the world's first electronic computers. These episodes are probably the most well known in GCHQ's history. As will be described later, though highly secret at the time, the wartime cryptology work carried out at Bletchley Park was eventually revealed in the 1970s. The motivation behind keeping this work secret is usually thought of in terms of allowing GCHQ to continue to use the same techniques to break the codes used by others during the Cold War. However, it has recently been argued that the primary purpose of hiding this work was simply to prevent knowledge of the existence of GCHQ (Moran 2013). In either case, since these activities were revealed, cryptology has been understood as a core activity of GCHQ. Furthermore, following the ground breaking nature of the work at Bletchley Park, the possibility that GCHQ might possess expertise that outstrips the expertise produced in other contexts has also found its way into the popular imagination.

After the Second World War, GC&CS was renamed the Government Communications Headquarters. In 1951, it began the process of moving from London to Cheltenham, Gloucestershire, where it remains to this day. During the war, GC&CS underwent a number of changes that set the trajectory for how GCHQ would operate during the Cold War and beyond. Whilst it may have entered the

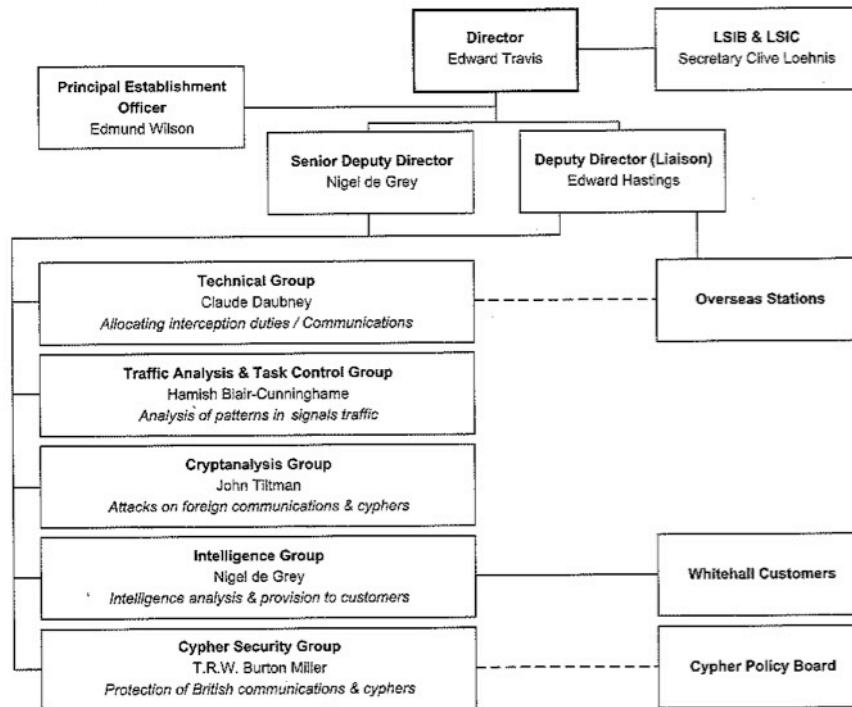


FIGURE 7.1: The Organization of GCHQ in 1946 (Aldrich 2010)

Second World War as a fairly disorganized collection of amateurs and eccentrics, in the years that followed, it emerged as a confident, professional and efficient, albeit smaller, organization (Aldrich 2010, p.69). Part of this process was the forming of intelligence alliances with other nations. The BRUSA (Britain-USA) agreement of 1946, followed by the UKUSA (United Kingdom-USA) agreement of 1948, essentially meant that all of the major English-speaking nations - including the UK, the US, Canada, Australia and New Zealand - would share intelligence (Rudner 2004). Despite some early disagreements over enciphering methods, the UKUSA agreement remains in place at the time of writing, although many of the details surrounding it remain classified.

The immediate priority for GCHQ, and indeed other UKUSA nations, following the war was monitoring the communications of the Soviet Union. Of particular interest were those related to the Soviet atomic bomb project. However, success in this endeavour was limited. High-level Soviet communications were encrypted using one-time pads - ciphers that were essentially unbreakable due to a lack of repeat use. The lack of success in this area prompted GCHQ, in the 1950s, to move away from attempting to derive intelligence through the use of cryptology, and towards the use of bugging equipment and electronic intelligence (ELINT). This

usually took the form of attempting to read unencrypted radar signals (Aldrich 2010, pp.108-110). Starting in the late 1960s, the use of satellites and other wireless technologies to transmit telephone calls prompted GCHQ and others to begin changing their methods once more. GCHQ began to build large domed receivers that collected unencrypted communications from the ether. The building of large, highly visible listening stations, like GCHQ Bude in Cornwall, damaged GCHQ's anonymity. GCHQ had also become much larger. By 1966 it was the largest of the three intelligence services in terms of budget, and given that it employed around 11,500 people, had more staff than MI5 and MI6 combined, and was larger than the entire British diplomatic service (including staff in overseas embassies). All of this contributed to the fact that GCHQ was reluctantly gaining public notoriety.

7.4 The Communications-Electronics Security Group

From the 1970s onwards, the branch of GCHQ that was most strongly associated with cryptology research was the Communications-Electronics Security Group (CESG). Before examining what is known about their work, it is important to distinguish more sharply between two key intelligence activities related to cryptology: signals intelligence (SIGINT); and communications security (COMSEC). SIGINT refers to activities geared towards the interception and interpretation of intelligence from signals transmitted by others. COMSEC refers to the activities geared towards the protection of one's own signals from other parties that may be trying to carry out SIGINT on their communications. Cryptology expertise is required for both SIGINT and COMSEC. Though SIGINT and COMSEC are clearly intertwined, throughout the period thus far described, they were somewhat separate organizationally.

7.4.1 Formation

From the early 1950s until the late 1960s, SIGINT and COMSEC were essentially handled by two different organizations. GCHQ were responsible for SIGINT, and an organization called the London Communications Security Agency (LCSA) were responsible for COMSEC. It was not until 1969 that it was decided that GCHQ should be formally responsible for both.

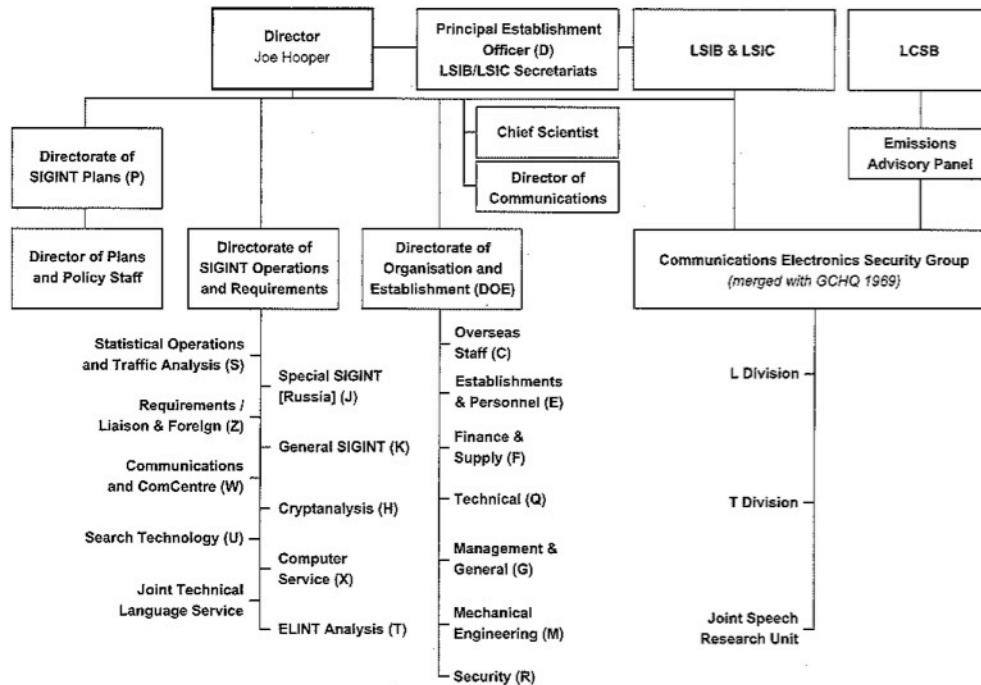


FIGURE 7.2: The Organization of GCHQ in 1970 (Aldrich 2010)

In the early 1950s, a review of the Cipher Policy Board's organisation and terms of reference led to the creation of a new agency, the London Communications Security Agency (LCSA). The LCSA had its own Director, but still remained administratively under GCHQ. In 1965, the LCSA became the Communications-Electronics Security Department (CESD), still based primarily in London, although parts were now co-located with GCHQ in Cheltenham, Gloucestershire. In 1969 CESD formally merged organisationally with GCHQ and was renamed the Communications-Electronics Security Group (CESG). In 1978 the last London elements of CESG moved to Cheltenham, where it has remained to the present day (Communications-Electronics Security Group 2012).

The original decision to separate SIGINT and COMSEC was taken when GCHQ began the process of moving to Cheltenham in the early 1950s. It was decided by senior intelligence officials that it would be better to make a fresh start in both areas, so responsibility for COMSEC was given to LCSA. Though the LCSA remained in existence for over twenty years, almost nothing is known about it (Aldrich 2010, pp.191-192). Richard Aldrich placed the 1969 integration of the CESG within the context of UK's relative economic decline - later highlighted

by the withdrawal from ‘East of Suez’. The integration of CESG was one consequence of a series of changes to the UK’s intelligence machinery that were designed to allow it to focus on obtaining economic and industrial intelligence, alongside military and diplomatic intelligence. It was also thought that a more tightly integrated government communications organization would ensure a more harmonious relationship with other UKUSA allies at a time when the relationship was showing signs of strain due to the emergence of American technical superiority (Aldrich 2010, pp.241-242).

Today, CESG is referred to as the ‘National Technical Authority’ for advice and services to protect governmental voice and data networks. As such, it remains the department of GCHQ responsible for COMSEC work. However, prior to the mid-1980s, its role was far less public. Very little is known about its activities prior to this. Though it has been possible for historians to construct a partial history of GCHQ during this period, a parallel history of CESG has not yet been produced. Almost no academic material exists that specifically deals with CESG, and very few internal CESG documents have been released. In one sense, this is surprising, given that CESG now has a much more public role than many other departments within GCHQ. However, historically, the secrecy that has surrounded COMSEC has often been much higher than that which surrounds SIGINT, and information related to CESG’s activity remains classified. Despite this, it is now known that it was within this department that ‘non-secret encryption’ was developed.

7.4.2 Research on Non-Secret Encryption

The only piece of cryptology research carried out by CESG that has been declassified and revealed to the public was that on non-secret encryption. Non-secret encryption was the term used to refer to a series of theoretical advances made within CESG that very closely resemble research carried out independently in the US under the heading of public-key cryptography - particularly the Diffie-Hellman key exchange and the RSA algorithm (Diffie & Hellman 1976, Rivest et al. 1978). Though kept secret for some thirty years, work done under the heading of non-secret encryption was revealed to the public in December 1997. The announcement was supported by the unusual release of five internal CESG documents (Ellis 1970,

Cocks 1973, Williamson 1974, 1976, Ellis 1987). Electronic copies of these documents are still available online from various sources.⁴ The physical documents - assuming they still exist - have not been made available. The work on non-secret encryption is the only major piece of post-Second World War cryptology research carried out within GCHQ that is now publicly known. How it came to be revealed, and the reaction it caused in the cryptology research community, will be discussed in chapter 8. The description that follows summarises the content of the documents.

The work on non-secret encryption was initiated by James H. Ellis. Ellis was born in Australia in 1924, but grew up in London. He studied for a degree in physics at Imperial College London, before working at the Post Office Research Station at Dollis Hill. Ellis joined GCHQ in 1952, and transferred to CESG (or CESD, as it was then known) in 1965. In 1969, Ellis was one of about half a dozen researchers working within CESG on long-range, ‘blue-sky’ projects. Ellis was working on what is known in cryptology as the key distribution problem. The problem, as it was then understood, centred on the fact that, if parties wished to communicate with one another in secret, they must all share the details of the process used to encrypt the message, and the reverse, which can be used to decrypt it. This symmetrical system left the parties with the problem of securely communicating the key. As telecommunications became more widespread, the key distribution problem grew. However, a solution was rarely sought, given that the sharing of keys between sender and recipient was considered to be one of the fundamental tenets of cryptology. Writing later, Ellis recalled that:

It was obvious to everyone, including me, that no secure communication was possible without a secret key, some other secret knowledge, or at least some way in which the recipient was in a different position from an interceptor. After all, if they were in identical situations, how could one possibly be able to receive what the other could not? Thus there was no incentive to look for something so clearly impossible (Ellis 1987).

Upon the discovery of a classified and unsigned paper produced by Bell Laboratories dating from the Second World War, Ellis began to change his view. The

⁴Electronic copies of these documents, and many other previously classified documents produced by intelligence organizations, may be downloaded from <http://www.cryptome.org>.

paper described how it might be possible for a sender to add analogue noise to a communication that could then be removed by the recipient if they knew exactly how it was generated. This demonstrated to Ellis that “secure communication was at least theoretically possible if the recipient took part in the encipherment” (Ellis 1987).

Ellis was subsequently able to conceive of a system whereby a communication can be secured by, first, the recipient generating a large number which is then transformed to a different number using a one-way mathematical function - a function that cannot easily be reversed. This new number is then sent to the sender, who uses it with a second function to scramble the message again before sending it back. The (original) recipient is then able to unscramble the message using the original number, which remains known only to them. Of course, such a system depends on there actually being a usable one-way mathematical function, and in 1969, it wasn't obvious to Ellis that such a function existed. Ellis produced an internal paper detailing his idea, and passed it on to Shawn Wylie, a chief mathematician at GCHQ. Although Wylie reported that the system appeared to be sound in principle, it was clear that it would need to be developed further to be of any use.

By 1971, the arrival of a new Chief Scientist at CESG had reignited interest in the system that Ellis had proposed. However, the search for usable functions remained unsuccessful, and was probably hampered by the fact that the system itself was still thought of as somewhat heretical, given that it violated one of the core assumptions of the discipline. In 1973, the problem found its way to new CESG employee Clifford Cocks. Before joining CESG, Cocks studied for an undergraduate degree in mathematics at the University of Cambridge, and a postgraduate degree at the University of Oxford. Upon arrival at CESG, Cocks was mentored by Nick Patterson. Patterson passed Ellis' idea onto Cocks. Although Patterson was aware that the problem had proved difficult in the past, he speculated that it might be useful to introduce it to someone who would approach it from outside of the context of key distribution. As Cocks would later suggest in a paper on the subject, the best way to produce a one-way function would be through the multiplication of prime numbers. Whilst it is straightforward to multiply large prime numbers, it is very difficult (although not impossible) to identify the prime numbers used if only their product is known (Cocks 1973). This method was almost identical to that

used in Rivest, Shamir and Adleman's RSA algorithm - developed independently in the US some years later.

Cocks' method became known throughout CESG, and came to the attention of their other mathematicians. Malcolm Williamson examined the solution, and after failing to find any flaws, developed his own scheme for the sharing of keys that fitted the non-secret encryption model. This method, which was written up in 1974, closely mirrors what would later be called the Diffie-Hellman key exchange, when it was also independently developed in the US (Williamson 1974).

In order to keep the work on non-secret encryption secret, and to comply with the Official Secrets Act, no material relating to it was disseminated outside of the intelligence community. This meant that the work did not appear in any form of publication, and was not presented at academic conferences. Furthermore, those involved in its development were prohibited from discussing it with colleagues outside of the UKUSA agreement. This was despite the fact that the work on non-secret encryption had the potential to completely subvert some of the classical tenets of cryptology. Indeed, it is perhaps a testament to how much secrecy informed research practices at CESG that they were able to keep this work secret until they chose to reveal it in the late 1990s.

Exactly what happened to the work on non-secret encryption immediately after it was developed remains unclear. It appears to have been left as an interesting idea, and was not put into practice by GCHQ or their UKUSA allies. Levy (2001, p.324) claimed that it shifted from being seen as impossible to impractical. Levy also claimed that GCHQ saw non-secret encryption as being something that could only potentially be used for transmitting messages, and as such, unlike those involved with work on public-key cryptography, did not anticipate that it could also be used for message authentication and data integrity. Aldrich (2001, p.491) claimed that non-secret encryption was later shared with Washington via Sean Wyllie, but that they were equally uninterested in developing it further.⁵ Again, because of the continued secrecy surrounding CESG and GCHQ, it is very difficult to place the work on non-secret encryption in any kind of context. However, it is unlikely that the documents relating to non-secret encryption are in any way 'representative',

⁵Aldrich and Levy disagree on how to spell this individual's first name and surname. This highlights the fact that, due to the secrecy surrounding CESG, even the most basic information relating to employees is both hard to source and hard to verify.

in the sense used by Scott (1990), of the totality of available documents related to CESG cryptology research.

7.4.3 Increased Public Awareness

During the 1970s, at around the same time as work on non-secret encryption was being carried out, the public were becoming increasingly aware of the existence of GCHQ, and the work of UK intelligence organizations more generally. Moran has argued that:

By the 1960s the state had concluded that maintaining absolute secrecy with respect to some of its work was not only impossible but also counterproductive ... With this, the state moved into the realm of 'offensive' information management, putting 'secrets' into the public domain on its own terms. The traditional 'defensive' approach of saying and releasing nothing was seen as too rigid. What was needed was flexibility (Moran 2013, p.5).

As a result, a certain tolerance of public information about GCHQ began to emerge. This tolerance even extended to information about GCHQ's past cryptology research. In 1974, F. W. Winterbotham - a former Royal Air Force (RAF) officer - published *The Ultra Secret*. This book provided the first public account of the code breaking efforts at Bletchley Park during the Second World War. Up to that point, this work had actively been kept secret, and was even effaced from authorized histories of the Second World War, such as those written by Winston Churchill. In 1976, Duncan Campbell and Mark Hosenball published an article in *Time Out* magazine entitled 'The Eavesdroppers'. This article was one of the first to publicly allude to GCHQ's current activities. However, details were sparse, and the article contains little more than speculative asides. Despite these flaws, both accounts undoubtedly contributed to an increased public awareness of GCHQ, and cemented its association with cryptology.

GCHQ returned to the public's attention again in 1984, when Margaret Thatcher's Conservative government successfully banned its employees from becoming union members in the interests of 'national security'. The New Labour government eventually overturned this decision in 1997. GCHQ also featured in Peter Wright's

notorious 1987 memoir *Spycatcher*. Wright, amongst other things, claimed that the intelligence services (particularly MI5) were secretly plotting against former Labour Prime Minister Harold Wilson when his government were in power in the 1960s and 70s. Many of Wright's claims have subsequently been discredited, but at the time, they again contributed to public awareness of GCHQ.

Following the end of the Cold War, the relationship between the UK's intelligence organizations and wider society began to change. This change can be seen as the result of processes like the Waldegrave Initiative - a process of legislative reform that was initiated in the late 1980s and early 1990s. Whereas, in the past, a position of absolute secrecy was adopted, the Waldegrave Initiative deliberately fashioned a culture of greater openness around intelligence organizations. The Conservative government's 1993 white paper on Open Government - resulting from an initiative launched by the Chancellor of the Duchy of Lancaster, William Waldegrave, and the Foreign Secretary, Douglas Hurd - led to a wide range of processes related to the re-reviewing of previously retained historical material. The public have since been granted access to a number of previously retained official documents related to, or created by, the UK's intelligence organizations. Furthermore, the policies that had underpinned some of the more extreme security measures were questioned. Under the 1994 Intelligence Services Act, GCHQ and SIS were placed on a statutory basis for the first time. This meant that the intelligence services and their directors could be legally referred to by name, and that records held by other departments that did so would no longer be eligible for retention on these grounds (Bennett 2002).

Whilst the Waldegrave Initiative has allowed greater access to recent records related to some public bodies and government departments, requests to release records related to GCHQ are either exempt or can be refused. In practice, this has meant that few records relating to GCHQ produced after the end of the Second World War have been released (Bennett 2002). Therefore, in contrast to many other topics, when examining the history of GCHQ, the closer one gets to the present the less is known. This is due to the fact that the information about their activities up to and including the Second World War are considered less secret than information that may pertain in some way to the present.

7.4.4 Public Role

These reforms occurred at the roughly same time as CESC started to take on more public duties. Whereas the work on non-secret encryption indicates that CESC had in the past been concerned with mathematical and theoretical research, their research in the 1990s included work geared towards the fulfilment of their role as the UK's National Technical Authority. In general, this resulted in the provision of cryptographic algorithms, but also included “generic cryptographic research”, the “design and development of general-purpose or bespoke cryptographic products” and the “development of cryptographic algorithms and integrated circuits”. CESC also used the expertise produced by this research to provide “evaluation and certification of products”, “setting up production contracts for licensed cryptographic products”, “updates and Post Design Services for CESC-designed cryptographic products, and technical support to users”, “[advice to] industry on commercial risk development of cryptographic products and systems for the UK official market”, and “independent advice on the suitability, application and integration of commercial off-the-shelf cryptographic products in public sector projects” (Communications-Electronics Security Group 1998*a*). However, it is important to re-iterate that these services were only available to public bodies, and knowledge of them still required security clearance.

Aside from these activities, CESC also played a role in the UK's official licensing schemes for computer security products. In 1985, CESC established facilities for evaluating the security of government computer systems. Then, in 1987, the DTI established the CESC-managed Commercial Computer Security Centre (CCSC). The CCSC was responsible for formally evaluating commercially available IT products. This resulted in the publication of a set of evaluation criteria known as the ‘Green Books’. In 1989, it was announced that a new nationwide scheme would be developed. This came to be known as the UK ITSEC scheme, and became fully operational in 1991. ITSEC, which was eventually harmonised across many European countries, was eventually incorporated into the Common Criteria. Although the development of the UK ITSEC Scheme is only partially dependent on cryptology, it does demonstrate how CESC acquired a more public role during this period.

7.4.5 Continued Secrecy

Despite the fact that CESG took on more public duties during the 1990s, given that it was a part of GCHQ both organizationally and physically, much of its activity remained secret. Those working there remained bound by the Official Secrets Act, and therefore carried out research in line with the practices that were designed to uphold it. Although those working at CESG would very occasionally publish the results of cryptology research carried out during this period (e.g. Cocks 1997, 2001), the vast majority remained classified. Furthermore, even though the research carried out there underpinned the cryptology services they provided to public bodies, information about the processes used and the motivations behind their cryptology research were not revealed as a part of this service. Although individuals who had worked for CESG during this period could not be interviewed, others who had come into contact with the practices used there in the 1990s were able to describe them to me. Practices related to security and secrecy were mentioned frequently when the topic of CESG or GCHQ was raised. In talking about the security of the CESG site, one interviewee explained:

Respondent: CESG is *within* GCHQ. You know the GCHQ site, and you know how secure that is, well CESG is fenced off within that. You're not allowed to take mobile phones, laptops, or anything into it. It is quite secure.

CESG also implemented practices to prevent individual names being associated with them, and vice versa:

Respondent: They were very strange. I went down to a meeting once at CESG, and it was for the whole security community, and you'd have the guest list there, and it would have your name and where you worked ... And then you'd have these people who were just names, and blank ... And if you sent them anything, because occasionally they would say "Oh, we want one of your reports", you had to double-envelope. You'd put it in an envelope, put their name and address on the outside, then you'd put it in another envelope and send it to this holding address, and then it would get sent internally.

CESG practices, in particular those associated with vetting and document marking, were also known to produce somewhat absurd situations:

Respondent: There were different levels of vetting. We were all vetted up to ‘Secret’. Most of our work, particularly work funded by the MOD, was classed as ‘Secret’. [Employee name], one of his reports, that he wrote, he wasn’t allowed to see it, because they classified it as ‘Top Secret’, and he wasn’t vetted up to ‘Top Secret’ . . . It was strange like that . . . He also wasn’t allowed to speak about it. Not entirely sure why. But you expect it if you work in that field.

This demonstrates that, although the stated role of CESG changed following the end of the Cold War, many of the practices that sustained it between 1990 and 2000 did not. When the political phase of the crypto wars began, CESG and GCHQ were still highly secretive organizations, and this characteristic was imprinted upon the expertise that they produced.

7.5 Conclusion

In this chapter, I have described the cryptology research that has been carried out at GCHQ. Unlike the research sites described in the preceding three chapters, the history of cryptology research of GCHQ extends back beyond 1970. Indeed, the most well known cryptology research carried out there came during the Second World War, with Alan Turing’s work on breaking ciphers generated by the German Enigma machine, and the construction of the Colossus computer to break the Lorenz cipher. Following the Second World War, as the use of one-time pads became more common, and later, as intelligence came to be increasingly derived from the analysis of unencrypted electronic communications, research into cryptology appears to have become less of a priority. Somewhat paradoxically, it was also during this period that mathematicians working within CESG carried out theoretical work on non-secret encryption. This work mirrored the work on public-key cryptography that was completed in the US nearly a decade later. During the 1990s, CESG took on a more public role, and carried out cryptology research commensurate with its role as the UK’s National Technical Authority for the official use of cryptology.

On the basis of what is known about the history of cryptology research at GCHQ, it could be characterized in a number of ways. However, it is important to note that compared to the data available on cryptology research at other sites in the UK, the available data on cryptology research carried out by GCHQ is very small. Although fragments of CESC's work have been revealed, it should be remembered that CESC is not the only department within GCHQ to carry out research related to cryptology. As can be seen in Figure 7.2, a 'Cryptanalysis' division existed outside of CESC, and it is very likely that this division both produced and possessed cryptology expertise. However, as this division lies outside of CESC, the work that it has carried out remains secret, and no information relating to it has been released. Despite processes like the Waldegrave Initiative, the secrecy surrounding GCHQ and CESC makes an explicit analysis of their cryptology research difficult. Although some records related to intelligence have been released, the majority remain withheld. The process of review is slow and resource-intensive, and ultimately, not a high priority for the intelligence agencies themselves (Bennett 2002). Furthermore, due to the fact that the information held within some records remains sensitive long after their immediate use (information pertaining to living individuals and their relatives, for example), as a general rule, GCHQ will not release records that were created after the end of the Second World War (Bennett 2002). Although some personal accounts have been published related to the inner-workings of intelligence agencies, individuals remain bound by the Official Secrets Act after they stop working for them. This makes conducting interviews with those who have worked for intelligence agencies very difficult from a practical, legal and ethical point of view - especially given that the government have in the past attempted to prosecute individuals for revealing sensitive information (Easter 2008, p.682).

These issues could be seen as methodological difficulties, and nothing more. Given the relative lack of data on cryptology research carried out within GCHQ, it may therefore be tempting to abandon a characterization of their cryptology research practices. However, this conclusion is rather unsatisfactory, as it fails to capture how the lack of public knowledge about the work of GCHQ impacted upon others working in the field, and does not reflect the curious pose that CESC struck between upholding secrecy and performing a clearly defined public role. Arriving at this conclusion also requires a certain betrayal of the approaches to documentary analysis that were discussed in chapter 3. In particular, the documentary realities

argument has discouraged researchers from seeing documents as sources of descriptions of reality, but rather as a part of the practices used to construct it (Atkinson & Coffey 2011). On this understanding, the inaccessibility and unavailability of documentary data can be seen as a result of practices that have been designed to uphold secrecy. For example, not publishing scientific work is not the absence or failure of practices used to publish in other institutional contexts. Rather, it is the presence and success of practices designed to minimize the extent to which that work is known. Therefore, on this understanding, it may be concluded that the practices employed by CESG were designed to produce secret contributory expertise in cryptology.

Chapter 8

The Political Phase of the Crypto Wars

8.1 Introduction

In the preceding four chapters, I have provided an overview of the nature of cryptology research practices at four different research sites. The research described can be thought of as constituting the technical phase of the crypto wars. I have also hinted at the nature of the expertise that these practices enacted. Up to this point, I have not described any of the developments that can be thought of as constituting the political phase of the crypto wars. That is the purpose of this chapter.

In this chapter, I will break the political phase of the crypto wars down into two debates over issues related to cryptology: a debate over Trusted Third Parties; and a debate over export controls. In contrast to previous descriptions, as outlined in chapter 1, I will place an emphasis on describing how the expertise produced during the technical phase was used during the political phase. By the end of this chapter, it will start to become clear that the expertise enacted at each cryptology research site was used during the crypto wars in different ways and for different purposes.

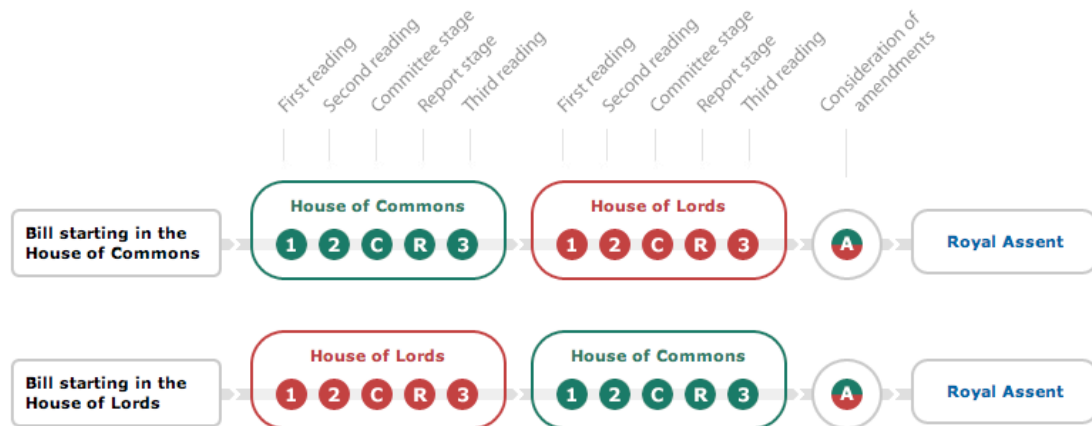


FIGURE 8.1: The Passage of a Bill in the UK Parliament (UK Parliament 2013)

8.1.1 A Note on Organization

In terms of both data collection and presentation, I have used the procedure used to pass a bill in the UK parliament as a general organizing principle (see Figure 8.1).¹ This has allowed me to identify separate debates by looking at the activity that surrounded the passage of different bills, and then to break that activity down into different phases based on government statements and publications. Though not part of the formal legislative process, I have also paid particular attention to the consultation periods that preceded the drafting of each bill.

If this process is used as an organising principle, two political debates featuring cryptology expertise can be identified. As I will describe later in this chapter, these debates are linked. Together, they can be thought of as constituting the political phase of the crypto wars in the UK. In broad terms, the first debate concerned proposals to implement Trusted Third Parties, and the second debate concerned changing the rules governing cryptography export. What follows is a chronological description of each debate. Within this structure, I will pay particular attention to the role of the cryptology research sites that were described in the previous four chapters. Thus, in contrast to other descriptions of the political phase of the crypto wars in the UK, I will describe political developments in relation to the cryptology expertise that was produced during the technical phase.

¹For a brief overview of the processes used to pass a bill in the UK parliament, see (UK Parliament 2013).

8.2 The Debate over Trusted Third Parties

The debate over Trusted Third Parties (TTPs) was framed by the government through questions over the best way to balance the potential economic benefits of widespread access to cryptology against the negative impact this might have on law enforcement capabilities. On the one hand, it was clear that cryptology could be used to promote electronic commerce. In the mid-1990s, there were concerns over a general lack of public trust in the security of the Internet for the transmission of financial information. Cryptography technologies could be used to encrypt this information, making it very difficult for potential fraudsters to obtain, thus removing one of the barriers to the adoption of electronic commerce. However, widespread access to cryptography technologies might also allow criminals to encrypt their communications and stored data, thus making it very difficult for law enforcement bodies to prevent crime or gather evidence. The government proposed a solution to this dilemma in the form of a system of Trusted Third Parties - organizations acting as state-licensed intermediaries that would provide cryptography services to the public, but would also be obliged to give law enforcement bodies the means to decrypt communications if requested. These proposals brought a number of propositional questions to the fore, such as: is the proposed nationwide TTP system, and the technologies and assumptions it's based on, secure? As will become clear, some with contributory cryptology expertise believed that it was, whereas others did not. However, this is only one dimension of the debate. The manner in which the proposals were formulated and then debated is also significant. As will become clear, those from the Security Group at the University of Cambridge in particular felt that the TTP proposals were being formulated without drawing on the expertise of the entire cryptology research community, including the expertise that they themselves had enacted. Although the debate over TTPs can be seen in terms of conflicting answers to propositional questions, when looked at through the lens of elective modernism, it can also be understood in terms of how expertise from the technical phase was used (or not) during the political phase.

8.2.1 Background

Before describing the debate over TTPs, it will be useful to provide some details about the wider national and international context. At the international level, by the time the political phase of the crypto wars in the UK was underway in

the early to mid 1990s, the European Commission (EC) - the executive body of the European Union (EU) - had already begun to address some of the issues that widespread access to cryptography technologies raised. In 1992, the Senior Officials Group - Information Systems Security (SOG-IS) was established to advise the EC on appropriate legislative steps. SOG-IS conducted pilot projects to investigate the challenges that cryptography posed. However, lengthy debates over the surrounding issues meant that by the time the crypto wars were underway, a clear Europe-wide policy framework had not been agreed upon.

By the time the crypto wars were underway, SOG-IS had overseen the development and harmonisation of the Information Technology Security Evaluation Criteria (ITSEC). As was described in previous chapters, ITSEC was a set of criteria that were implemented in several European countries - including France, Germany, the Netherlands and the UK - for the purposes of evaluating computer security products. The UK Information Technology Security Evaluation and Certification Scheme, which organized the testing of security products and technologies against ITSEC, was managed by CESG and accredited by NPL. The scheme had been running for around five years by the time the debate on TTPs started.

There also existed international bodies that provided guidance to nation states that planned to legislate on matters related to the use of cryptography. Typically, these bodies did not have the power to legislate directly, but could produce guidelines for nations to use as a framework for their laws. The Organisation for Economic Co-operation and Development (OECD) was founded in 1961 to stimulate economic progress and world trade. Unsurprisingly, the advent of electronic commerce was a key issue for the OECD in the 1990s. In 1995, the OECD held a conference for representatives from industry and government to discuss the impact of cryptography. This led to the formation of an Expert Group that met four times in 1996. The group produced a paper on 'OECD Guidelines for Cryptography Policy' (Organisation for Economic Co-operation and Development 1997). It laid out eight principles for national legislation. They were, that:

1. Cryptographic methods should be trustworthy in order to generate confidence in the use of information and communications systems;
2. Users should have a right to choose any cryptographic method, subject to applicable law;

3. Cryptographic methods should be developed in response to the needs, demands and responsibilities of individuals, businesses and governments;
4. Technical standards, criteria and protocols for cryptographic methods should be developed and promulgated at the national and international level;
5. The fundamental rights of individuals to privacy, including secrecy of communications and protection of personal data, should be respected in national cryptography policies and in the implementation and use of cryptographic methods;
6. National cryptography policies may allow lawful access to plaintext, or cryptographic keys, of encrypted data. These policies must respect the other principles contained in the guidelines to the greatest extent possible;
7. Whether established by contract or legislation, the liability of individuals and entities that offer cryptographic services or hold or access cryptographic keys should be clearly stated;
8. Governments should co-operate to co-ordinate cryptography policies. As part of this effort, governments should remove, or avoid creating in the name of cryptography policy, unjustified obstacles to trade (Organisation for Economic Co-operation and Development 1997).

Whilst open to interpretation, the wording of principle number six appeared to suggest that, though government access to encryption keys could be lawful, this concern should not be prioritised over the other principles.

At the national level, in terms of law enforcement and communications, the most important piece of legislation prior to the start of the debate over TTPs was the Interception of Communications Act 1985. This piece of legislation made it an offence to unlawfully intercept communications sent by post or by a public telecommunications system. It also established a procedure for law enforcement bodies to obtain a warrant from the Secretary of State to lawfully intercept communications in certain circumstances. Section 2(2) of the act stated that the Secretary of State will not issue a warrant under this section unless he considers that the warrant is necessary:

1. In the interests of national security;

2. For the purpose of preventing or detecting serious crime; or
3. For the purpose of safeguarding the economic well-being of the United Kingdom (Interception of Communications Act 1985, p.2).

In summary, prior to the debate on TTPs, there already existed powers that allowed the legal interception of electronic communications for certain law enforcement and economic purposes. There also existed guidance for European nations on cryptography policy, as well as frameworks for assessing the conformance of security technologies to European standards. The cryptology expertises of those working at CESG and NPL were put to use in the day-to-day running of the latter. Finally, it is worth making clear that, in contrast to the situation in the US, there's no evidence that law enforcement or intelligence agencies attempted to subvert the course of cryptology research in universities or in industry.² However, the Clipper Chip proposals - which were discussed in chapter 1 - were known to many of those involved in the technical and political phases of the crypto wars in the UK, and this may have altered the way they viewed debates over cryptology regulation.³

8.2.2 The Announcement of the TTP Proposals

In this subsection, I will describe the way in which the government's TTP proposals were announced. As will become clear, the TTP proposals were announced suddenly, and were immediately controversial because of the nature of what they proposed. Amongst some cryptology experts, they were also controversial because of the lack of technical detail they contained, their finalised tone, and the narrowness of the expertise they appeared to be based upon.

By 1994, the crypto wars in the US were prompting concerns over whether the UK government had similar intentions. These concerns filtered through to a small group of MPs. This prompted David Shaw MP on 21st April 1994 to formally ask the Department of Trade and Industry (DTI) whether they were planning on

²One interviewee from the Information Security Group at Royal Holloway told me that no-one from CESG or GCHQ had ever tried to stop them carrying out cryptology research.

³There is some evidence to suggest that those broadly opposed to the TTP proposals saw them in terms of attempts to introduce the Clipper Chip in the US. For example, in an early discussion of the proposals, they were referred to informally as "HMS Clipper" (Gladman 1996a).

introducing any legislation related to cryptology. The DTI replied that they had “no current plans to introduce legislation relating to data encryption” (HC 1994).

This position was officially reversed just over two years later with a statement from the DTI on June 10th 1996 (Department of Trade and Industry 1996*b*). The statement was fronted by Ian Taylor MP - a Conservative minister for Science and Technology. In short, the ‘Paper On Regulatory Intent Concerning Use Of Encryption On Public Networks’ described how the government would attempt to juggle the desire to engender trust in electronic commerce whilst maintaining the possibility of effective law enforcement. The paper stated that the government intended to introduce legislation to:

Facilitate the development of electronic commerce by the introduction of measures which recognise the growing demand for encryption services to safeguard the integrity and confidentiality of electronic information transmitted on public telecommunications networks (Department of Trade and Industry 1996*b*).

However, the government also stated an aim to:

Preserve the ability of the intelligence and law enforcement agencies to fight serious crime and terrorism by establishing procedures for disclosure to them of encryption keys, under safeguards similar to those which already exist for warranted interception under the Interception of Communications Act (1985) (Department of Trade and Industry 1996*b*).

The government proposed to do this through the use of TTPs. TTPs would be commercial or non-profit organizations that would act as intermediaries between two parties - say, an online vendor and a customer - that wished to communicate securely using cryptography. In terms of who the government saw as appropriate potential TTPs, they stated that they “would expect organizations with existing customers, such as banks, network operators and associations (trade or otherwise) to be prime candidates” (Department of Trade and Industry 1996*b*).

Trust was identified as the key issue for TTPs. The government believed that TTPs would be trusted in their role, because they would each be required to obtain a license from the Secretary of State for Trade and Industry:

By their nature, TTPs, whatever services they may provide, will have to be trusted by their clients. Indeed in a global trading environment there will have to be trust of, and between, the various bodies fulfilling this function. To engender such trust, TTPs providing information security services to the general public will be licensed. The licensing regime would seek to ensure that organisations and bodies desiring to be TTPs will be fit for the purpose. The criteria could include fiduciary requirements (eg appropriate liability cover), competence of employees and adherence to quality management standards (Department of Trade and Industry 1996*b*).

Crucially, the government stated that encryption keys - the vital information required to encrypt and decrypt communications - would be held by TTPs in escrow. As encrypted communications were seen by the government as having potential law enforcement implications, “TTPs would also be required to release to the authorities the encryption keys of their clients under similar safeguards to those which already exist” (Department of Trade and Industry 1996*b*). Once the TTP had handed over an individual’s encryption keys, law enforcement bodies would be able to use them to decrypt intercepted communications or stored information. This, unsurprisingly, turned out to be particularly controversial.

The initial statement from the government announcing their proposals for TTPs, and many of the documents that followed, were short on technical details. They typically did not provide specific details about the technologies that would be used. However, the TTP system would clearly need to be built on top of an appropriate technological infrastructure, comprised of specific protocols and algorithms, and decisions about what technologies to use would be important. Some of those following the developments believed that the so-called ‘Royal Holloway’ protocol - developed by the Information Security Group at Royal Holloway - would form the basis of the proposed TTP system.⁴ Though this was not clearly stated by the government, the belief stemmed from details about two public-sector information security schemes that were already underway. The first of these proposed the use of the Royal Holloway protocol to secure emails sent between government departments (Communications-Electronics Security Group 1996). The second scheme

⁴The Royal Holloway protocol was also referred to as ‘CASM’ or the ‘GCHQ’ protocol at various points during the crypto wars. This is because it was modified slightly at various points to suit the context in which it would be applied.

proposed the use of the Royal Holloway protocol - in conjunction with a classified CESG-developed algorithm called Red Pike - to secure the flow of clinical data around a proposed National Health Service (NHS) network (Zergo 1996). Ross Anderson and Michael Roe (1997) from the Security Group at the University of Cambridge believed that the technologies earmarked for use on these projects would also be used for the TTP scheme to ensure compatibility when they became part of the same government network.

The Royal Holloway protocol was developed by Nigel Jefferies, Chris Mitchell, and Michael Walker (1995).⁵ It was first presented at the Cryptography Policy and Algorithms Conference in Queensland, Australia, in July 1995. The title of the paper was ‘A Proposed Architecture for Trusted Third Party Services’. The paper was also presented at subsequent conferences in America and Switzerland in the following year. The protocol emerged out of the aforementioned UK DTI/EPSRC research project entitled ‘Security Studies for Third Generation Telecommunications Systems’ (GR/J17173/01), carried out by the Information Security Group in conjunction with Vodafone, GEC Plessey Telecommunications (GPT), and the DTI.

The paper proposed “a novel mechanism that will enable TTPs to perform the dual role of providing users with key management services and providing law enforcement agencies with warranted access to a particular user’s communications” (Jefferies et al. 1995). As such, it matched the requirements that would later be outlined in the government’s TTP proposals. The paper also described how it would be possible for two parties - using two different TTPs - to communicate in secret. Furthermore:

Should there be a warrant for legal interception of this communication, an intercepting authority can retrieve the private key of one of the users from the associated trusted third party within its jurisdiction and use this with the public key of the other user (which is transmitted along with the encrypted message) to find the session key for the encryption. There is no requirement for the intercepting authority to retrieve the private keys of both users (Jefferies et al. 1995).

⁵At the time, only Mitchell held a position within the Information Security Group. However, all three had previously worked at Racal and maintained strong links with the group. On the original paper, the Information Security Group at Royal Holloway was listed as the institutional affiliation of all three authors.

A unique feature of the protocol was that it allowed law enforcement bodies to decrypt both incoming and outgoing traffic. In the set of assumptions about the requirements of TTPs, the paper stated that “access must be provided to the subject’s incoming and outgoing communication, where a warrant is held” and that “this is clearly achieved for the proposed scheme, as the subject’s TTP can provide the appropriate send and receive private keys” (Jefferies et al. 1995). This was in contrast to most other protocols, where decryption was permitted, but only in a single direction. However, this feature also matched with what would later be outlined in the government’s TTP proposals.

The first public indication that the government intended to make use of the Royal Holloway protocol came when CESG published a report on another proposed system to secure government emails (Communications-Electronics Security Group 1996). The system, known as the ‘CESG Architecture for Secure Messaging’ (CASM), was designed to:

Facilitate pan-government secure inter-operability of electronic mail, by simplifying the implementation of secure electronic mail within government, ensuring secure electronic mail between departments is possible, attempting to facilitate future inter-operability with commercial users, maximising the use of commercial technology in a controlled manner, whilst allowing access to keys for data recovery or law enforcement purposes if required (Communications-Electronics Security Group 1996).

The report clearly stated that CASM was “based upon a proposal by the Royal Holloway College [RHC] for trusted third party services” (Communications-Electronics Security Group 1996).

This was followed by the release of a report on proposals to use cryptography to secure the flow of clinical data around NHS networks. Prior to the publication of this report, it had long been believed that the networking of electronic clinical data within the NHS could offer benefits, and that concerns over privacy formed one of the main obstacles. Some, including the British Medical Association (BMA), were particularly concerned over the security of the planned NHS network. Specifically, they were worried that non-medical personnel working for the NHS might be able to use the system to access patient data. Cryptography appeared to offer

a solution. The Information Management Group (IMG) of the NHS Executive commissioned a report to examine “the ramifications of using encryption and related services across the NHS-Wide Network (NWN)”. The report was produced by Zergo Ltd. Zergo - as was described in chapter 5 - was a cryptography solutions firm founded by Henry Beker in 1988. As such, Zergo had links with the Information Security Group at Royal Holloway and Racal. The Zergo report recommended that “the NHS will need to develop a key management infrastructure to sit across the many different systems that will become encryption-enabled. This infrastructure will require one or more Trusted Third Parties (TTP) management centres” and that “the NHS’s needs should be addressed by a family of related encryption products built on the Red Pike encryption algorithm. This algorithm has recently been made available to the NHS by CESG, the National Technical Security Authority within HMG” (Zergo 1996). This could be read so as to suggest that plans were already being put in place for a system that would be compatible with a future national TTP infrastructure.

Because of the stated intention to use it for the NHS network, it was believed that the proposed nationwide TTP system would also make use of the Red Pike algorithm. Little is known about Red Pike. It was developed in secret at CESG, and information relating to it remains classified. In 1996, CESG commissioned an analysis of Red Pike by Codes and Ciphers Ltd - a cryptography consultancy company set up by Fred Piper which also drew on the expertise of individuals from the Information Security Group at Royal Holloway (Mitchell et al. 1996).

As was mentioned earlier, the BMA were critical of the security of the proposed NHS network. To aid their challenge, they asked Ross Anderson to act as a security consultant, so that they could propose their own security policy. Anderson was particularly critical of the decision to use a modified version of the Royal Holloway protocol (which Anderson preferred to refer to as the GCHQ protocol following their small modifications to it) in a system to secure government emails and in a system to secure clinical data, given a mismatch between their respective requirements:

The heart of the matter is that the IMG cryptography strategy appears to encourage the NHS to adopt protection mechanisms very similar to those designed by CESG (a department of GCHQ) to protect government electronic mail . . . However, the GCHQ protocol mechanisms

have different goals from those of the clinical professions. They attempt to keep a message between two officials secret from third parties, but available to both their superiors (and to the police and intelligence services) by ensuring that each official's departmental security officer has a spare copy of the key used to encrypt it . . . Clinical professionals, on the other hand, require safety and privacy. The origin and content of messages should be indisputable, whether for the purposes of immediate clinical decision making or for litigation many years later (Anderson 1997a, p.3).

Anderson (1997a) was also critical of Red Pike, and in particular, its proposed use for the NHS network. Along with a general criticism of the NHS strategy, Anderson questioned the choice of a classified algorithm - rather than algorithms such as SAFER, WAKE, and Blowfish - all of which had been strengthened by academic scrutiny.

The proposals were also criticised from a technical point of view. The most prominent critique of the technical aspects of Royal Holloway protocol, and GCHQ's modifications to it, came from a paper written by members of the Cambridge group. The paper, written by Ross Anderson and Michael Roe, used the HMG email and NHS reports to make a case for why the technical details of the protocol mattered:

If an unsound protocol were to be adopted across Europe, then this could adversely affect not just the secrecy of national classified data, the safety and privacy of medical systems, and the confidentiality of tax returns and government grant applications. It could also affect a wide range of commercial systems too, and make Europe significantly more vulnerable to information warfare. If the protocols were sound but inefficient, then they might not be widely adopted; or if they were, the costs imposed on the economy could place European products and services at a competitive disadvantage (Anderson & Roe 1997).

Here, Anderson and Roe argued that the adoption of the Royal Holloway protocol might prompt other nations within Europe to adopt it for standardization purposes. What followed was a detailed technical analysis, with the conclusions summarized in four key points:

1. The key management scheme gives us all the disadvantages of public key crypto (high computational complexity, long key management messages, difficult to implement on cheap devices such as smartcards), and all the disadvantages of secret key crypto (single point of failure, little forward security, little evidential force, difficulty of ‘plug and play’ with shrink-wrapped software). It does not provide any of the advantages that one could get from either of these technologies; and its complexity is likely to lead to the subtle and unexpected implementation bugs which are the cause of most real world security failures.
2. It is designed for tightly hierarchical organisations, and cannot economically cope with the more complex trust structures in modern commerce, industry and professional practice. Its main effect in government may to perpetuate rigid hierarchies and frustrate the efficiency improvements that modern management techniques might make possible.
3. It goes about establishing trust in the wrong way. To plan to bootstrap signature keys from a ‘national public key infrastructure’ of escrowed confidentiality keys shows a cavalier disregard of the realities of evidence and of safety-critical systems.
4. There are a number of serious technical problems with the modifications that have been made to the US Message Security Protocol, which underlies the UK government’s offering. Quite independently of the key management scheme and trust hierarchy that are eventually adopted, these modifications are unsound and should not be used (Anderson & Roe 1997).

The paper argued strongly that the protocol was “very poorly engineered” (Anderson & Roe 1997).

The reports detailing the TTP proposals, the NHS network, and the system to secure government emails, taken together, suggested to some that the government were rushing to implement a number of secure systems using the same technologies. Unusually, CESG director Andrew Saunders released a statement that openly addressed this issue. Saunders (1997) stated that elements of the HMG email and NHS proposals were mistakenly conflated:

Some commentators have confused the relationship between the recommendations and a National Health Service project for a secure network. The two are similar but distinct. In early 1996 Zergo Limited produced a study on the use of encryption services for an NHS-wide network. It recommended that the NHS should adopt X.509 Authentication Framework, Certification Authorities, X.509 version 3 certificates, Trusted Third Parties, Diffie-Hellman, Red Pike, DSA etc, but did not refer to CESC's recommendations, only to CESC. It has been incorrectly assumed that the recommendations are the same as the solution proposed for the NHS. However, CESC's programme is aimed only at HMG and has no connection with the Zergo proposal (Saunders 1997).

Furthermore, Saunders accused critics of the proposals of inferring too much from the information given in the two documents, in particular, the desire to promote the CESC-designed Red Pike algorithm:

Another common misconception is that the CESC Red Pike algorithm is being recommended for use in the public arena. No confidentiality algorithm is mandated in the recommendations: for HMG use, however, approved algorithms will be required; Red Pike was designed for a broad range of HMG applications (Saunders 1997).

These denials did little to allay the concerns that some had over CESC's input into the scheme, and what appeared to be the proposed use of classified technology (e.g. Gladman 1997, Shepherd 1997). Importantly, Saunders' response ignored the more technical criticisms levelled by Anderson and Roe.

The technical problems that some saw in the Royal Holloway protocol, and in other technologies supposedly earmarked for the TTP proposals, were exacerbated by the fact that it appeared to them that the government had already 'made up its mind' without consulting the wider cryptology community. As Brian Gladman, former scientist and computer security specialist at the Royal Radar Establishment in Malvern and the Ministry of Defence, wrote in response to the tone of the original DTI announcement:⁶

⁶Further details about how Gladman acquired his cryptology expertise are provided on his personal website: <http://www.gladman.me.uk>

Note here that the policy HAS BEEN decided - that's it folks - we know what is best for you - you know that you can trust us and we certainly don't want you to think for yourselves - heavens above, that how revolutions start! What did you say, it's all about democracy, governments serving their citizens, achieving open government, freedom of information and all that? You mean that we should actually seek views BEFORE we set our policy? No, no, we can't possibly do that - if we did that people might disagree with us and we may then find it difficult to have the policy we want - OOPS, I mean the policy that we have decided is best for you (Gladman 1996*a*).

As this quote highlights, views were not publicly sought on the TTP policy before it was announced. Furthermore, given that the protocol had been partly paid for by the DTI, and would possibly be used in conjunction with classified technology provided by CESG, some felt there had not been adequate opportunity to scrutinize the technology. It therefore appeared that the government had decided that the answer to the question "can a nationwide TTP system be secure?" was "yes". It was therefore felt by some, particularly the Cambridge group, that the government had not made proper use of the available cryptology expertise, and that if it had consulted more widely, it would have been compelled to consider an alternative view. To compound this, there was also a distinct lack of clarity about the technologies that the government were proposing to use. Though, by piecing together information from different documents, some were apparently convinced that the Royal Holloway protocol and the Red Pike algorithm would be used, this had not been made absolutely clear. Some, including members of the Cambridge group, clearly felt that the lack of clarity over the TTP proposals, and the lack of consultation, was due to the fact that the impetus for them was coming from GCHQ, and that the secrecy surrounding them prevented the policy, and associated technologies, from being properly scrutinized (Anderson 1997*a*).

To sum up how cryptology expertise had been used up to this point in the political phase, we can see expertise from different sites being put to use in different ways. Expertise from Royal Holloway was used to produce a protocol for use in the proposed TTP system. Expertise from Cambridge was not drawn upon during the formulation of the initial proposals. However, given that their work in the technical phase was at odds with the reasoning behind the proposals, it was used to critique the protocol and the broader assumptions it was based upon. The BMA

also used this expertise to challenge the principles upon which the NHS proposals were based, and to develop their own security policy arguments. Expertise from CESG appears to have fed into the political phase through their input on related schemes, and the proposed use of Red Pike algorithm. However, the lack of clarity surrounding CESG's activity also started to fuel speculation over their motives. CESG also appeared to draw on the expertise of cryptography consultancy with links to Royal Holloway in order to legitimise their proposals and diffuse speculation about their intentions. Expertise from NPL does not appear to have been visibly drawn upon by the government or used to influence the political phase in any other way.

8.2.3 The First Stage of Consultation

The announcement of the TTP proposals was marked by a lack of clarity over how cryptology expertise had been used to inform them. In common with many policy announcements, it was later followed by a government consultation exercise. This ostensibly offered an opportunity for those who had not been involved in the initial formulation of the proposals - or who had otherwise disagreed with them - to offer up their expertise. Prior to the start of the consultation process, several interested individuals decided to establish a mailing list dedicated to discussion of this issue. This list became known as the UK Cryptography Policy Discussion Group - but was informally referred to as the 'ukcrypto' mailing list. The mailing list became a hub of activity related to the political phase of the crypto wars, and was particularly popular with those who had specialist cryptology expertise. The list was established by Brian Gladman in July 1996. It had an initial membership of around 25 subscribers. Members came from a variety of backgrounds, including industry, the public sector, and academia. Though it is difficult to be sure, owing to the technical nature of much of the discussion it would appear that most of the members had some experience of cryptology research, or experience of research in an adjacent field. As such, membership of the list appears to have been made up of people who either had contributory or interactional cryptology expertise.

From reading through the mailing list's archives, it is clear that the majority of members were opposed to the government's proposals. However, the mailing list also featured occasional contributions from members who were somewhat supportive, so constructive debates were not ruled out. Nonetheless, particularly judging

from the early messages, it is clear that most saw the list as a platform from which to organise their opposition to the proposals and to communicate this opposition to policymakers. In establishing the mailing list, Gladman wrote:

My own view here is that we need to avoid any hidden agendas which will lay us open to accusations from both crypto fascists and anarchists that we are conspiring to undermine their undoubtedly correct positions. Hence I believe that it is important that we are open and honest about the group, its intentions and its membership. I do not think we can expect to be listened to if we are not prepared to be completely open about what we are doing. My own experience (and I have plenty!) is that the best defence against the extremists in this dirty political arena is one of being completely open and honest about positions and intentions - conspiracies and closed 'behind the scenes' activities are what the extremists get up to (Gladman 1996*b*).

Issues related to the technology to be used in the TTP proposals continued to be discussed on the ukcrypto mailing list some years after the publication the two aforementioned CESG and NHS reports. In particular, it continued to be debated whether the wording of the report indicated that Zergo had indeed advised the NHS to adopt a key escrow scheme. The broadly oppositional stance taken by the members of the Cambridge group tended to be mirrored by other list members. In addition to repeated criticisms of the decision to base systems on the use of the Royal Holloway protocol and Red Pike, there persisted a general feeling that much of what had happened had been 'orchestrated' by CESG, and that it was all a pretence to force the adoption of some form key escrow scheme for the benefit of GCHQ (e.g. Back 1998). However, a lack of concrete evidence on both sides meant that debates often had a rather circular quality.

The TTP policy officially surfaced again in March 1997, when the DTI published a consultation document on the 'Licensing of Trusted Third Parties for the Provision of Encryption Services' (Department of Trade and Industry 1997*a*). In the UK, the purpose of consultation documents (sometimes called Green/White Papers) is to canvass the opinions of industry, academics and the wider public on a policy before attempts are made to draft a bill. Although not a formal part of the UK's legislative process, the consultation process is covered by the Cabinet Office Code of Practice. Though consultation should be viewed as a positive feature of the

legislative process, it is still open to criticism. Common criticisms include the fact that: political pressure often means that there is little time to consult properly; the process may be overly influenced by larger organizations; consultation often focuses on broad issues at the expense of detail; and governments can ultimately ignore the opinions of others and proceed regardless of contrary opinion (Rogers & Walters 2006, p.194). Despite these limitations, the publication of this document still marked the start of a consultation process, during which those opposed to the TTP proposals were able to offer up their expertise.

The timing of this consultation document is perhaps noteworthy, given its close proximity to the 1997 general election. This election resulted in a majority victory for the Labour party, ending an 18 year period in opposition. The Labour victory was significant because their election manifesto had explicitly stated that, if elected, their government would not follow the same path as the US when it came to cryptography regulation:

It is important that privacy is rigorously protected over the new networks, for both personal and commercial reasons. We do not accept the “clipper chip” argument developed in the United States for the authorities to be able to swoop down on any encrypted message at will and unscramble it . . . The only power we would wish to give to the authorities, in order to pursue a defined legitimate anti-criminal purpose, would be to enable decryption to be demanded under judicial warrant (in the same way that a warrant is required in order to search someone’s home) . . . Attempts to control the use of encryption technology are wrong in principle, unworkable in practice, and damaging to the long-term economic value of the information networks. There is no fundamental difference between an encrypted file and a locked safe. A safe may be effectively impregnable in that the effort taken to open it would destroy the contents. An encryption algorithm, similarly, may be effectively unbreakable (Labour Party 1997).

This appeared to suggest that a Labour government would not pursue the Conservative TTP policy. However, it also indicated that they believed that access to keys for law enforcement purposes would still be desirable.

Returning to the government’s consultation paper, which was published before the 1997 general election, the central aim was still to facilitate electronic commerce

without undermining law enforcement. The document reiterated that anyone wishing to provide ‘encryption services’ would be legally required to become a TTP, and would thus require a license from the government. This license would only be granted if the would-be TTP stored a copy of their client’s decryption keys. The public were not required to actually use TTPs, but if they did choose to do so, the government would have access to their decryption keys. The consultation document also stated that there was a possibility that, for ease of access, a central repository would be established that would interface between law enforcement bodies and TTPs. It was argued that, if established, “it is envisaged that it should take no more than an hour for a TTP, once presented with a validated warrant request, to deposit the appropriate client encryption key(s) with a central repository” (Department of Trade and Industry 1997*a*).

The government reiterated that the TTP proposals, although controversial, merely aimed to update the existing laws regarding the interception of communications to incorporate encryption technologies, in line with the Interception of Communications Act 1985:

A critical issue presented by cryptography is the possible conflict between privacy and law enforcement. While the use of cryptography is important for the protection of privacy, it can also be put to improper use such as hiding the illegal activities of criminals and terrorists. Consequently, there is a requirement to establish appropriate mechanisms for lawful access to encrypted information. In the UK, security, intelligence and law enforcement agencies can lawfully intercept communications under certain conditions in accordance with the Interception of Communications Act 1985 (IOCA). Unfortunately, the use of cryptography has the potential to seriously hamper this important law enforcement tool, by making legally intercepted messages unreadable, to the detriment of all law abiding citizens (Department of Trade and Industry 1997*a*).

It was stated that although the UK was in a sense “taking the lead” with these proposals, they believed they were nonetheless in line with EU and OECD guidelines (Department of Trade and Industry 1997*a*).

As it often the case, the consultation document specified topics on which feedback would be most welcome. The consultation document requested responses to the following thirteen issues/questions:

1. Whether the suggested scope of an exclusion from licensing for intra-company TTPs is appropriate?
2. Whether, in the short term, it would be sufficient for business to rely on agreements under contract regarding the integrity of documents and identification of signatures; or whether it would be helpful for legislation to introduce some form of rebuttable presumption for recognition of signed electronic documents?
3. The appropriateness of the proposed arrangements for licensing and regulation.
4. Views on the proposed [licensing] conditions.
5. What if any, specific exemptions for particular organisations offering encryption services would be appropriate depending on the nature of the services offered?
6. Whether it is thought desirable to licence the provision of encryption services to businesses and citizens wholly outside the UK?
7. Should electronic methods for the delivery of electronic warrants by the central repository and the subsequent delivery of keys by the TTP be introduced?
8. Does the legislation specifically need to refer to other forms of legal access including a civil court order for access to cryptographic keys used to protect information relating to civil matters such as bankruptcy?
9. Should deliberate (and perhaps wilfully negligent) disclosure of a client's private encryption key be a specific criminal offence, or would existing civil and criminal sanctions suffice?
10. Whether the principle of strict liability is appropriate in these circumstances?
11. Whether, in principle, an independent appeals body (such as a Tribunal) should be created?

12. Whether the proposed duties of an independent Tribunal are appropriate?
13. Would mandatory ITSEC formal evaluation be appropriate? (DTI 1997b)

These issues were undoubtedly important. But, it is notable that they have little to do with the technical or cryptological side of the proposals, and were instead primarily focussed on legal issues. It may be inferred from this that the DTI remained uninterested in soliciting the views of cryptology experts on technical matters, and instead sought expertise on matters related to law, policy, and electronic commerce. Despite this, the start of the consultation processes signalled the wider involvement of cryptology experts in the political phase. Although some had reacted to previous announcements informally using various methods of on-line communication, and others had responded by writing technical reports, it had largely been a one-way process. The start of the consultation process offered an opportunity for those who felt their expertise had been marginalised to formally introduce it to the political phase.

At the start of the consultation process, the ukcrypto mailing list was a key platform for organizing the activity of those with technical expertise. Charles Lindsey - a computer scientist at the University of Manchester and a key figure in the development of the ALGOL 68 programming language - offered the following advice in a message to the group following the publication of the first TTP consultation document:

I would suggest that for large numbers of people immediately to write off to the DTI at the address given is NOT the right way. We have to learn to do it their way, using their language. And most notably, they are always more impressed by submissions coming from well known individuals, or from bodies that have (or appear to have) some standing and support from within the community. Language must be clear and polite. Sarcasm is out (the driest of dry irony is possibly in). I therefore suggest three phases:

1. Study the document, and establish exactly what it means.
2. Decide which issues are the essential ones to raise.
3. Prepare specific commentaries and counter proposals (Lindsey 1997).

Whether stemming from this message or not, this is quite a good summary of the approach used by individuals to transfer their expertise. The first TTP consultation resulted in 102 responses from organizations - including Intel, Hewlett-Packard and British Telecom - and 158 responses from individuals. This reflected the fact that many felt that the initial policy did not adequately reflect the available expertise. The DTI published a summary of these responses in February 1998. The full responses themselves remain confidential, and were only made available to members of parliament. However, at the time, Hosein (1997) solicited for full responses using mailing lists and Internet searches, and produced a summary report of the 25 responses he received.

Issue	Description
Legal recognition of digital signatures	Almost universal consensus that the legal recognition of digital signatures for the purposes of e-commerce was a positive development
Distinguishing between e-commerce and surveillance.	General consensus that e-commerce and law enforcement are separate issues that require separate policies
Mandatory licensing	Frequently expressed concern over mandatory licensing of TTPs. Some felt that licensing should be voluntary, whilst others felt that licensing should be mandatory, but under different terms.
International consistency	Frequently expressed concern over the commercial impact that would result from the UK having a policy that differed from their trading partners.
Security risks	Concern over the security risks associated with placing keys in a centralised system outside of their owner's control.
Technical specifics	Concern over general lack of technical detail about the proposed scheme, and concern over technical competence where detail was provided.
Criminal circumvention	Concern over the potential for criminals to be able to bypass the system and undermine the law enforcement aspects.

TABLE 8.1: Summary of Responses to the First TTP Consultation (Hosein 1997)

From Table 8.1 it is clear that the legal status of digital signatures and the need to separate e-commerce and law enforcement policies were of primary concern to those who had submitted responses. However, despite not being solicited, comments were also received relating to the technical side of the proposals. Anderson used the consultation exercise to reiterate his view that the use of the Royal Holloway protocol was unacceptable, and to question the way in which the DTI had sourced technical expertise:

The GCHQ protocol is extremely poorly engineered and should not be used. The detailed reasons for this are contained in the paper “The GCHQ Protocol and its problems” which is appended and is hereby included . . . Quite apart from the specific problems with the GCHQ protocol, I would suggest that the DTI acquire some source of technical advice on cryptography that is of substantially higher grade than has clearly been made available to date (Anderson 1997*b*).

Anderson combined these specific technical criticisms, and criticisms of how the DTI had sourced its expertise, with a broader critique of the policy:

The previous government’s proposals for introducing key escrow via a scheme of compulsory licensing of so-called ‘Trusted Third Parties’ are founded on mistaken assumptions. They are probably illegal under European law; they will place a significant cost burden on British business; they will decrease public confidence in information systems; and they attempt to centralise the trust structures in our society in a way that will have many unpleasant consequences – only some of which we can now predict. No case has been made for any law enforcement gain to offset these financial and social costs (Anderson 1997*b*).

Responses from industry tended to be similarly negative, but were usually less damning in their criticism, thus striking a more conciliatory tone. For example, the submission from the computer hardware manufacturer Intel stated:

We doubt that:

1. The proposals represent the surest and least burdensome way of building public trust in the supply of encryption services.
2. There is a strong likelihood of their adoption, in principle at least, across the European Community and in other OECD countries.
3. The proposals would stimulate the growth of electronic commerce in the UK.

4. There is no other viable option for enforcement agencies to have access to encrypted data in readable form and suggest that a fresh look is taken of the issues, with further consultation of interested parties (Intel 1997).

In summing up the responses that he had received, Hosein observed that:

While there is no clear consensus on the outcome of any policy that is to be introduced, what must be acknowledged is that within the UK we have a wealth of knowledge, a wealth of experience, and a wealth of opportunity. Any future development in the field of cryptography policy needs to capitalise on this, embrace the participation and invite the knowledge. We must release the obscurity, stop the scrambling, and introduce the sunlight (Hosein 1997).

This again reflected the view held by many that cryptology expertise had not been adequately drawn upon in the formulation of the policy.

Then, during the first part of the consultation process, there was a rather unexpected development. As was described in chapter 7, there exists strong evidence that the ideas that underpin public-key cryptography were first discovered by researchers working within CESG in the 1970s. These ideas were explored under the heading of work on non-secret encryption. However, all information about this discovery was kept secret from the public. In December 1997, in-between the publication of the first consultation document and the publication of the summary of responses, CESG took the decision to declassify and release information relating to this work.

As was described earlier, five documents were released in total (Communications-Electronics Security Group 1998*b*). The first of these was an historical reflection on the discovery written by James H. Ellis, the individual now credited with the initial idea (Ellis 1987). The release of Ellis' historical account was later followed by the uploading to the CESG website of four further documents related to non-secret encryption: a report outlining Ellis' original idea (Ellis 1970); and three reports outlining Clifford Cocks' and Malcolm Williamson's mathematical implementations (Cocks 1973, Williamson 1974, 1976).

Despite the implications that this might have had on how CESG was perceived during the crypto wars, a full explanation for the release of the documents relating to non-secret encryption was not given. This is perhaps strange given that the release of internal documents by CESG is rare, and questions about CESG's technical competence had already been raised by members of the Cambridge group. Although this revelation may appear to only be tangentially connected to the crypto wars, at the very least it evidenced that CESG had conducted important work on cryptology in secret, and that they believed there would be some value in informing people about this work. It should also be remembered that this revelation occurred during a period of 'offensive' information management by intelligence organizations - based on putting secrets in the public domain on their own terms (Moran 2013). Despite this, when the revealing of non-secret encryption was discussed on the ukcrypto mailing list, the timing of the release was not linked to the wider context created by the TTP proposals.

To sum up, returning to the TTP debate, following the initial announcement, because some felt that their expertise was being ignored, they decided to organize themselves in order to be better able to contribute during the political phase. When the first consultation document was published, an opportunity arose to formally feed their expertise into the political phase. During this part of the consultation process, it became clear that the initial TTP policy did not reflect expert opinion. Despite earlier denials from CESG, it was still believed that the Royal Holloway Protocol and the Red Pike algorithm would be used. The DTI neither officially confirmed nor denied whether this was the case. The expertise of the Cambridge group remained partially marginalized, and does not appear to be reflected in the TTP policy. None of the technical criticisms levelled by the Cambridge Group elicited a response. Additionally, the lack of clarity that plagued the earlier stage of the TTP debate remained.

8.2.4 The Second Stage of Consultation

The consultation process then progressed into a second stage. The publication of the summary of responses to the first consultation was followed by a Secure Electronic Commerce Statement from the new Labour Parliamentary Under Secretary of State at the DTI, Barbara Roche MP (Department of Trade and Industry

1998*a*). Significantly, the statement explained that the mandatory approach to licensing favoured by the previous government would be replaced with a voluntary scheme:

The measures the Government plan to introduce take account of these differing aspects of cryptography and also the responses to the consultation process on the licensing of Trusted Third Parties initiated by the previous Administration. In respect of the latter, the Government has responded to business concerns and criticisms of the previous “mandatory” approach to licensing. Thus, as will be explained below, the new proposals will neither oblige service providers to obtain licences nor to use any particular encryption products or technologies (Department of Trade and Industry 1998*a*).

Though not entirely consistent with the claims made in Labour’s election manifesto, this statement represented a relaxing of the rules surrounding the TTP proposals. The statement also indicated that the Labour government saw digital signatures as separate from the debate, and stated that they would be dealt with using a separate policy (Department of Trade and Industry 1998*a*).

With the consultation process picking up speed, and faint signs that their expertise was being listened to, those opposed to the government’s proposals acted to mount an even stronger presence during the political phase. This resulted in the establishment of the Foundation for Information Policy Research (FIPR) in May 1998. FIPR was (and still is) a think-tank for technology policy, with a particular emphasis on policies related to information technology. Founded by Caspar Bowden, Ross Anderson, Roger Needham, and others, FIPR went on to become an important actor in the crypto wars, and clearly indicated that, rather than fading, the strength of the opposition to the TTP proposals was likely to grow. The FIPR launch press release stated that:

Too often, policy issues relating to information technology are separately debated by two distinct groups: technology experts and those focused on social concerns. Policy makers face the challenge of reconciling the separate debates in areas where technology is often evolving very quickly. [FIPR] aims to provide clear advice that spans this gap

and is independent of vested interests (Foundation for Information Policy Research 1998*a*).

With the founding of FIPR, it was clear that the problems associated with the way the government had in the past sourced their technical expertise had been recognised. FIPR, though made up of scientists and other experts, clearly aimed to have a foot in both the technical and political phases of the crypto wars, and aimed to bridge the gap between the two if the government was unwilling to draw on all of the available technical expertise. The press release also explained how FIPR would be funded:

Microsoft has contributed a six-figure sum to cover the launch costs. Internet Service Providers Poptel and Demon are also providing support. The Foundation's independence will be guaranteed, however, by a board of trustees. In the medium term it will be supported by subscriptions from a range of firms in commerce and industry (Foundation for Information Policy Research 1998*a*).

This highlighted that, as well as failing to draw upon academic cryptology expertise, there was a belief that the government had also neglected to take full advantage of the expertise within industry. Following its inception, many of the public activities of those critical of the government's proposals were co-ordinated through FIPR. In a sense, it can be seen as the next organizational incarnation of the ukcrypto mailing list, in that it often represented the same viewpoint, but was able to do so in a more co-ordinated way. Once established, responses to consultation documents and independent surveys tended to be conducted through FIPR, rather than through individuals.

A second consultation document entitled 'Building Confidence in Electronic Commerce' followed in March 1999 (Department of Trade and Industry 1999*a*). In many ways, the second consultation document mirrored the first. Striking a balance between promoting trust in electronic commerce, and upholding law enforcement capabilities was still seen as the central issue, and the way in which the government planned to legislate for this remained broadly the same. However, the document also reflected the changes outlined in the Secure Electronic Commerce statement. Firstly, it was now proposed that obtaining a TTP license would be

voluntary. The stated reason for this change was that fact that “electronic signatures and encryption will also be provided in ‘closed’ environments where trust already exists between the counterparties and ‘trust’ in the provider may be less important” (Department of Trade and Industry 1999a). Secondly, it was proposed that digital signatures would be treated separately from other cryptographic services. In response to the feedback from the first consultation, it was recognised that digital signatures had a different commercial application than other encryption services. So-called Certification Authorities (CAs) would be responsible for electronic signatures, rather than TTPs, and as such, would be prohibited from escrowing keys and then disclosing them for law enforcement purposes. In separating electronic signatures - typically used for authenticity, from encryption - typically used for confidentiality, the government was beginning to separate electronic commerce from law enforcement.⁷

The government’s proposals then came under the scrutiny of the Trade and Industry Select Committee. The summary of responses to the second consultation document was published just before the publication of the Select Committee’s report (Select Committee 1999). In the UK, there are two main types of Select Committee: Commons Select Committees; and Lords Select Committees. Commons Select Committees are responsible for overseeing and scrutinising the work of government departments. They are made up of government-appointed members of parliament. When a committee meets, they gather written and oral evidence relating to a particular piece of government business, and produce a report of their findings. The government then typically has sixty days to provide a response. From January 26th to March 17th, the Trade and Industry Select Committee met six times to discuss the government’s proposals for ‘Building Confidence in Electronic Commerce’. They examined oral and written evidence from a number of parties from both sides of the debate, many of which submitted evidence during the consultation exercises.

The committee examined written and oral evidence from a wide range of sources, including scientists, lawyers, the Post Office, companies associated with cryptology such as Baltimore (who had merged with Zergo in January 1999), and companies perhaps not traditionally associated with technology, such as Tesco. As with the

⁷As a side issue, although not fully discussed in the first consultation, the second consultation document stated that ITSEC approval would be required for technologies used for TTPs and CAs in order for them to receive a license from the government.

previous consultation exercises, most expressed concerns about the TTP proposals, particularly over what were seen as attempts to promote a system based on key escrow. In his submission to the committee, Ross Anderson drew upon a paper he had recently co-authored with a group of high-profile US cryptology scientists, entitled ‘The Risks of Key Recovery, Key Escrow, and Trusted Third-Party Encryption’ (Abelson et al. 1997). The paper directly challenged the key escrow proposals and the TTP system more broadly. It was argued that:

The deployment of a global key-recovery-based encryption infrastructure to meet law enforcement’s stated specifications will result in substantial sacrifices in security and greatly increased costs to the end-user. Building the secure infrastructure of the breath-taking scale and complexity demanded by these requirements is far beyond the experience and current competency of the field. Even if such an infrastructure could be built, the risks and costs of such a system may ultimately prove unacceptable (Abelson et al. 1997, p.250).

This paper, which clearly had an international focus, and envisioned the integration of national encryption systems within a global network, became the most widely cited on the issue of key escrow, and was also presented as evidence to the Senate during the crypto wars in the US. In their memorandum submitted to the committee, Baltimore saw a role for TTPs in electronic commerce:

There is a clear role for the trusted third party (TTP) within the global electronic market place. Indeed, some global trusted third party agreements are already in place in order to meet existing market demands. Examples include the VISA and MasterCard organisations that set technical standards for bank and credit card systems. They also arrange clearing and settlement between retailers’ banks and credit card issuing banks that would otherwise have no commercial agreement between them. There are many other examples of such “trusted third parties” in the world of automotive component purchasing, shipping and so on. “Trust” is a small but important part of a wider range of services that ensure that such organisations add value to an electronic commerce transaction and thus have a viable business role (Select Committee 1999).

However, they were concerned about some of the potential economic consequences of key escrow, and an overall weakening of security:

A variety of key recovery, key escrow and trusted third party encryption requirements have been proposed, primarily by the police and military services of various countries seeking the ability to monitor and conduct surveillance on illegal transactions which may take place via electronic means. Whilst recognising the needs of law enforcements agencies, we are concerned that the deployment of key-recovery based encryption infrastructures will lead to an increase in costs for the user and / or a weakening of the overall levels of information security (Select Committee 1999).

After hearing the evidence, the report produced by the Select Committee urged that “the Government’s proposals to facilitate trust in electronic commerce must not interfere with existing, and often long-standing, electronic commerce relationships”. Furthermore, that:

The Government’s proposals are tied, perhaps unduly, to the creation of a regulatory regime based on one particular technology - public-key cryptography - and a specific market model, which, although they could be considered attractive at present, may not be optimal bases for electronic commerce carried out over the internet in the future (Select Committee 1999).

Finally, the Select Committee suggested that:

... the Government think twice about the content of its forthcoming Electronic Commerce Bill and only include in the Bill measures which will promote electronic commerce, rather than measures discarded from the previous key escrow policy which are concerned with controlling, not facilitating, electronic commerce (Select Committee 1999).

It was clear from the tone of the Select Committee report that the government’s policy would need to be refined, and in all likelihood, the core assumptions of the

TTP proposals would need to be rethought. By the end of the Select Committee process, those who felt that their expertise had been initially excluded from the political phase had successfully managed to transfer their expertise, and have it feed into formal political processes. However, it was also clear that being successful in this had required substantial effort, and without this effort, their expertise was unlikely to have been heard by the government. Furthermore, the changes to the proposals that this resulted in were modest. Most of the arguments that were made during the consultation process were ignored. In terms of how expertise from the four sites was used, the expertise of the group at Royal Holloway is hardly visible at this stage, either in the form of new technologies that could underpin the TTP system, or in the debates during the formal consultation exercises, select committee meetings, or on the ukcrypto mailing list. In contrast, the expertise of the Cambridge group is highly visible during the consultation process, on ukcrypto, and during the Select Committee process. The arguments made by Ross Anderson in particular appear to define the more general opposition to the government's proposals. In contrast to the earlier stages of the debate, CESG is far less visible. Whilst they did not involve themselves in either the informal or the political debates, their presence was clearly noticeable in their decision to reveal details of their work on non-secret encryption. Again, expertise from NPL does not appear to have been drawn upon. At this point, the conformance of any TTP technologies to ITSEC was a minor concern, as the debate had largely moved away from a technical discussion.

8.2.5 The End of the TTP Debate

Shortly after the Select Committee report was published, the Prime Minister's new Performance and Innovation Unit (PIU) published a report entitled 'Encryption and Law Enforcement' (Performance and Innovation Unit 1999). The PIU - as part of the Cabinet Office - was essentially a high-profile government think tank. Their involvement signalled that the TTP debate was now an important issue for the government. The PIU put together a 'task force' to examine it. Once again, the perceived conflict between electronic commerce and law enforcement was central. However, the report also signalled that the government was losing faith in the TTP concept:

The task force welcomes the intention to include in the Electronic Commerce Bill provisions to allow lawful access to decryption keys and/or plain text under proper authority. The task force also recommended that further attention should be given in the Bill to placing the onus on the recipient of a disclosure notice to prove to the authorities that the requested keys or plain text are not in his possession, and to state to the best of his knowledge and belief where they are (Performance and Innovation Unit 1999).

With hindsight, this report can therefore be seen as marking a transition from the TTP model to a policy that sought direct and lawful access to encryption keys without the use of intermediaries (Hosein 2003).

In July 1999 the DTI published their draft Electronic Commerce Bill (Department of Trade and Industry 1999*b*). The draft bill built upon the commitment made in ‘Encryption and Law Enforcement’ in that it proposed powers to allow for the lawful access of decryption keys under warrant. This meant that, if passed, law enforcement bodies would have the power to ‘force’ an individual to hand over their decryption keys through the threat of further legal action. As such, the law enforcement aspect of the original TTP proposals had been approached in a different way, and was no longer in direct conflict with the desire to promote electronic commerce. With a balance between law enforcement and electronic commerce priorities no longer required, TTPs lost their *raison d’être*.

Although still largely unpopular with those that opposed the TTP proposals, the wider consensus over the change to the law enforcement aspects of the bill, particularly the removal of key escrow, was powerful. The DTI agreed to transfer the law enforcement issues to the Home Office, who were in the process of making amendments to the Interception of Communications Act 1985 and drafting a bill on the Regulation of Investigatory Powers. In the Queen’s Speech of November 1999 - the point at which the government announces the laws they intend to bring forward - the Electronic Communications Bill contained only a sunset clause referring to the establishment of TTPs, which was set to expire if no action was taken before May 2005. This essentially marked the end of the debate over TTPs. FIPR issued a press release claiming that “the ‘crypto wars’ are finally over - and we’ve won” (Foundation for Information Policy Research 2005). This, in a sense, was misleading, as cryptography continued to be thought of by the government

as a threat to law enforcement during the intervening period. When the law enforcement issues were effectively passed to the Home Office, they proceeded with the controversial Regulation of Investigatory Powers bill, which eventually became the Regulation of Investigatory Powers Act 2000 (RIPA). RIPA gave powers to law enforcement bodies to intercept and lawfully decrypt communications under warrant, and although it has been modified slightly since it received royal ascent, the powers it established are largely in effect at present.⁸

8.3 The Debate over Export Controls

Running parallel to the debate over TTPs was another debate that featured cryptology expertise - a debate over export controls. The debate over export controls concerned the best way, given the developments in electronic communications, to manage the intangible transfer of technology by electronic means, without damaging international trade or academic freedom in the process. By the 1990s, the Internet could effectively be used to export software that could be used for military purposes, and it was argued that new powers were needed to regulate it. However, if new powers were too drawn widely, they had the potential to impact negatively upon international trade and academic freedom. During this debate, cryptology was not seen as a potential solution to this problem, but one of the technologies and research fields that might be adversely affected by new laws. As such, cryptology had a different relationship to the debate over export controls than to the debate over TTPs. However, it is worth examining, because it again highlights how those with expertise in cryptology - in particular the Cambridge group - worked to have their expertise heard during the political phase. Given that, as will become clear, the TTP and export control debates were running in parallel, and the debates featured many of the same actors and processes, the export control debate will not be described in as much detail. The main purpose of describing the export control debate, aside from it being an important aspect of the crypto wars, is to further highlight how cryptology experts worked to transfer their expertise to the political phase.

⁸The debate surrounding RIPA was large, detailed and complex, and as a result, would require a separate study to analyze fully.

8.3.1 The Background

As in many other countries, controls exist in the UK to regulate the export of certain goods. Export law is very complex, and what follows is only a brief overview. As a general rule, in most countries, goods designed specifically for military purposes - such as weapons - are usually subject to export controls. Therefore, those wishing to export them require a license from the government. However, goods that are designed for non-military purposes, but could potentially be used for military ends, can be classified as 'dual-use'. Those wishing to export goods that are classified as dual-use also require a license from the government. Additionally, some nations are subject to international sanctions that forbid other nations from exporting weapons or dual-use goods to them.

At the international level, the primary agreement concerned with export was (and still is) the Wassenaar Arrangement (see Evans 2009). The Wassenaar Arrangement - named after the Dutch town in which it was agreed - effectively replaced the Coordinating Committee for Multilateral Export Controls (CoCom) - a Cold War agreement not to export weapons to the Eastern Bloc and other socialist states. Like CoCom, Wassenaar is an export control regime agreed upon by nation states. Wassenaar is not legally binding, but instead provides guidelines for the member states to align their legislation with. According to the Wassenaar website:

The Wassenaar Arrangement has been established in order to contribute to regional and international security and stability, by promoting transparency and greater responsibility in transfers of conventional arms and dual-use goods and technologies, thus preventing destabilising accumulations. Participating States seek, through their national policies, to ensure that transfers of these items do not contribute to the development or enhancement of military capabilities which undermine these goals, and are not diverted to support such capabilities ... The decision to transfer or deny transfer of any item is the sole responsibility of each Participating State. All measures with respect to the Arrangement are taken in accordance with national legislation and policies and are implemented on the basis of national discretion (Wassenaar Arrangement 2013).

Thirty-one countries signed the original Wassenaar Arrangement in 1996. The signatories included most of the world's developed countries. Like CoCom before it, the arrangement classified cryptography as dual-use, and as such, countries that had signed the arrangement, including the UK, licensed the export of cryptography. This meant that UK manufacturers and developers, prior to the start of the crypto wars, required a license to export goods that used cryptography. However, the Wassenaar Arrangement also contained a note on 'General Software'. This note exempted mass-market or public domain cryptography software from export controls. There also existed a personal-use exemption that permitted the effective export of products that accompanied the user (on a laptop, for example) as they crossed borders. However, the arrangement did not address the export of intangible goods, including the export of software using the Internet.

The EU also passed laws about the export of goods. These laws tended to mirror the Wassenaar Arrangement. Amended versions of EU Council Regulation (EC) No. 3381/94 and EU Council Decision No. 94/942/CFSP - both of which came into force in 1995 - regulated the export of dual-use goods (including cryptography) outside of the EU. As with the Wassenaar Arrangement, these rules did not apply to mass-market or public-domain cryptography, and no attempt was made to address the issue of intangible export.

In common with most other countries, the UK aligned its export laws with the Wassenaar Arrangement.⁹ This meant that the export of many cryptography products required a license. Prior to the start of the crypto wars, the Export Control Organization of the DTI was responsible for issuing licenses to those who wished to export cryptography. However, an assessment of the application for an export license was made by CESG. In general, the DTI would follow the advice of CESG when it came to issuing licenses (Parviainen 2000, p.88).

The primary piece of legislation relating to the export of goods from the UK was the Import, Export and Customs (Powers) Defence Act (1939). This act, passed at the outbreak of the Second World War, allowed the Board of Trade to regulate the passage of goods in and out of the country in light of the present emergency. Given the age of the act, and the circumstances under which it was passed, by the 1990s it was seen as out-dated. The emergence of computers and the Internet, allowing as they did the transfer of goods by intangible means, brought this into even sharper

⁹As was alluded to in Chapter 1, the US was the only signatory to deviate from this arrangement.

focus. However, this, according to the government, was not the main impetus for changing the UK's export laws. This came from the so-called Arms-to-Iraq affair, and the resultant 'Scott Report of the Inquiry into the Export of Defence Equipment and Dual-Use Goods to Iraq and Related Prosecutions' (Scott 1996). The Arms-to-Iraq affair concerned the alleged sale of machine tools by the British firm Matrix-Churchill to Saddam Hussein's Iraqi government. The subsequent trial collapsed after it was revealed that Matrix-Churchill had received advice from the government on exporting to Iraq - a country that was then subject to blanket sanctions. As a result, Conservative Prime Minister John Major commissioned an enquiry headed by Lord Justice of Appeal Sir Richard Scott. Scott's 1,800-page report, published in 1996, was heavily critical of the export regime then in place, and recommended a complete reform of government powers to control exports. Scott also criticized the level of government transparency on the issue, and recommended a more open and accountable export regime.

8.3.2 The Announcement of the Export Control Proposals

Following on from this in July 1996, the government published a Green Paper on Strategic Export Controls (Department of Trade and Industry 1996*a*). Responding to the Scott Report, the Green Paper proposed tighter regulation of the export of so-called weapons of mass destruction, the passage of scientific knowledge related to weapons of mass destruction, and the trafficking and brokering of arms deals. Additionally, the Green Paper briefly addressed the new issue of the intangible transfer of technology by electronic means, noting that "there are currently no comprehensive controls on such activities" (Department of Trade and Industry 1996*a*). This concern was not raised in the Scott Report, and had not yet been addressed by any of the relevant international bodies. As such, it was the first public indication that the UK government had identified it as an issue.

As was mentioned earlier in this chapter, a new Labour government replaced the Conservative government after the General Election of May 1997. Shortly afterwards, Foreign Secretary Robin Cook MP issued a statement about the changes to the UK's export regime that the government were proposing. The statement claimed that Labour planned to build on their manifesto commitment to "not permit the sale of arms to regimes that might use them for internal repression or international aggression". Furthermore, that:

The Scott Report on the supply of arms equipment to Iraq revealed the dangers of such decisions being taken in secret. In order that parliament and public can observe that the new policy is being enforced, I am today committing the Government to an annual report on the application of arms exports. An informed public debate is the best guarantee of responsible regulation of the arms trade (Department of Trade and Industry 1997b).

Although the announcement did not specifically mention cryptography, it signalled that the new government were planning significant changes to export control legislation.

In the months after this announcement, changes were also made to some of the aforementioned international frameworks. The transfer of technology by intangible means was slowly emerging as a concern. In October 1997 the EC published ‘Towards A European Framework for Digital Signatures And Encryption’. This document recommended that, when the time comes for the EU to review its dual-use regulation, it could be improved by:

1. Progressively dismantling intra-Community controls on commercial encryption products (i.e. not necessarily for very advanced encryption).
2. Launching a discussion on the scope and interpretation of certain provisions, such as the so-called “General Software Note” (stipulating that public domain software is not subject to controls).
3. Dealing with problems like intangible means of transmission (e.g. transmission of technology by fax or e-mail) (European Commission 1997).

It therefore became clear that the export of intangible goods was something that the EU might address in the future. In December 1998, the Wassenaar Arrangement was amended. Goods employing relatively weak cryptography became free from export regulation, as did cryptography, such as that on DVDs, used to protect intellectual property. However, once again, no mention was made of the export of goods through intangible means.

To sum up, unlike the government’s TTP policy, the government’s policy on export controls did not have cryptography technologies or expertise at its heart. However,

the classification of cryptography as dual-use had long been seen as a matter of concern by those with cryptology expertise. This was particularly true in the period following the release of PGP, and the subsequent attempts to prosecute Phil Zimmerman for exporting it from the US without a license. It's therefore unsurprising that broader changes to export controls were a matter of concern for cryptology experts, and an issue on which they might want their expertise to play a part.

8.3.3 The Consultation Process

In July 1998 the DTI published a White Paper on 'Strategic Export Controls' (Department of Trade and Industry 1998*b*). This marked the start of the consultation process. In the UK, White Papers are more detailed follow-ups to Green Papers, and are produced by the government to gather feedback before drafting a bill. The White Paper on Strategic Export Controls included a foreword from the then president of the Board of Trade, Margaret Beckett MP. In this foreword, the changes to export controls were framed against the backdrop of the Scott Report, and also, the recent technological advances related to electronic communications:

The Government believes that there is also a need to ensure that its strategic export control powers are brought up to date to enable it to deal with modern means of trading, such as transferring information via the Internet, and brokering deals involving the transfer of goods between two other countries (Department of Trade and Industry 1998*b*).

The White Paper then expanded upon how it would deal with the issues of transparency and accountability identified by the Scott Report. It also indicated that the government intended to legislate for what it called the transfer of technology by intangible means. The government believed that the existing Import, Export and Customs Powers (Defence) Act 1939 was outdated because it was worded in such a way that it could only deal with tangible transfer. The government proposed that, as was the case with tangible transfer, the export of goods through intangible means should require a license from the DTI. The government therefore aimed to address the issue of export using email and fax:

Given the ever increasing ease with which information can be transferred across national boundaries by electronic means, i.e. by fax or e-mail, the Government proposes to provide that documents transferred abroad containing controlled technology should be subject to export licensing requirements, whether exported physically or in electronic form (Department of Trade and Industry 1998*b*).

The government also reserved additional proposals for the dissemination of goods and materials related to weapons of mass destruction. Here, the proposals also considered information uploaded to the World Wide Web:

The Government has also given consideration to possible controls on the publication of controlled technology on electronic networks such as the World Wide Web (WWW). Even the comprehensive controls on electronic transfers proposed above would not cover a situation in which sensitive information (which, if exported, would require a licence) was posted on the electronic networks (at which point it would move into the public domain) where it would become accessible to potential proliferators. A possible solution to this would be to add a provision to the weapons of mass destruction-related offences proposed in section 3.1 above, making publication of controlled technology relevant to the development of weapons of mass destruction an offence. This would apply whatever the medium of publication (Department of Trade and Industry 1998*b*).

But most drastically, the proposals also considered the regulation of information related to weapons of mass destruction that was 'exported' through oral communication:

Information can also be passed on in non-documentary form (e.g. orally or through personal demonstration). The proposal to make it an offence to do something which it was known or suspected could assist a weapons of mass destruction or long range missile programme, described in paragraph 3.1.4, would catch transfers of information in non-documentary form. This offence would be implemented under the power to control the transfer of technology by any means. While this

power would enable the Government, if need arose, to introduce the same controls on other types of technology, we propose for the time being, to limit this wider offence to technology related to weapons of mass destruction and long-range missiles. The Government considers that it is right that controls on the transfer of information orally or through personal demonstration should be limited to the areas of greatest concern, in view of the difficulties of licensing such transfers, both for applicants and for the licensing authority, and given also that there are sensitivities in relation to free speech and academic freedom (Department of Trade and Industry 1998*b*).

Clearly, then, the government saw a wide remit for their new export control regime, and it was this broad scope that became the most controversial element of their proposals.

The DTI published two further consultation documents requesting feedback on their export control proposals. The first received 38 responses. The second received 54 responses. Just over half of these responses were from industry and the rest were from Non-Governmental Organizations (NGOs). However, the government did not publish a summary of these responses, and only those voluntarily published by the author are available to view.

Although the White Paper did not specifically mention cryptography or encryption, some of the responses identified this as an issue, given the belief that it would be affected by the proposals. FIPR submitted a response entitled ‘Strategic Export Controls: The Impact on Cryptography’ (Foundation for Information Policy Research 1998*b*). Here, they argued that commercial uses of cryptography should be separated from military uses:

Military cryptography tends to be custom built and expensive. It could be distinguished in law from the new growing area of commercial civil cryptography. But the rules make no such distinction. The controls apply, contrary to the spirit of Wassenaar, to what is emerging as a vital area for civil commerce (Foundation for Information Policy Research 1998*b*).

FIPR also speculated on how the proposed legislation would affect research into cryptography, as well as other scientific fields:

The new proposals would effectively cover all cross-border research into cryptography. As proposed, they would make an export licence necessary for every fax and email on the subject (and those granting the licence would hardly be likely to understand the messages in question). It might control the education of non-UK residents including, for example, three quarters of Cambridge science and technology research students. UK participants would need a licence before submitting a contribution about cryptography to Internet mailing lists and news groups (Foundation for Information Policy Research 1998*b*).

Here, then, FIPR were using cryptography as an example to illustrate wider negative implications for the education of, and research contribution of, foreign students.

As had occurred with the TTP proposals, on November 10th 1998 the White Paper on Strategic Export Controls was scrutinized by the Commons Trade and Industry Select Committee (Select Committee 1998). Given the primary goals of the White Paper, most of the written and oral evidence that the committee heard came from defence manufacturers and anti-war/weapons NGOs. Perhaps unsurprisingly, defence manufacturers - such as the Society for British Aerospace Companies - tended argue that the proposals, particularly in relation to intangibles, went too far, and anti-war NGOs - such as the Campaign Against Arms Trade - while generally supportive, tended to think that the proposals didn't go far enough.

A written memorandum was submitted by Ross Anderson in the form of a letter to the Secretary of State for Trade and Industry, Peter Mandelson MP. Anderson reiterated the arguments made in the earlier FIPR consultation response. Anderson argued that the proposed licensing scheme for intangible export would have a negative impact upon the science and technology departments of UK universities:

The majority of our research students, being foreign nationals, would appear to require personal export licenses in order to get access to high-tech equipment they use routinely in their work. The proposed requirements are likely to result in a considerable waste of time and of public money, and will give enormous scope for acrimony. They will also impose significant costs on our high-tech industries and harm our collaboration with them (including collaborations funded by the DTI).

Their effect will be to undermine the DTI's relationship with the UK science and technology community (Anderson 1998).

However, in this instance, cryptography technologies were only referred to briefly. They were listed as just one of many examples of science and technology research fields that would be affected.

The report published by the Select Committee, whilst acknowledging that there was no reason for intangible transfer to be completely exempted, expressed serious doubts about the viability of seeking to license in this way:

There is an irrefutable logic in the proposal which nearly all involved accept; there is no argument in favour of the principle of explicitly exempting such transfers from the licensing regime. Grave doubts have however been expressed as to the practicality of the proposals and fears raised as to their consequences if implemented (Select Committee 1998).

Their report concluded that there could be no question of licensing export through intangible means until a consensus was reached. The report also dealt with some of the concerns raised in the responses to the White Paper that were specific to cryptography, and linked them to attempts during the crypto wars in the US to seek to limit its public use:

It was also implied by some respondents to the White Paper that the proposed extension of controls was a somewhat sneaky attempt to control the transfer of "strong" cryptography. The USA has apparently tried and failed to prevent such transfers. It has also been implied that the new controls sought will be unenforceable unless the Government's agency has access to the necessary decryption, and that this will be used as a justification to obtain sweeping decryption powers: an implication firmly rejected by the Minister (Select Committee 1998).

It's clear from this that the export control debates and the TTP debates are were linked in the minds of those involved. It was also clear that, as in the TTP debate, the Select Committee had serious concerns about the regulation of cryptology.

As well as contributing to the consultation processes formally, those broadly opposed to the government's proposals also aimed to be heard through other channels. In 1998 and 1999, Ross Anderson co-authored a book called *The Global Trust Register* (Anderson et al. 1999). This book was ostensibly a directory of important public keys that could be used to encrypt communications. At the time, the export regulations surrounding cryptography prohibited the book from being legally published electronically, so instead, the rather cumbersome keys were distributed in print. Although the stated purpose of the book was “to cut through the current chaos of public key certification by printing the important keys in a widely distributed paper book and thus providing a kind of phone book or trade directory for security on the net”, it could also be used to impart a clear political message. As the preface to the 1999 edition stated, “the 1998 edition of this book already played a role in history: the British Government decided to tone down the previous government's proposed legislation on cryptography after we visited the minister responsible for culture, gave him a copy and complained that the new law would result in it being banned” (Anderson et al. 1999, pp.viii-ix).

In summary, during the consultation process, the expertise of the Cambridge group, particularly Ross Anderson, was clearly visible. This expertise was again used to form an opposition to the government's proposals. Although it would be difficult to argue that this expertise was ignored in the same way as it was during the TTP debate, the processes used to transfer it to the political phase were the same, and it is unlikely to have been heard if these processes were not utilized. Expertise from CESG, though absent from the visible consultation process, was used during the process of granting licenses for export under the previous export regime, and nothing in the proposals suggested that this would change. Expertise from Royal Holloway and NPL does not appear to have played any part in the debate over export controls, though when the matter was discussed during interviews, a general view was expressed that the export of cryptography should not be licensed.

8.3.4 The Export Control Bill

On March 29th 2001, the government published its draft Export Control and Non-Proliferation Bill (Department of Trade and Industry 2001). The bill was eventually renamed to the Export Control Bill, and was introduced to the House

of Commons on June 26th, 2001. Despite the concerns that were raised during the consultation process, the bill contained clauses to license the export of goods by intangible means. Section 2(6) of the Export Control Bill stated that:

In this Act -

“transfer”, in relation to any technology, means a transfer by any means (or combination of means), including oral communication and transfer of goods on which the technology is recorded or from which it can be derived, other than the exportation of such goods;

“technology” means information (including information comprised in software) that is capable of use in connection with -

1. The development, production or use of any goods or software;
2. The development of, or the carrying out of, an industrial or commercial activity or an activity of any other kind whatsoever (Department of Trade and Industry 2001).

In response, through FIPR, Anderson and others campaigned for an amendment relating to academic freedom when the bill reached the House of Lords (Foundation for Information Policy Research 2003). The Lords placed this section under scrutiny, and eventually proposed that a clause be added to the bill that exempted the transfer of any information that was in the public domain, and the transfer of information “in writing or electronically in the ordinary course of academic teaching or research”. This amendment was approved, and the bill was sent back to the House of Commons. The Export Control Bill - complete with this amendment - became the Export Control Act 2002 in the July of that year. As a result, the campaigning of FIPR and the Cambridge group succeeded in narrowing the scope of the government’s export controls, and succeed in limiting their impact upon academic freedom. At present, the export of cryptography is still regulated in accordance with EU regulations, the Wassenaar Arrangement, and the Dual-Use Items (Export Control) Regulations 2000. These regulations cover both the tangible and intangible export of goods. To be clear, the amendment approved by the House of Lords, only exempted cryptography if bound up with academic freedom.

8.4 Conclusion

My description of the political phase of the crypto wars in the UK described two debates over cryptology: a debate over TTPs; and a debate over export controls. Although distinct in terms of the legislative process, and their relationship to cryptology expertise, they overlapped in terms of some of the issues discussed and the actors involved.

The debate over TTPs was framed by the government as a debate over the best way to promote electronic commerce by using cryptography to make electronic transactions safer, whilst at the same time retaining the ability of law enforcement bodies to monitor communications for the purposes of preventing crime. Eventually, the government abandoned these proposals in favour of a system that permitted lawful government access to keys under the Regulation of Investigatory Powers Act 2000. Expertise enacted by the Information Security Group at Royal Holloway was used to develop technologies that could be used in the proposed TTP system, and which were also to be used in a system to secure government emails, and patient data within an NHS network. Expertise enacted by the Security Group at the University of Cambridge was used to construct arguments to oppose the government's proposals. These arguments were deployed informally on the ukcrypto mailing list, and formally during consultation processes and Select Committee Meetings. Expertise enacted at CESG was used to inform the government's TTP proposals, and to assist with the development of the government email and NHS systems. Furthermore, CESG expertise was used to develop technologies that could be used in these systems. Expertise enacted within the Data Security Group at the NPL does not appear to have been drawn on.

The debate over TTPs can also be understood in terms of disagreements over how the government had arrived at their proposals, and thus the motivations behind them. It was felt by some - particularly those working within the Cambridge group - that their cryptology expertise had been ignored during the formation of the proposals. There also existed a view that the expertise that had been used - particularly that of CESG - was not up to the standard required to address the complexities of the TTP issue. Furthermore, the lack of detail in the proposals, particularly the lack of clear information about the nature of the involvement of CESG (and thus, possibly GCHQ), fuelled speculation about whether the proposals were a way of introducing a key escrow-based system that intelligence agencies

had ultimate control over. Given that their expertise was being ignored, those opposed to the proposals organized themselves in order to mount a stronger opposition. Despite the fact that many of their arguments continued to go unanswered and largely unrealised in the proposals, those opposed to them consistently attempted to engage with the formal processes within the political phase, in order to get their expertise heard. When the TTP proposals were eventually dropped in 2005, FIPR clearly believed that they had been instrumental in achieving this, given that they claimed to have “won the crypto wars” (Foundation for Information Policy Research 2005).

The debate over TTPs unfolded in parallel with a debate over export controls. The debate over export controls was framed by the government as a debate over the best way, given the developments in electronic communications, to manage the export of technologies through intangible means - like cryptographic software - that could be used to threaten national security, without damaging trade or academic freedom in the process. Although the expertise enacted by Royal Holloway and NPL was absent from this debate, expertise enacted by the Cambridge group was again used to construct an opposition to the government’s proposals. This expertise was later used to underpin arguments about the wider issues of academic freedom and the state of scientific research in the UK. Expertise enacted by CESG was used by the DTI in the process of granting licenses under the existing export control regime - a process that continued following the Export Control Bill.

As in the debate over TTPs, the cryptology expertise of those opposed to the proposals was initially ignored. However, in contrast to the debate over TTPs, the debate over export controls did not revolve around technical questions related to cryptology. As such, the way in which the DTI sourced its cryptology expertise was not as important. However, generally speaking, those opposed to the TTP proposals were also those opposed to the export control proposals, and the way in which they made their expertise known to policymakers during the political phase was much the same. Efforts were organized through ukcrypto and FIPR, and arguments were made during consultation processes, and through engagement with the press. These efforts could be described as successful in that they succeeded in preserving a particular strand of academic freedom. Therefore, aside from their relationship to cryptology, the debate over TTPs and the debate over export controls have in common the fact that their initial proposals were not based on the best available expertise, and if it were not for a considerable effort on the part

of a determined group of individuals, it is likely that the UK law would now be different.

Chapter 9

Multiplicity and its Consequences

9.1 Introduction

This chapter will be split into two halves. In the first half, I will describe how the third wave concept of sociological discrimination can be used to demonstrate that multiple historical and institutional research practices enacted multiple cryptology expertises during the technical phase of the crypto wars. The process of sociological discrimination described will be based on Philippe Larédo and Philippe Mustar's criteria for assembling laboratory activity profiles. Moving from the specific descriptions of expertises described in chapters 4 through 7, I will categorise the activity of each research group in terms of the more general categories of: the production of certified knowledge, education and training activities, public research and the innovation process, the participation in public or collective goods and finalities, and public debates about science and technology. In the second half of this chapter, I will describe the consequences that the multiplicity of contributory expertises had for the political phase of the crypto wars. I will describe how expertises were used for different and often distinct purposes during the political phase, and that these uses can be linked to the nature of the practices that enacted them during the technical phase. I will also begin to describe in more detail how the political phase came to know about the expertise that was enacted during the technical phase. I will bring all of this information together to highlight some of the contradictions and problems that acknowledging multiplicity brings to the fore, and suggest some ways in which these might be addressed. In particular, I

will describe what I have called the ‘Minimum Transfer Requirement’. This requirement aims to specify the minimum conditions under which it can be claimed that the political phase has considered the available expertise enacted during the technical phase. I will also identify what I have called the ‘problem of expert discrimination’. The problem of expert discrimination acknowledges that it will not always be possible for domain-specific discrimination to be adequately carried out or to be public demonstrated, and therefore suggests that in order to preserve a clear separation between the technical and political phases, there are circumstances under which expertise from the technical phase should not be transferred to the political phase.

9.2 Contributory Expertise During the Technical Phase

As was described in chapter 2, the periodic table of expertise, elective modernism, and the third wave more generally, does not account for the different types of contributory expertise that can exist within a discipline. However, chapters 4 to 7 have pointed towards some of the ways in which contributory expertise can vary according to the multiple institutional and historical research practices used to enact it.

To recap briefly, in chapter 4, it was described how the Data Security Group at NPL enacted their cryptology expertise during the technical phase of the crypto wars. Due to the increasing influence of New Public Management, the activities of the Data Security Group were initially focussed on the production of standards rather than original, open-ended research exemplified by Donald Davies’ pioneering projects. By the 1990s, increased financial pressures were forcing management at NPL to place even greater emphasis on sourcing work using the customer-contractor model. Between 1992 and 1994, the budget for standards research was drastically cut, and almost all IT standards-based work was terminated. Within the Data Security Group, this was replaced with work on the accreditation of testing centres for the UK ITSEC scheme. NPL became ‘Government Owned, Contractor Operated’ in 1995, resulting in fewer members of staff and a reduced work programme. The Data Security Group initially survived this change, but

with an increased emphasis on formal methods, strict conformance testing, and accreditation.

In chapter 5, it was described how the expertise of the Information Security Group at Royal Holloway was enacted during the technical phase through collaborative projects with industry and government. In particular, the group maintained a long and fruitful relationship with Racal-Comsec. We saw how key members of the group became involved in cryptology through consulting for Racal-Comsec, and how - when the research group at Royal Holloway was established - a clear emphasis was placed on acquiring funding through industrial collaboration. The other distinctive feature of the expertise enacted by the group was its mathematical nature. Many of those working within the group had backgrounds in mathematics, and much of the academic output of the group tended towards the pure or the theoretical. Finally, it should also be remembered that the group also ran an MSc course in Information Security that aimed to train students for employment as industrial data security managers.

In chapter 6, it was described how the cryptology expertise of the Security Group at the University of Cambridge was enacted through practices based on interdisciplinary methods designed to probe the real-world use of cryptology systems. Research into computer security began with Roger Needham's work on the Needham-Schroeder protocol and BAN logic, and flourished as the laboratory expanded through the 1980s. The key actor during the 1990s was Ross Anderson. By the time the UK crypto wars were underway in 1996, Anderson had already published work asking why crypto systems fail - based on studies of how systems behaved in the real-world, and the consequences that this might have for systems design. It should also be remembered that, again, this was only one strand of the work of the group. They also ran a doctoral programme for the training and education of PhD students, developed technologies, collaborated with industry, and carried out a number of other activities one would typically associate with a university department.

In chapter 7, a description was provided of the cryptology expertise enacted by CESG. CESG enacted their expertise through pioneering research into cryptology in the 1970s under the heading of non-secret encryption, and later, through the development of cryptographic algorithms such as Red Pike. CESG also advised governments on the use of cryptography for their communications systems, and

managed schemes such as ITSEC that approved security products and technologies developed by others. However, these specific details also fail, in important ways, to paint a complete picture of CESH's expertise during the technical phase. This is because they capture little of the style in which CESH went about their activity. CESH, as a branch of GCHQ, was very secretive. Because of the links between CESH and the intelligence world, details about much of what they did - including basic details such as the names of their employees and the results of their research - were not made public. This created methodological difficulties for my research. Analysable documents were very hard to come by, and interviewees were either impossible to locate or unwilling or unable to participate. However, the methodological difficulties encountered should be seen as being caused by the very same practices that enacted their expertise. On this understanding, it can be argued that CESH enacted a secret cryptology expertise.

9.2.1 Identifying Multiplicity Using Sociological Discrimination

Though the third wave does not distinguish between different types of contributory expertise, it may be possible to use a form of sociological discrimination to arrive at an appreciation of it. In discussing sociological discrimination, Weinel argued:

The application of STS methodologies and theories to instances of science under Wave 2 has resulted in a particular understanding of the nature of science. It is argued here that it is this particular understanding of the nature of science that constitutes a specific type of transmuted meta-expertise, when it is used to inform science-related judgements such as judgements of the authenticity of scientific controversies or, which is also possible, of the credibility of particular scientific claims (Weinel 2010, p.198).

Weinel claimed that, because STS had developed an understanding of the nature of scientific controversies, this constituted a form of meta-expertise that could be used to arrive at a judgement of their authenticity. I argue that the same reasoning can be applied to other aspects of the nature of science that STS has developed an understanding of. Work within STS is undoubtedly sensitive to difference within

science. Therefore, it should be possible to use existing concepts within STS to underpin sociological discrimination. This subsection will therefore consider some of the available options.

Other strands of STS have attempted to deal with difference and variation within scientific disciplines. Tony Becher and Paul R. Trowler (2001), in their broad-based study of the nature of academic disciplines, have drawn attention to the myriad specialities that can exist within them. In particular, they noted that “there is no single method of enquiry, no standard verification procedure, no definitive set of concepts that uniquely characterizes each particular discipline” and that it is “in some contexts more meaningful to speak about the identifiable and coherent properties of subsidiary areas within one disciplinary domain or another” (Becher & Trowler 2001, p.65). Though the study of disciplines and specialization within science was an important area of enquiry for STS scholars in the 1960s and 1970s, studies were mainly concerned with developing a sociological understanding of how specialisms were formed, and did not aim to develop logical categories or types (Wray 2005). However, some studies have still touched on these ideas. For example, Karin Knorr Cetina (1999) would later examine “machines of knowledge construction” in high-energy physics and molecular biology to argue that the differences in the way they produce scientific knowledge revealed disunity amongst the sciences, and the possibility for disunity within scientific disciplines.

Given that it is rooted in practices, an alternative way of approaching an understanding of different types of contributory expertise may be derived from a consideration of different laboratory ‘types’. As Arjan van Rooij (2011) has observed, past descriptions of laboratory activity, particularly those of a historical nature, are predominantly based on contrasting laboratories situated within a university with those situated within an R&D environment:

Taking a broad view of the historical literature makes clear that several types of laboratory have existed side by side. Yet the literature is fragmented. Typically, one type of laboratory is central to a particular strand of the literature. Crucially, academic and research and development (R&D) laboratories have attracted much more attention than other laboratories in business and government or laboratories run as

stand-alone businesses. Only very few studies have pursued a comparative perspective, but these studies typically focus on only a few countries and only a few types. Because of the fragmentation, it seems that one type of laboratory has replaced, or has dominated, other types, and that one particular type embodies a superior mode of knowledge production. Particularly the history of R&D laboratories has been written from the perspective that such labs are inevitable for (big) business (van Rooij 2011, pp.427-428).

This can create problems for studying the multiplicity of contributory expertise produced in laboratories, because in the absence of an integrated account of different types, the understanding of university laboratories has dominated:

The university laboratory often appears as the basic type from which other labs are derived. In this perspective, research is equated with university science, but the context of business or government changes it. A strong current in the literature views R&D labs as halfway between the university laboratory and the world of business, as academic research with another goal (van Rooij 2011, p.428).

This also serves to reinforce a naive binary distinction between academic laboratories that are concerned with knowledge, and commercial laboratories that are concerned with profit. As a way of beginning to address this problem, van Rooij attempted to differentiate between eight different types of laboratory (see Table 9.1).

Laboratory Type	Origin	Features
University	1800s	Linked to science education
Works	1800s	Linked to production plants
Company	Late 1800s	Mix of testing and consulting
Internal Government	Late 1800s	Service function for government departments
Regulative Government	Late 1800s	Support of enforcement of standards
R&D	Late 1880s	Improvement of firm's position
Normative Government	After 1900	Support of policy goals
Research Association	1910s	Organized co-operation

TABLE 9.1: Eight Types of Scientific Laboratory (van Rooij 2011)

Van Rooij (2011) classified laboratories, in the first instance, based on the type of knowledge produced, their orientation, and their ownership. In terms of knowledge

production, van Rooij claimed that laboratories can be concerned with research ('why' questions), development ('how' questions), or testing ('what' questions). Though these classifications perhaps lack nuance in that they fail to account for the blurring of different activities, they do provide an interesting starting point for thinking about laboratories and the enacting of multiple contributory expertises, given that, at the very least, they acknowledge that different types of laboratory can produce different types of knowledge. Furthermore, van Rooij (2011) linked his types of knowledge production to historical and institutional practices, thus acknowledging the possible causal relationship between the two.

In a similar vein, Philippe Larédo and Philippe Mustar (2000) have attempted to assemble 'activity profiles' for scientific laboratories. Their starting point was the 'research compass card model' - co-developed with Michel Callon (Larédo et al. 1992). This model identified five categories of activity that scientific laboratories typically engage in. They were:

1. The production of certified knowledge: the production of open scientific knowledge that is reviewed by colleagues and resistant to controversies;
2. Education, training activities and embodied knowledge: equipping scientists with the tacit knowledge required to undertake laboratory work;
3. Public research and the innovation process: the creation of competitive advantages and products through innovation.
4. The participation in public or collective goods and finalities: the production of standards and research that can be used by the public sector (defence, health, etc.);
5. Research and public debate about science and technology: contributing to controversies or public debates about the role of science and scientific research (Larédo & Mustar 2000, pp.517-521).

Larédo and Mustar (2000) used the research compass card to construct activity profiles for laboratories based on a quantitative analysis of the extent to which they carried out each activity. They based their assessment of whether laboratories engaged in the production of certified knowledge through an assessment of: the average amount of academic articles produced by each researcher; indirect forms

of individual recognition (such as conference organization, conference initiations, prizes, and membership of journal editorial boards); indirect forms of laboratory recognition (such as participation in national and international programmes); and official national forms of recognition. They based their assessment of education and training on: the number of postgraduate students in the laboratory; and the number of staff qualified to supervise postgraduate students. They based their assessment of public research and innovation on: the level of industrial funding; and the nature of activities the laboratory entered into. They based their assessment of the production of public or collective goods on: official participation in national and international programmes; and the importance of national priorities to individuals at the laboratory. Finally, they based their assessment of involvement in public debate on: the involvement of the laboratory in controversies over technical decisions; and the involvement in controversies over research policy.

As Larédo and Mustar (2000) acknowledged, it is clear that any approach that attempts to use metrics to arrive at an assessment of the degree to which a laboratory engages in, say, the innovation process compared to the production of certified knowledge, will contain flaws. Nonetheless, it demonstrates that, in principle, the research compass card model can be used as a useful way of operationalizing sociological discrimination to identify how practices at laboratories can vary. With this in mind, I have decided to use the research compass card as a starting point for developing a qualitative understanding of the work of the four laboratories described in the previous chapters, and thus, the expertise they enacted.

I will now attempt to use this to characterize the activity of the research groups described in previous chapters. I will do this through the aid of a thematic conceptual matrix. More specifically, I will use the approach to the thematic conceptual matrix described by Stuart Henderson and Eden Segal (2013). Under this approach, the first column of the matrix states the principle, the second column defines that principle, the third column states the individual elements that make up the principle, the fourth column states whether each element was confirmed or refuted by the data, and the fifth column uses the data from the fourth column to arrive at an overall judgement of whether the principle itself was evident (see Tables 9.2-9.5). Based on the descriptions that I have developed in earlier chapters, I will assign a category of either: confirming evidence; partial confirming evidence;

refuting evidence; or unknown, to each principle.¹ The advantage of this approach is that it is descriptive enough to show the disparities sometimes evident in the data, whilst also eschewing inappropriate or unsuitable attempts at quantification.

During the technical phase of the crypto wars, the Data Security Group at NPL was most closely aligned with the van Rooij's 'regulative government' classification. However, as van Rooij (2011, 2013) specifically pointed out, more broadly the laboratory can probably be seen as a lying somewhere between the 'regulative government' and 'normative government' laboratory types. In terms of their activity during the technical phase, the group were probably most closely aligned with 'public research and innovation' and 'participation towards public or collective goods and finalities' on the research compass card, but were also partially engaged in the 'production of certified knowledge' and 'education and training activities' (see Table 9.2).

The Information Security Group at Royal Holloway is clearly most closely aligned with van Rooij's 'university' laboratory classification. However, this alone tells us little about the nature of the work that the group undertook. In terms of activity during the technical phase of the crypto wars and the research compass card, the group were engaged in work characteristic of the 'production of certified knowledge', 'education, training activities and embodied knowledge', 'public research and the innovation process', and 'the participation to public or collective goods and finalities' (see Table 9.3).

The Security Group at the University of Cambridge is also clearly aligned with van Rooij's 'university' classification. In this case, the 'university' classification provides a good starting point for thinking about the research practices that the group developed during the technical phase of the crypto wars. In terms of the research compass card, the group were engaged in work characteristic of the 'production of certified knowledge', 'education, training activities and embodied knowledge', 'research and public debate about science and technology', and partially characteristic of 'public research and the innovation process' (see Table 9.4).

CESG appears to be closely aligned with van Rooij's 'internal government' classification. However, CESG had an unusual role, and does not fit easily into this

¹To be absolutely clear, the category of 'partial confirming evidence' is applied to situations where there exists confirming evidence of partial involvement in an activity, rather than partial evidence.

classification scheme. The group could equally be slotted into van Rooij's 'regulative government' and 'normative government' classifications. In terms of activity during the technical phase of the crypto wars and the research compass card, there is evidence that the group carried out some work related to the 'participation to public or collective goods and finalities', but this makes for an awkward fit. This perhaps reflects the fact that the scientific work of intelligence organizations is rarely discussed. However, what is most noticeable about the attempt to characterize the work of CESH is the amount of information that is unknown (see Table 9.5).

This information can be summarised in a final table (see Table 9.6). This method of operationalizing sociological discrimination is neither flawless, nor the only way of going about the task. Whilst it does not capture the intricacies of the descriptions provided in chapters 4 through 7, or identify a clear descriptive label for the expertise enacted in each case, it makes up for this by associating the expertise with general categories that may be present in other fields of scientific research. Furthermore, the way in which research practices (and thus contributory expertise) has been characterized clearly demonstrates divergence, and thus multiplicity. In addition to what it reveals about conformance to the activities listed on the research compass card, it is also able to show instances where information about certain practices is absent.

This suggests additional ways in which sociological discrimination could be used in real-time during the political phase of a controversy over technological decision-making. For example, assuming that the data is readily available, it could be used during the political phase to make decisions about where contributory expertise should be sought. Given that a controversy can hinge upon a wide range of issues, there may be a disparate range of propositional questions that require answers. Though a laboratory that specializes in education and training activities may be said to possess contributory expertise in a field, it is unlikely to be the best source of contributory expertise during a controversy that hinges on questions about the appropriateness of a particular technology. Thus, recognising that the contributory expertise produced within a discipline is multiple, and that this multiplicity can be identified and characterized through sociological discrimination, may allow those active during the political phase to make better use of the expertise available from the technical, given that it offers an alternative to basing decisions on

Principle	Definition	Elements	Evidence of Elements	Evidence of Principle
1. The production of certified knowledge	The production of open scientific knowledge that is reviewed by colleagues and resistant to controversies	1.1 Production of academic articles 1.2 Indirect individual recognition (e.g. prizes, conference invitations, editorial board membership) 1.3 Indirect laboratory recognition (e.g. participation in national and international programmes) 1.4 National recognition (e.g. RAE)	✕ ✕ ✕ X X	✕ ✕ ✕ X ✕
2. Education, training activities, and embodied knowledge	The equipping of scientists with the tacit knowledge required to undertake laboratory work	2.1 Presence of postgraduate students	X	✕
3. Public research and the innovation process	The creation of competitive products and advantages through innovation	2.2 Presence of individuals qualified to supervise postgraduate students 3.1 Receipt of funding from industry. 3.2 Patenting of work 3.3 Direct commercialization of work 3.4 Provision of consulting services to industry	✕ ✓ ✓ ✕ ✓	✓ ✓ ✕ ✓
4. Participation in the production of collective or public goods	The production of standards and research that can be used by the public sector	4.1 Participation in official national or international programmes or schemes	✓	✓
5. Research and public debate about science and technology	The contribution to controversies or public debates about the role of science and scientific research	5.1 Production of knowledge related to social impact of scientific work 5.2 Production of policy arguments based on production of certified knowledge	X X	X X

✓ - Confirming Evidence
 ✕ - Partial Confirming Evidence
 X - Refuting Evidence
 ? - Unknown

TABLE 9.2: Research Practices of the Data Security Group at NPL

Principle	Definition	Elements	Evidence of Elements	Evidence of Principle
1. The production of certified knowledge	The production of open scientific knowledge that is reviewed by colleagues and resistant to controversies	1.1 Production of academic articles 1.2 Indirect individual recognition (e.g. prizes, conference invitations, editorial board membership) 1.3 Indirect laboratory recognition (e.g. participation in national and international programmes) 1.4 National recognition (e.g. RAE)	✓ ✓ ✓ ✓ ✓	✓
2. Education, training activities, and embodied knowledge	The equipping of scientists with the tacit knowledge required to undertake laboratory work ✓	2.1 Presence of postgraduate students	✓	✓
3. Public research and the innovation process	The creation of competitive products and advantages through innovation	2.2 Presence of individuals qualified to supervise postgraduate students 3.1 Receipt of funding from industry. 3.2 Patenting of work 3.3 Direct commercialization of work 3.4 Provision of consulting services to industry	✓ ✓ ✓ ✓ ✓	✓
4. Participation in the production of collective or public goods	The production of standards and research that can be used by the public sector	4.1 Participation in official national or international programmes or schemes	✓	✓
5. Research and public debate about science and technology	The contribution to controversies or public debates about the role of science and scientific research	5.1 Production of knowledge related to social impact of scientific work 5.2 Production of policy arguments based on production of certified knowledge	X X	X

✓ - Confirming Evidence
 ✕ - Partial Confirming Evidence
 X - Refuting Evidence
 ? - Unknown

TABLE 9.3: Research Practices of the Information Security Group at Royal Holloway

Principle	Definition	Elements	Evidence of Elements	Evidence of Principle
1. The production of certified knowledge	The production of open scientific knowledge that is reviewed by colleagues and resistant to controversies	1.1 Production of academic articles 1.2 Indirect individual recognition (e.g. prizes, conference invitations, editorial board membership) 1.3 Indirect laboratory recognition (e.g. participation in national and international programmes) 1.4 National recognition (e.g. RAE)	✓ ✓ ✓ ✓ ✓	✓ ✓ ✓ ✓ ✓
2. Education, training activities, and embodied knowledge	The equipping of scientists with the tacit knowledge required to undertake laboratory work	2.1 Presence of postgraduate students	✓	✓
3. Public research and the innovation process	The creation of competitive products and advantages through innovation	2.2 Presence of individuals qualified to supervise postgraduate students	✓	✓
4. Participation in the production of collective or public goods	The production of standards and research that can be used by the public sector	3.1 Receipt of funding from industry.	✗	✗
		3.2 Patenting of work	✗	✗
		3.3 Direct commercialization of work	✗	✗
		3.4 Provision of consulting services to industry	✓	✓
5. Research and public debate about science and technology	The contribution to controversies or public debates about the role of science and scientific research	4.1 Participation in official national or international programmes or schemes	X	X
		5.1 Production of knowledge related to social impact of scientific work	✓	✓
		5.2 Production of policy arguments based on production of certified knowledge	✗	✗

✓ - Confirming Evidence
 ✗ - Partial Confirming Evidence
 X - Refuting Evidence
 ? - Unknown

TABLE 9.4: Research Practices of the Security Group at the University of Cambridge

Principle	Definition	Elements	Evidence of Elements	Evidence of Principle
1. The production of certified knowledge	The production of open scientific knowledge that is reviewed by colleagues and resistant to controversies	1.1 Production of academic articles 1.2 Indirect individual recognition (e.g. prizes, conference invitations, editorial board membership) 1.3 Indirect laboratory recognition (e.g. participation in national and international programmes) 1.4 National recognition (e.g. RAE)	X ✕ ✕ X	✕
2. Education, training activities, and embodied knowledge	The equipping of scientists with the tacit knowledge required to undertake laboratory work	2.1 Presence of postgraduate students	?	?
3. Public research and the innovation process	The creation of competitive products and advantages through innovation	2.2 Presence of individuals qualified to supervise postgraduate students 3.1 Receipt of funding from industry. 3.2 Patenting of work 3.3 Direct commercialization of work 3.4 Provision of consulting services to industry	? ? ✕ ? X	? ?
4. Participation in the production of collective or public goods	The production of standards and research that can be used by the public sector	4.1 Participation in official national or international programmes or schemes	✓	✓
5. Research and public debate about science and technology	The contribution to controversies or public debates about the role of science and scientific research	5.1 Production of knowledge related to social impact of scientific work 5.2 Production of policy arguments based on production of certified knowledge	? ?	? ?

✓ - Confirming Evidence
 ✕ - Partial Confirming Evidence
 X - Refuting Evidence
 ? - Unknown

TABLE 9.5: Research Practices of CESC at GCHQ

ubiquitous discrimination.

	Certified Knowledge	Education	Public research and innovation	Collective or public goods	Debates and science or science policy
Data Security Group	✖	✖	✓	✓	X
Information Security Group	✓	✓	✓	✓	X
Security Group	✓	✓	✖	X	✓
CESG	✖	?	?	✓	?

✓ - Confirming Evidence
 ✖ - Partial Confirming Evidence
 X - Refuting Evidence
 ? - Unknown

TABLE 9.6: Summary of Multiple Research Practices

9.3 Contributory Expertise During the Political Phase

In the previous section, it was pointed out that the third wave does not distinguish between different types of contributory expertise. However, an examination of the technical phase of the crypto wars, as described in chapters 4 to 7, showed that historical and institutional practices varied across research sites. A method of sociological discrimination that can be used to demonstrate this multiplicity was therefore described. The third wave also makes few rules about the ways in which contributory expertise ought be used during the political phase of a controversy. As was described in chapter 2, under elective modernism, the technical phase can have a constraining effect on the justifications used during the political if the Minimum Default Position is applied, and more generally, it is stipulated that the political phase should make no attempt to subvert the findings of the technical (Weinel 2010). But, in order to avoid technocracy, any kind of argument is permitted during the political phase, except for the advancing of quasi-religious/populist arguments by those with specialist expertise. In order to develop further guiding principles, elective modernism should incorporate an appreciation of how contributory expertise can be used during the political phase. This is because, it is conceivable that, if contributory expertise is misused during the political phase, or if it is absent, then the technical phase could have a minimal or negative influence on the political. This is clearly to be avoided, given that it would undermine the crucial relationship between the technical and political phases that underpins

elective modernism. With that in mind, I will now use the description provided in chapter 8 to characterize the ways in which the contributory expertise produced at each site was used during the political phase.

The most important thing to say about the expertise of the Data Security Group during the political phase of the crypto wars is that there is almost no evidence of its presence. Neither the Data Security Group nor NPL were involved as institutions, and individuals working within the group did not involve themselves as private citizens. None of the documents analysed that originated from NPL discussed the debates over TTPs or export controls. Similarly, none of the documents produced by the government or other actors during the political phase mentioned NPL or the Data Security Group at any point. This is despite the fact that, historically, the laboratory had built up a considerable body of contributory expertise related to cryptology. Furthermore, despite changes to the way NPL was funded, its stated role was still to support industrial and commercial activity through scientific research, and questions about the best way to aid electronic commerce were at the heart of the TTP debate. The only partial link that can be made between the Data Security Group and the crypto wars relates to early TTP consultation documents that asked whether it would be appropriate for the technologies used in TTP systems to receive ITSEC approval. However, due to the way in which the debate unfolded, the proposals did not progress far enough for this to be a pressing concern.²

Those working within the Data Security Group did not involve themselves in the political phase of the crypto wars as private citizens. This was particularly noticeable when I interviewed former members of the group. Most were only vaguely aware of the crypto wars, and of those that were, most demonstrated little interest in them:

Interviewer: Did you have a take on [the crypto wars]?

Respondent: No [laughs] . . . to put it bluntly.

Interviewer: Was it something that you followed at all?

²Even if the TTP proposals had progressed further, and ITSEC approval was eventually seen as a requirement, the Data Security Group would not have been the ones actually carrying out the testing. This would have been carried out by one of the CLEFs, who would have in turn been accredited by the Data Security Group. As it turned out, the CASM system that was proposed for use in securing emails between government departments, did receive ITSEC approval in 1998 (Communications-Electronics Security Group 1998c). However, this ceased to be of relevance to the crypto wars when the TTP proposals were dropped.

Respondent: No, I don't think I did, really.

Interviewer: Ok.

Respondent: To be honest, the reason I moved across to software was because I didn't really like security, actually. I quite liked the maths, all the algorithms and stuff, but I didn't really like the whole security area.

Others did not involve themselves directly in the crypto wars, but were interested in related issues, such as liability for banking fraud:

Interviewer: The impression that is given sometimes is that people were crusading for something. Is that something that was thought about?

Respondent: I think the only sense I got of crusading was against the banks - to make the banks responsible. There were people out there who really did want to make the banks responsible for their security - and rightly so.

This refers to the fact that, in the UK in the early 1990s, if a bank fraud case went to court, it was up to the bank's customer to demonstrate that they were not in some way responsible for the fraud having taken place, as opposed to the bank having to demonstrate that their system was as secure as they claimed. Although an important concern, questions about banking fraud and liability were only tangentially connected to the debates over TTPs and export controls.

On this basis, it can be concluded that the contributory expertise of the Data Security Group was absent from the crypto wars. This alone is significant, but it also begs the question of whether it is possible to further understand this absence. At present, a comprehensive sociology of the 'absent' does not exist. On the one hand, this is unsurprising, given that it is clearly difficult to place a consideration of the absent on a solid empirical footing. Furthermore, given that, as the well-known aphorism goes, "absence of evidence is not evidence of absence", those considering the absent can quickly find themselves on shaky logical ground. One way in which to think about absences has been provided by Boaventura de Sousa Santos. He has argued that:

The sociology of absences consists of an inquiry that aims to explain that what does not exist is, in fact, actively produced as non-existent, that is, as a non-credible alternative to what exists. The objective of the sociology of absences is to transform impossible into possible objects, absent into present objects. The logics and processes through which hegemonic criteria of rationality and efficiency produce non-existence are various. Nonexistence is produced whenever a certain entity is disqualified and rendered invisible, unintelligible, or irreversibly discardable. What unites the different logics of production of non-existence is that they are all manifestations of the same rational monoculture (Santos 2004).

On this understanding, an explanation of the absence of the contributory expertise of the Data Security Group during the crypto wars can be found rooted in the institutional and historical practices used to produce it. Given that, by the time the crypto wars started, the practices of the group were largely shaped by the deliberate imposition of the customer-contractor principle, the DTI (as the group's main customer) was able to specify a focus on accreditation and ITSEC. As the contributory expertise required to carry out this role had little bearing on the issues being discussed in both the debate over TTPs and the debate over export controls, there was no obvious route for the Data Security Group into the political phase. Furthermore, in order to carry out a program of research that would have allowed the group to enact a body of contributory expertise relevant to the political phase of crypto wars, they would have had to secure a particular type of external contract. On this basis, I argue that by thinking about the absence of the contributory expertise of the Data Security Group in this way, absences can be thought of as being enacted in the same way as the presences of contributory expertises from other sources.

During the political phase of the crypto wars, the contributory expertise of the Information Security Group at Royal Holloway was present in the proposed use of a protocol that could be used to underpin the TTP system. As was described in chapter 5, the protocol was produced as part of a UK DTI/EPSRC research project entitled 'Security Studies for Third Generation Telecommunications Systems' (GR/J17173/01). The £160,000 project was carried out in conjunction with Vodafone (which had recently emerged from the Racal group) and GPT. The Royal Holloway protocol matched the government's stated requirements for the

TTP network, and was linked to proposed systems to secure government emails and NHS data. In addition to the provision of technical solutions, members of the group were also commissioned to provide an independent assessment of the CESG-developed Red Pike algorithm. On this basis, it can be argued that the expertise of the group was also used to underpin technological solutions and to approve the solutions developed by others.

Despite using their expertise to produce the technologies at the heart of the TTP proposals, the group do not appear to have been involved in any other respect. For example, there exists no evidence of them being involved in the public debates, either informally at the level of ukcrypto mailing list exchanges, or formally during public consultations or Select Committee meetings. When interviewed, the members of the group appeared comfortable in making a distinction between their technical work and its wider policy implications. Some, when asked, did not appear to connect the two:

Interviewer: What was your position during the 90s encryption policy debates, export regulations et cetera, if indeed you took one?

Respondent: My involvement in cryptology has been concerned with technical matters, and I have not been involved in such policy issues.

Interviewer: Did you make any contribution towards policy or legislation?

Respondent: No.

Interviewer: Did you contribute in other ways, such as sitting on government committees?

Respondent: No.

Others recognised that their technical work and policy views might conflict, but were still able to separate them. Here, the key escrow aspect of the Royal Holloway protocol is discussed:

Respondent: Key escrow was a proposal based on, crudely, and I mean very crudely, on the assumption we, we being governments, were happy for our companies to be secure against anyone but us, but we need to control encryption somehow so that's where encryption came in. We did not, as an emotional concept, support key escrow, but we

did design a key escrow system that could work. We saw no contradiction in terms in that. Basically, we designed a key escrow management system. The trust model was clear, and people didn't like the trust model, which is fine. I'm not sure we liked the trust model either, but that wasn't the point. The key escrow model could work. Under key escrow you were secure against your enemies, provided your enemies didn't infiltrate the key escrow agent. But, technically it was quite an interesting problem.

Despite not supporting the politics of the key escrow concept, the group were able to design a system that they believed could work. Furthermore, by using their expertise to provide the technology that was to be used in the controversial proposals, there is a sense in which this precluded the group from using that same expertise to express any dissatisfaction with the principles upon which the policy was based. This may explain why those from the group did not involve themselves in the crypto wars in other ways, and why they were absent from formal and informal debates. At the very least, it demonstrates that contributory expertise enacted during the technical phase may not be present during the political phase in the form of symbolic arguments - it may also be found embedded in technologies.

Again, there appears to be a link between the practices used to enact expertise within the Information Security Group, and what this expertise was used to accomplish during the crypto wars. The group enacted their expertise through collaborating with government and industry to create technologies, and they were able to use this expertise to produce a technology to be used in the proposed TTP system. That a research group within a university should design their research practices in this way is no longer unusual or surprising. A number of theoretical frameworks have been developed to describe such trends in scientific research practices, including: 'Mode 2' (Gibbons et al. 1994); 'Triple Helix' (Etzkowitz & Leydesdorff 2000); and 'Academic Capitalism' (Slaughter & Leslie 1997). Of these, the theory of academic capitalism appears to be the most relevant to this particular case. The theory of academic capitalism:

... sees groups of actors - faculty, students, administrators, and academic professionals - as using a variety of state resources to create new circuits of knowledge that link higher education institutions to the

new economy. These actors also use state resources to enable interstitial organizations to emerge that bring the corporate sector inside the university, to develop new networks that intermediate between private and public sector, and to expand managerial capacity to supervise new flows of external resources, investment in research infrastructure for the new economy, and investment in infrastructure to market institutions, products, and services to students. Expanded managerial capacity is also directed toward restructuring faculty work to lower instructional costs (although not costs generally) (Slaughter & Rhoades 2004, p.1).

The broad consequence of this trend, assuming it is accurate, is a shift from universities that enact expertise exclusively for the ‘public good’, and towards universities that enact expertise that can also be used as a commercial resource (or a ‘private good’). This, of course, does not imply that there is no overlap between the two, and that they cannot in some sense be co-produced. However, studies of academic capitalism have also highlighted the failure of some academic institutions to meet the expectations of society in terms of economic growth, employment, and student training. Simon Marginson (2011), drawing on Jürgen Habermas’ ideas about the development of the ‘public sphere’, has argued that universities, as recognized producers of public goods, play an important role in criticizing policy, and that status competition and marketization can undermine this enterprise. But, as the Royal Holloway example shows, expertise can be embedded in technologies that support public goods, as well as in advice offered to policy makers in the form of arguments (Calhoun 2006). Therefore, the consequences that a trend towards academic capitalism might have in terms of how the expertise it produces effects controversies have not been fully developed. However, if research practices increasingly result in the enacting of expertise that produces technologies in collaboration with commercial or governmental partners, then role of the contributory expertise that emerges from certain university laboratories needs to be reconceptualized accordingly.

During the political phase of the crypto wars, the contributory expertise enacted by the Security Group at the University of Cambridge was absent from the initial export control and TTP proposals. This resulted in the construction of an opposition to them. The form that this opposition took changed as the crypto wars progressed. At the beginning of the debate over TTPs, Ross Anderson and Michael Roe used their contributory expertise to level explicit technical criticisms at the Royal Holloway protocol. These criticisms demonstrated that the TTP

proposals were potentially insecure. Anderson also used his expertise to develop criticisms of how the protocol was to be put to use by government departments and the NHS. These criticisms were expressed at meetings to discuss the proposals, and informally on the ukcrypto mailing list - with the thrust of the arguments being largely accepted by those using the latter. When the TTP consultation process began, the opportunity to register a formal opposition arose, and was taken by members of the Cambridge group, and others persuaded by their arguments. Anderson in particular contributed during every stage of the consultation process, and gave evidence in person at Select Committee meetings. The Cambridge group also played a key role in creating the FIPR think-tank, which would mount a more coordinated opposition later on in the debates. This led to those opposed to the government's proposals becoming more involved in the political phase of the crypto wars. In addition to rising to the ever more formal requirements that effective opposition demanded, Anderson and other FIPR contributors were able to use their cryptology expertise to contribute to the wider export control debate, during which cryptology was only one of many fields that might have been affected. This culminated when they successfully used their cryptology expertise to push for the amending of the Export Control Bill to exclude the export of academic ideas, in an arena that was partly made up of groups pushing for tighter regulation. By the end of the crypto wars, it was clear that the Cambridge group and others had transferred their expertise from the technical phase to the political phase.

It has been acknowledged within STS for some time that scientific expertise is often enlisted by others in order to add weight to political claims during controversies (see Nelkin 1995). However, the use of scientific expertise by scientists to underpin their own participation during a political phase appears to be a different exercise, and one that is less commonly referred to. An extreme form of this activity has been referred to as 'scientist activism'. Scott Frickel (2004) has described how the field of genetic toxicology emerged from the political activism of scientists concerned about the dangers of chemical mutagens. Although this differs slightly from how the expertise of the Cambridge group was used to underpin their own participation in the political phase, it is perhaps something that elective modernism should account for, given it prescribes that society should "always aspire to keep the technical and the political phase separate even where they are combined in institutions or individuals" (Collins et al. 2010). However, interventions such as this, though arguably a positive development during crypto wars, may be undesirable in other cases if a clear separation between phases is to be maintained.

Again, there appears to be a link between the research practices the Cambridge group used to enact their expertise, and what this expertise was used to accomplish during the crypto wars. The group designed research practices to understand how security systems behaved in the real world. To do this they used interdisciplinary methods then unorthodox in security and cryptology research. The conclusions from this research - based as they were on appraisals of threat models used by systems designers - enacted a body of expertise that could be translated into policy arguments. Furthermore, those that possessed the expertise were also sensitive to some of the issues associated with secure system design that they felt others may not have been aware of. When the government proposed their TTP system, the arguments that the group had developed could be used to critique the proposals, and formed the basis for an opposition. By contributing to the formal consultation processes during the TTP debates, it could be argued that Anderson and other members of FIPR also enacted a body expertise related to contributing towards scientific controversies. This expertise was put to use during the export control debates, where cryptography was a relatively minor issue.

During the political phase of the crypto wars, the contributory expertise enacted by CESG was partially evident in their proposals for government security schemes. However, the secret nature of CESG's involvement in the crypto wars was one of the most important aspects of their contribution. To understand the impact of their secret expertise on the political phase of the crypto wars, the first step is to consider their historical association with cryptography. CESG (and GCHQ) have long been associated with cryptology expertise. Up until the 1970s, cryptology expertise had been seen as something that was the preserve of militaries and intelligence organizations. This is exemplified by the fact that Alan Turing's now famous code breaking efforts at Bletchley Park during the Second World War were undertaken whilst working for the GC&CS - which morphed into GCHQ after the end of the war. Therefore, it became natural to assume that they would continue to have expertise in this area. For example, when Donald Davies first started serious cryptology research at NPL in the 1970s, he thought that this might conflict with work being done at GCHQ. Davies recalled that he "had been in touch with GCHQ about data security and realized that it would be very difficult for NPL to get involved in this area, because of the way they regarded any government work in this field to be very much their province" (Campbell-Kelly 1986). This was further compounded by the fact that, by the time the crypto wars in the UK were

underway in the mid-1990s, the involvement of the NSA (usually thought of as GCHQ's American counterpart) in the crypto wars in the US was widely known.

Despite this reputation, by the time the crypto wars in the UK were underway, the secrecy that surrounded CESG meant that very little was known about their expertise in the field. This made an accurate assessment of their expertise problematic. The combination of CESG's historical reputation, and the lack of information about their expertise, appears to have resulted in a tension. This tension primarily manifested itself as anxieties about CESG's lack of expertise and the motivations behind their involvement. In terms of their motivations, over time, the view that CESG and GCHQ were 'pulling the strings' for the TTP and export control proposals gained wider currency. This view was initially made evident in discussions on the ukcrypto mailing list. When the documents relating to the NHS network and the Cloud Cover system were discussed, contributors speculated on behind-the-scenes involvement and motivations of CESG and GCHQ. There was a repeatedly expressed belief that the NHS were being persuaded to adopt the CESG-developed Red Pike algorithm to improve the chances of the wider adoption of a key escrow-based system over which CESG and GCHQ had control, thus allowing them to continue to use telecommunications networks as a source of intelligence.

There also emerged a view that, despite being at the forefront historically, CESG had fallen behind with developments in cryptology. Again, this view appears to have emerged from discussions on the ukcrypto mailing list. One contributor based at a UK university stated that:

I have spoken with people from CESG (and GCHQ) on various occasions here at the University and elsewhere concerning research programs etc. My impression is that they had a "golden age" which started with the stunning successes at Bletchley and lasted up to about 10-15 years ago but that they have now really "lost it" in all respects. They were often unaware of key papers in many areas and were very grateful for a photocopy to take away with them!! Far and away the best cryptographers are in the academic world and CESG is becoming more and more sidelined in the main stream of events. This was very evident at the CompuSec exhibition last year from the way they presented themselves on their stand. If they want to be taken seriously as

a consultancy to whom the general populace will come for IT security advice then they really need to get their act together in a VERY big way first.

Although initially expressed informally, views such as these came to be used formally during the political phase, as the debates heightened. This is best demonstrated by quoting an oral response given by Ross Anderson to a question from Helen Southworth MP during a Select Committee meeting held on March 2nd, 1999:

Helen Southworth MP: You are suggesting that the intelligence agencies are pushing key escrow and that they are being left behind with developments in cryptography. Can you justify those two points of view?

Ross Anderson: That the intelligence agencies are pushing the escrow agenda is I think evident from the [Zergo NHS report] written by Henry Beker and Chris Amery, who have been long-term suppliers to GCHQ, have the clearances and so on. The kinds of problems that one is having with GCHQ's falling behind the curve with commercial cryptography can be seen, for example, in Cloud Cover which is a key management system that has been promoted within the Civil Service and which CESG has tried to get the NHS to adopt. I have ended up on the opposite side to that because I advised the BMA on safety and privacy of clinical systems and we found quite a number of things wrong with Cloud Cover (Select Committee 1999).

It could be argued that these are examples of those with contributory expertise carrying out domain-specific discrimination, given that they were forming judgements about others who have contributory expertise on the surface, but can be distinguished from 'true' experts by virtue of internal technical and social criteria (Weinel 2010, Collins et al. 2010). It may be that their assessments were correct. However, given the secrecy surrounding CESG, it is clear that these assessments are unlike other attempts at domain-specific discrimination. For instance, they are based on incomplete information about the expertise being judged, and, as was acknowledged by those making the judgements, they are attempting to judge

individuals and institutions whose motivations they are unsure of. More specifically, given that CESC rarely published details of their work, it would have been difficult to assess - and even harder to prove - how favourably their work compared to that of other contributory experts in terms of criteria like the journals they published in, and their overall readership. Indeed, it could be argued that the lack of information relating to CESC's cryptology expertise, in part, allowed such arguments to be made and gave them greater currency, given that, in order to maintain secrecy, they could not be readily falsified or refuted.

Uncertainty about CESC's expertise and motivations also allowed for arguments to be made about the broader relationship between surveillance and the state during the political phase. Caspar Bowden - at the time the director of FIPR - and Yaman Akdeniz - founder and director of Cyber Rights & Cyber Liberties - discussed this in a paper on 'Cryptography and Democracy'. Here, it was argued that the implementation of the government's proposals would result in a "slippery slope" towards a surveillance state, reminiscent of George Orwell's *Nineteen Eighty-Four* (Bowden & Akdeniz 1999). At one point, it is claimed that representatives from GCHQ attempted to persuade the OECD to adopt the CASM system as an international standard. However, the source for this claim was "private information from those present at OECD meetings" (Bowden & Akdeniz 1999, p.110). Whilst there is nothing to suggest that this claim is untrue, it is typical of the way in which claims about GCHQ and CESC were often made in the crypto wars, and indicates how secrecy can both distort attempts at assessing expertise, and distort attempts to publicly demonstrate claims.

On this basis, I argue that the secret nature of the expertise of CESC produced uncertainty during political phase of the crypto wars. As such, it can be linked with an emerging STS literature on secret science. As was described in chapter 2, although little has been written about secret science from an STS angle, an emerging view urges that secret science be thought of as something more than just open science done behind closed doors. Much scientific work involves an element of secrecy. It is not uncommon, for example, for scientific work to be kept secret in order to properly establish priority claims or to gain a commercial advantage. However, the level of secrecy employed by institutions like GCHQ clearly exceeds this (counter-)norm, given that the upholding and maintaining of secrecy dictates most of their activity. Balmer has argued that secret science can have consequences for those whom information is kept from. For example,

secrecy can actively construct uncertainty, gossip and rumour that can be used both positively and negatively by actors on all sides of the secrecy divide (Balmer 2012, pp.76-77). On this understanding, the effects of secrecy are linked to the practices used to enact secret expertise.

This is important to bear in mind when considering what was enacted by CESG's expertise during the crypto wars. Whilst CESG's expertise undoubtedly produced technological artefacts in the form of Red Pike, technical assistance and advice in the form of the Cloud Cover system, and 'collective goods' in the form of their management of ITSEC, the secret nature of their research practices produced uncertainty during the crypto wars that effected the nature of many of the other contributions. This uncertainty made assessments of expertise problematic - even during attempts at domain-specific discrimination - and encouraged speculation about underlying motives. This was perhaps expressed best by one interviewee who, when asked about CESG's role during the crypto wars, simply said that "they were either everywhere or nowhere".

To sum up, the descriptions provided in this section indicate how varied the uses of contributory expertise during a political phase can be. It also shows that the links between the practices used to enact the contributory expertise during the technical phase, and its use during the political phase, are intelligible.³ This can be used to suggest that some of the assumptions built into elective modernism may need refining. For example, elective modernism assumes that "political decisions should not be made without considering as much as possible of the technical knowledge which bears upon the decision" (Collins et al. 2010). Yet, the use of the contributory expertise enacted by CESG during the political phase appeared to have had a negative influence on the debate, given that it introduced expertise that was problematic for other experts to accurately assess. Although the operationalization of domain-specific discrimination has not been precisely laid out, it is clear that no matter how it is done, if it is based on incomplete information then it becomes difficult (or even impossible) to know whether the conclusions are flawed. It is also likely that this lack of information on which to base domain-specific discrimination went some way towards licensing speculation about the motives of CESG and GCHQ, and had the potential to legitimize the use of quasi-religious/populist arguments by those with contributory expertise. Although private citizens are

³It should be noted that these links were made intelligible with hindsight. It is not yet clear whether it is possible to use sociological discrimination to predict exactly how a particular type of contributory expertise will be used, or the effect it might have, in real-time.

permitted to make these arguments during the political phase, those with specialist expertise are not. However, they may be the best placed to make them, given that it would be difficult to argue that their attempts at domain-specific discrimination - though compromised - were inferior to those based on ubiquitous or local discrimination. During the crypto wars, this served to undermine the separation between the formative intentions of the technical phase and the political phase, because technical questions became entangled with political questions. This is further evident in the fact that, over the course of the political phase, the Cambridge group's involvement in the political phase blurred with their technical work. This can be seen in their particular form of 'scientist activism', such as during the formation of FIPR, and some of their more overtly political publications, such as the Global Trust Register. This isn't a criticism of how the Cambridge group, and other activists, involved themselves in the crypto wars. However, it made the aspiration "to keep the technical and the political phase separate even where they are combined in institutions or individuals" (Collins et al. 2010) unrealistic and potentially undesirable. These contradictions suggest that additional elective modernist principles may be required in order to chart a path through controversies of this nature.

9.4 The Transfer of Contributory Expertise

Before moving on to making some suggestions for what those principles might be, it is worth considering another dimension of the relationship between the technical and political phases, and multiple contributory expertises - namely, how contributory expertise is 'transferred' from the technical phase to the political phase. As was described in the literature review, Evans and Plows (2007) have argued that the relationship between the phases is circular, with the outputs from one informing the other. Weinel (2010) has argued that the conclusions of the technical phase should have a constraining effect on the justifications used during the political. However, these arguments - though useful - are not based on a prior distinction between different types of contributory expertise. Making this distinction may be helpful in understanding the transfer of contributory expertise between phases, and its impact upon the decisions made during the political phase.

Evans and Plows (2007) described the relationship between the technical and political phases as a whole. I will focus on one particular 'transfer' within that broad

relationship - namely, how the political phase comes to know about the output of the technical phase. Given that the relationship between the phases is reciprocal, this is only one dimension of the transfer process. However, it is an important one for elective modernism to address, given that the quality of the debate during the political phase is dependent on it.

Before proceeding, it is necessary to devote some time to discussing what ‘transfer’ means in this context, and following on from this, how a transfer might be identified and characterized. In simple terms, a transfer can be said to have occurred when the expertise enacted during the technical phase is used during the political phase. However, expertise ‘use’ can potentially mean a number of different things. In order to flesh out this idea, it is useful to draw on insights from fields such as ‘Knowledge Transfer’, ‘Research Utilization’, and ‘Knowledge Mobilization’, that have focused on how research is used to inform policy.⁴ In a survey of research use, Julius Court and John Young (2003) found that it could be used for a wide variety of ends, including pushing issues onto the policy agenda, influencing key policy decisions, enhancing knowledge, supporting the development of networks, and changing ways of working. It is therefore important to acknowledge that research use during a political phase can be either ‘instrumental’ - when it has a direct impact upon policy decisions, or ‘conceptual’ - when it has an indirect impact upon knowledge, attitudes, and practices (Caplan 1979, Webber 1986). On this understanding, when aiming to identify a transfer of expertise, it is important to remain aware that its use may take many different forms.

Whilst it is important to acknowledge that expertise can be used during the political phase in a number of different ways, some have nonetheless attempted to identify stages of research use. Jack Knott and Aaron Wildavsky (1980) conceived of research use as occurring in seven linear stages:⁵

1. Reception: The research is received (the research is considered ‘used’ even if it is never actually read);
2. Cognition: The research is read and understood;

⁴Of course, expertise, knowledge, and research are not precisely the same thing. However, the study of the relationship between knowledge, research, and policymakers offers a useful starting point for considering the transfer of expertise, given that the terms undoubtedly overlap.

⁵A number of models also aim to describe the process of research use in this way. In some cases, the first stage is similar to Knott and Wildavsky’s ‘Reception’, in that they describe the ‘Transmission’ of research findings, or practitioners becoming ‘Aware’ of them (Landry et al. 2001, Glasziou & Haines 2005).

3. Reference: The research changes ways of thinking;
4. Effort: The research has shaped actions;
5. Adoption: The research has a direct influence on policy;
6. Implementation: The policy containing the research has been implemented;
7. Impact: The implemented policy is successful.

Though Knott and Wildavsky's model was designed to understand how research is used, it can also be used to underpin ideas about the transfer of expertise. It enables us to say that expertise has been transferred from the technical phase to the political phase if there exists evidence of any of the above stages having occurred. For example, 'Reception' can be said to have occurred if a document produced during the course of the enacting of expertise during the technical phase - such as a research paper - appears on the list of documents submitted as evidence to a Select Committee, and 'Adoption' can be said to have occurred if that document is cited by a government policy report. In both cases, it is clear that the expertise has been transferred from the technical phase to the political phase.

Once a transfer of expertise has been identified in this way, it becomes possible to characterize the nature of that transfer. Once again, ideas about research use can be used to illuminate the nature of the expertise transfer process. As Sandra M. Nutley, Isabel Walter, and Huw T. O. Davies (2007) pointed out, policymakers can encounter research in a number of different ways, and are able to access it from a variety of different sources. It may be encountered through direct communication with those that produced it, or, the transfer process may be mediated through knowledge brokers such as research centres and government organizations. Furthermore, as Åke Bergmark and Tommy Lundström (2002) highlighted, it is possible to distinguish between research that is encountered actively - if the policymaker seeks out research to help support their work, and research that is encountered passively - if the policymaker is presented with research. Finally, it is also possible to consider the form in which research is transferred. Research findings may be embodied in a number of different ways. For example, they may be embodied in technologies, objects, actions, practices, and symbolic arguments.

With this in mind, I will now return to the descriptions of the technical phase of the crypto wars from chapters 4 through 7, as well as the description of the

political phase in chapter 8, to identify and characterize the transfer processes used. The contributory expertise enacted during the technical phase in the Data Security Group at the NPL was absent from the political phase. Therefore, we can say that no transfer of expertise occurred. There may be interesting and important reasons for why this was the case, but for the time being, it is sufficient to note that no transfer took place. Given this, there is nothing that can be said about how the political phase came to know about the conclusions of the technical, or about the embodiment of the expertise.

In contrast, the contributory expertise enacted during the technical phase by the Information Security Group at Royal Holloway was transferred to the political. This expertise was embodied in a technology that underpinned the controversial TTP proposals. Therefore, we can say that the expertise was transferred in the form of a technology, rather than through the use of symbolic arguments. The ‘Royal Holloway’ protocol was produced in conjunction with the DTI as part of a SERC-funded project. Therefore, the expertise of the Information Security Group was actively ‘commissioned’ from those working in the technical phase.

The contributory expertise enacted during the technical phase in the Security Group at the University of Cambridge was initially ignored, but was then used to construct an opposition to the controversial proposals. Therefore, we can say that the expertise was embodied in technical and political symbolic arguments. It was transferred from the technical phase to the political phase through active engagement with the consultation procedures, engagement with the press, and other political activity. Therefore, in contrast to the contributory expertise of the Information Security Group, from the point of view of the policymakers in the political phase, the expertise of the Security Group was ‘delivered’ by those working in the technical phase.⁶

The contributory expertise enacted during the technical phase within CESC at GCHQ was embodied in technologies that underpinned the controversial proposals. Furthermore, it was believed that CESC’s expertise was used to provide symbolic arguments that underpinned how the systems should function, and what

⁶It is important to remember that, whilst the expertise of the Cambridge group was delivered by the technical phase to the political, this was facilitated through formal consultation procedures, one of the purposes of which is to provide an opportunity for experts to make their views known. However, participation in the consultation procedures is optional. Therefore, although the expertise of the Cambridge group was in a sense ‘solicited’ by the political phase, it was still up to the Cambridge group to actually take steps to ‘deliver’ it.

their relationship with other policies should be. However, the conditions under which this expertise was enacted were secret. Furthermore, and as a partial consequence, the way in which the expertise was transferred from the technical phase to the political phase was not made clear. Unlike in the case of the expertise enacted by the Information Security Group and the Security Group, it was unclear whether the expertise was commissioned by, or delivered to, the political phase. Therefore, the transfer of expertise can be described as ‘invisible’.

Research Site	Transferred?	Embodiment	Direction	Transfer Type
Data Security Group	No	N/A	N/A	N/A
Information Security Group	Yes	Technology	Technical <- Political	Commissioned
Security Group	Yes	Symbolic argument	Technical ->Political	Delivered
CESG	Yes	Technology; Symbolic argument	Technical ?? Political	Invisible

TABLE 9.7: The Transfer of Contributory Expertise

As with the other descriptions in this chapter, these descriptions of contributory expertise transfer can be summarized in a table (see Table 9.7). This information shows that we can use sociological discrimination to distinguish between the embodiment of the expertise, and the direction of the transfer from the technical to the political phase. The consequences that making these distinctions might have for elective modernism will be suggested in the next section.

9.5 Consequences for Elective Modernism

In this short final section, I will use the information from the previous three to briefly consider some of the consequences that acknowledging a multiplicity of contributory expertise might have for how elective modernism should function during controversies over technological decision-making. I will outline a number of problems that emerge from the descriptions of what happened during the crypto wars. I will then suggest what I have called the ‘Minimum Transfer Requirement’, and identify what I have called the ‘problem of expert discrimination’. These are by no means the only insights that can stem from an acknowledgement of multiple contributory expertises, and nor are they fully-developed enough to be considered final or comprehensive. However, they will provide an indication of both why it is important to take multiplicity seriously, and how elective modernism might be refined in order to confront the issues that it raises.

9.5.1 Minimum Transfer Requirement

In the first section, it was shown that it is possible to characterize the different types of contributory expertise that can be enacted within a discipline during the technical phase of a controversy. In the second section, it was shown that the uses of contributory expertise during the political phase are also multiple. Furthermore, plausible links can be established between the practices used to enact the expertise during the technical phase, and how it is used during the political. Elective modernism requires that “political decisions should not be made without considering as much as possible of the technical knowledge which bears upon the decision” (Collins et al. 2010). However, this raises the question of whether it is possible to assess whether or not this has actually taken place. A way of beginning to address this problem is by understanding more about what I have called the ‘transfer’ process. I use this term to refer exclusively to how the political phase comes to know about the conclusions from the technical. In the third section of this chapter, I have shown that it is possible to arrive at a distinction between expertise that is: not transferred; commissioned; delivered; or transferred invisibly. I suggest that these categories can be used to define a minimum requirement that those in the political phase must fulfil in order to be able to claim that they have attempted to draw upon an adequate level of expertise from the technical phase. I call this the Minimum Transfer Requirement.

The Minimum Transfer Requirement states that decisions made in the political phase cannot be made solely upon consideration of expertise that has been commissioned, or expertise that has been transferred invisibly. Decisions should not be made based solely upon consideration of expertise that is commissioned because this would allow decision makers to draw on expertise that they either know to be already aligned with preconceived policy objectives, or expertise that they had an active role in enacting. Decisions should not be made based solely on expertise that has been transferred invisibly as it would, in extreme cases, be impossible for those observing the political phase (including those working in the technical phase) to use sociological discrimination (or other forms of meta-expertise) to determine whether it had been commissioned or delivered. It therefore also follows that decisions made in the political phase cannot be made after considering only a combination of invisible and commissioned expertise. Of course, as elective modernism stipulates, there is no requirement for the political phase to act upon

expertise that has been transferred in a particular way. It may be that the political phase chooses to act upon expertise that has been commissioned, rather than expertise that has been delivered, assuming they are in contradiction. However, in order for the political phase to be able to claim that it has considered as much as possible of the technical knowledge which bears upon the decision, at the very least it must fulfil the Minimum Transfer Requirement.

9.5.2 Problem of Expert Discrimination

In the first section, although it was shown that it is possible to use sociological discrimination to characterize the different types of contributory expertise that can be produced within a discipline during the technical phase of a controversy, in some cases, the adoption of practices designed to uphold secrecy can actively prevent the acquisition of the information required to make such a characterization. In this case, the expertise itself can be characterized as secret. In the second section, it was shown that the uses of contributory expertise during the political phase are also multiple. Furthermore, plausible links can be established between the practices used to create the expertise during the technical phase and how it is used during the political. In the case of the secret expertise, this created uncertainty during the political phase because the secrecy surrounding it made an assessment of its quality problematic. Even domain-specific discrimination could not be carried out in a satisfactory way because it was based on incomplete information, and the construction of secrecy around it may have meant it was also based on deliberately misleading information. What would have been domain-specific discrimination in other circumstances, was potentially reduced to a form of transmuted expertise. The lack of information upon which to base an assessment of this expertise was itself used by experts to criticise it during the political phase, and as a platform for broader arguments about the circumstances which were used to produce it. This ultimately had the potential to undermine the crucial separation between the technical phase and the political phase.

The suggests that, in some cases, there may be a ‘problem of expert discrimination’. The tension at the heart of this problem is that elective modernism states that political decisions should not be made without considering as much as possible of the technical knowledge which bears upon the decision. However, the use of some contributory expertise to inform decisions can be controversial because

it cannot be adequately assessed by others - even other experts. The problem of expert discrimination carries consequences. If this expertise is not drawn upon, then political decisions may not be taken on the basis of the best technical expertise. If it is used, it may not be trusted by either experts or the public. It may undermine the separation between the technical and political phases, as it blurs the concerns of both. Furthermore, it may lead to decisions that are viewed as ‘undemocratic’. Given that the separation between the technical and political phases lies at the heart of elective modernism, it follows that principles should be designed to uphold it. There may be numerous ways of addressing this problem. For example, answers to propositional questions that cannot be publicly assessed using domain-specific discrimination, and by implication other forms of non-transmuted expertise, should not be permitted to transfer from the technical phase to the political. This would strike a balance between aiming to draw upon as much expert advice as possible, whilst also preserving the crucial distinction between the technical and political phases. An alternative way of addressing the problem of expert discrimination may be to conclude that contributory expertise that cannot in principle be assessed using domain-specific discrimination should not be considered contributory expertise at all. If this view is taken, then not permitting this ‘expertise’ to transfer from the technical phase to the political phase could also be justified. However, taking this view may require an additional reconceptualization of the technical phase, given that it may still be possible to enact expertise that cannot be assessed using domain-specific discrimination using practices that are aligned with the formative intentions of science.

9.6 Conclusion

In this chapter, I have shown how historical and institutional research practices enacted a multiplicity of contributory expertises during the technical phase of the crypto wars. To demonstrate this, I have used Larédo and Mustar’s criteria for assembling laboratory activity profiles to underpin a process of sociological discrimination. This has been used to show how, to varying degrees, the Data Security Group at NPL, the Information Security Group at Royal Holloway, the Security Group at the University of Cambridge, and CESG at GCHQ, engaged in the production of certified knowledge, education and training, public research and innovation, the production of collective goods, and debates over policy. I have

further shown that these multiple enactments of expertise were used for specific purposes during the political phase of the crypto wars. On this basis, I argue that it is possible to differentiate between sources of contributory expertise based on the specific characteristics of practices, rather than at a disciplinary level.

I have also made a case for why it is important to take the multiplicity of contributory expertise seriously. In a broad sense, I have described a method that could be used to inform the processes used during the political phase in order to make the best use of the expertise available. Given that a controversy over technological decision-making can hinge on a disparate range of propositional questions, an appreciation of the multiplicity of contributory expertises that can exist within a discipline can offer guidance to decision-makers about how to go about sourcing informed answers. Acknowledging this multiplicity also highlights the problems and contradictions associated with some of the existing principles of elective modernism. During the process of characterizing the various enactments of expertise, it became clear that the secrecy surrounding some of the practices used at CESG meant that there was neither the presence of confirming or refuting evidence for particular indicators of research activity. The contributory expertise produced by CESG during the technical phase was therefore characterized as secret. When this expertise came to be used during the political phase, it introduced further uncertainty. Other experts came to question the motivations behind the contributions made by CESG, and speculated about the quality of their expertise. This arose out of the fact that the secrecy surrounding CESG's research during the technical phase could not be properly assessed using domain-specific discrimination. The assessments of other experts were necessarily based on incomplete information, and it was acknowledged that CESG might be engaged in practices that were deliberately designed to deceive. This, in part, prompted some experts to become active during the political phase, and to act in such a way that the boundaries between the technical and political phases became blurred. Given that political decisions should not be made without considering as much as possible of the technical knowledge which bears upon the decision, this raises questions about whether it is desirable for the political phase to draw upon types of contributory expertise that have the potential to undermine the other principles of elective modernism.

In order to develop answers to these questions, I characterized elements of the relationship between the technical and political phases during the crypto wars. I focussed on one particular process in this relationship - namely, how the political

phase came to know the conclusions from the technical phase. This allowed for a simple distinction to be made between expertise that was not transferred, expertise that was commissioned, expertise that was delivered, and expertise that was transferred in a way that was not visible to those not directly involved in the transfer process. This information can be used to think more carefully about how much information from the technical phase should be available to the political phase. Given that elective modernism states that as much from the technical phase as possible should be considered by the political phase, it appears useful to think more carefully about what can be classed as an appropriate basis for decision-making. Therefore, I suggest the use of the ‘Minimum Transfer Requirement’ principle. The principle states that, for it to be claimed that the political phase has made its decisions in light of the conclusions from the technical, it must at least be able to demonstrate that it has not made decisions solely on the basis of expertise that was commissioned or transferred invisibly. Secondly, the problem of expert discrimination recognises that, although decisions in the political phase should be made in light of as much expert advice as possible, this is complicated by the fact that it is not always possible for experts to adequately assess the expertise of other experts, thus throwing doubt on the expert status of some contributions.

Chapter 10

Conclusion

10.1 Introduction

In this concluding chapter, I will re-state the research questions that formed the backbone of this thesis, followed by my answers to them. I will then reflect on the limitations of the arguments made in this thesis and suggest some ideas for future work. What emerges from these reflections is a need for more empirical case studies viewed through the lens of the third wave.

10.2 Research Questions and Answers

In chapters 1 and 2, I used the current literature on the crypto wars, controversies over technological decision-making, the third wave, and the ontological framework, to pose three research questions:

1. How was contributory cryptology expertise produced during the technical phase of the crypto wars in the United Kingdom?
2. Can the ontological framework and the third wave be used in conjunction to develop a reconceptualization of the production of this contributory expertise as ‘multiple’?
3. What were the consequences of this multiplicity of contributory expertise during the political phase of the crypto wars?

My answer to the first research question was provided in chapters 4 through 7, and was summed up in chapter 9. My answer was given in the form of four individual descriptions of the practices used to enact contributory cryptology expertise at four different research sites: the Data Security Group at the National Physical Laboratory; the Information Security Group at Royal Holloway, University of London; the Security Group at the University of Cambridge; and CESG at the Government Communications Headquarters.

In the case of the Data Security Group at NPL, practices were designed in accordance with New Public Management priorities. Although the Computing Division had been associated with pioneering computing research following the Second World War, by the time the Data Security Group was founded in the late 1970s, the move towards New Public Management had increasingly commercialized the work of the division. In response to the requirements of commercialization, the Data Security Group initially enacted expertise in cryptology standards, but later, their focus turned to enacting expertise on cryptology accreditation and testing.

In the case of Royal Holloway, following a series of reforms to the University of London that culminated in the 1980s, practices across science departments were altered to foster industrial collaboration. The Information Security Group was founded because cryptology was seen as a useful niche for Royal Holloway to occupy. Given that the key founder members of group had mathematical backgrounds, the Information Security Group was able to provide the mathematics required to underpin commercial products and government technologies. As a result of their strong industrial partnerships with companies like Racal-Comsec, those working within the Information Security Group enacted expertise that met with the requirements of external partners.

In the case of the Security Group at the University of Cambridge, following an earlier emphasis on providing a usable computing service to the university, and several large projects to build early computers, the computing laboratory diversified its research programme during the 1980s - in part thanks to a period of generous funding. Computer security was one of the new areas of enquiry that emerged from a traditional focus on computer systems. Although the group carried out work on the development of technologies, other work attempted to understand cryptology as a component in larger security systems. This strand of research was characterized by the use of interdisciplinary methods, and engagement with fields

such as psychology and economics. The group therefore used practices to enact expertise that pertained to the real-world behaviour of cryptology systems.

Finally, in the case of CESG at GCHQ, research into cryptology had a much longer history. From its founding in 1919, cryptology research practices were designed in accordance with intelligence priorities. The nature of the research carried out at GC&CS in this early period is exemplified by the work of Alan Turing during the Second World War. As a result of intelligence priorities, the nature of the cryptology research carried out by CESG has typically been a closely guarded secret. With the exception of the research on non-secret encryption - carried out in the 1970s and revealed in the late 1990s - little of their modern cryptology research has been made public. Although there have been attempts to reform some aspects of the UK's intelligence organizations since the late 1980s, and although CESG has a more public role than most bodies within GCHQ, practices designed to uphold the Official Secrets Act continued to dominate. Therefore, during the technical phase of the crypto wars, CESG enacted a secret cryptology expertise.

Some of the research practices employed at each of these research sites undoubtedly overlapped with one another. However, they also clearly exhibited divergence. The second research question essentially asks whether this divergence can be understood through the use of the ontological framework, and whether contributory expertise can be usefully thought of as multiple. My answer to this question is a qualified 'yes'. Although usually kept separate in the STS literature, the third wave understanding of contributory expertise shares key tenets with the ontological framework. They are both rooted in practices, and they both consider what is enacted by those practices to be real. Within the STS literature, one of the most popular strands of the ontological framework has come to think of the enactments that result from practices as multiple. In particular, prominent sociologists such as Annemarie Mol and John Law have argued that divergent practices will result in multiple 'realities'. Though the consequences of insisting on the enacting of multiple realities renders the adoption of a full philosophical ontology problematic, Michael Lynch's ontography can still be used as a useful way of adjusting the analytical focus.

It is possible to discern multiplicity in the descriptions of the practices used during the technical phase of the crypto wars, offered as part of the answer to the first research question. But, additional steps can also be taken to identify multiplicity more clearly. Again, in order to do this, ideas from the third wave and

the ontological framework can be used in conjunction. In particular, a process of sociological discrimination can be employed. Martin Weinel has argued that the knowledge that STS analysts have acquired about the nature of science constitutes a particular type of non-transmuted meta-expertise. Weinel argued that given that those working within STS have built a considerable knowledge of controversies over technological decision-making, sociological discrimination can be used to demarcate authentic and inauthentic controversies. I have argued that, given that STS has also built knowledge of how scientific work can differ within a discipline, sociological discrimination can also be used to delineate multiple enactments of contributory expertise.

I decided that the most appropriate of the available options for delineating different enactments of contributory expertise during the technical phase of the crypto wars was Philippe Larédo and Philippe Mustar's criteria for assembling laboratory activity profiles. Larédo and Mustar argued that activity profiles could be assembled based on a quantitative assessment of the extent to which laboratories engaged in: the production of certified knowledge; education, training activities and embodied knowledge; public research and the innovation process; the participation in public or collective goods and finalities; and research and public debate about science and technology. I argued that a qualitative assessment of whether there is evidence of engagement with each of these activities could be used to underpin the delineation of enactments of contributory expertise. When applied to the expertise enacted during the technical phase of the crypto wars, it was clear that, though the practices of some research sites overlapped, there were clear areas of divergence. Furthermore, when analysed using a thematic conceptual matrix (constructed on the basis of whether there exists: evidence; partial evidence; refuting evidence; or no evidence of particular practices) it was possible to characterize contributory expertise in other ways. In particular, by making a distinction between refuting evidence obtainable through interviews and absence of evidence noticeable in documents, absences could also be used to indicate the enacting of secret expertise given that it was symptomatic of practices designed to limit access to information.

Though it was possible to use sociological discrimination to identify the existence of multiple enactments of contributory expertise during the technical phase of the crypto wars, it may not be immediately clear why it was worth going to the trouble of doing so. The response to this forms the answer to the third research question.

In short, in the case of the crypto wars, considering the multiplicity of contributory expertise from the technical phase was worthwhile because it could be used as a platform for an analysis of the political phase. In contrast to the one previous description of the political phase of the crypto wars in the UK, I described events with respect to the expertise that was enacted during the technical phase. This revealed how each enactment was used during the political phase. My description of the political phase of the crypto wars was structured around two sets of government policy proposals related to cryptography. The first was a set of proposals for a nationwide network of TTPs. In light of developments within the field of cryptology, and the increasing use of large-scale electronic communications networks, the TTP proposals were framed as a way to balance the desire to promote electronic commerce, whilst preserving the capabilities of law enforcement bodies. Under the proposals, cryptography would be both regulated, and used as a means of regulation. The second policy proposed reforms to the UK's export control regime. Again, in light of the increasing use of large-scale electronic communications networks, the changes to export controls were framed as a way of addressing export through intangible means. Unlike the TTP proposals, cryptography was not instrumental to the proposed regime, but would be one of the many scientific fields affected.

The multiple cryptology expertises produced during the technical phase were used in a variety of different ways during the formal and informal processes that characterized the activity of the political phase of the crypto wars. The expertise enacted within the Data Security Group was absent from the political phase. This absence can be seen as being actively constructed, given that research practices within the Data Security Group were designed to enact expertise on matters that did not feature prominently in the controversy. Furthermore, practices were sufficiently constraining so as to prohibit the enacting of expertise outside of testing and accreditation. Expertise enacted within the Information Security Group was used to produce a technology to underpin the TTP proposals. During the technical phase, the group had used practices to enact contributory expertise that was sensitive to the requirements of industry and government, and were able to align their research with the goals of external partners. Expertise enacted within the Security Group was initially ignored during the formation of the TTP proposals, but was later used to criticise them and the underlying technologies. The group had previously enacted expertise that related to the real-world use of security systems, and were able to translate this research into arguments suited to a policy

debate. Members of the group were also successful in influencing the export control debate by using their cryptology expertise to construct arguments about the preservation of academic freedom. Expertise enacted within CESC was used to both produce technologies that could be used in the TTP proposals, and to advise the government on the formation of their policies. However, the expertise enacted within CESC also caused uncertainty during the crypto wars. The level of secrecy surrounding their cryptology research during the technical phase meant that other experts and the public were prevented from knowing details about the quality of their expertise, and the motivations behind its enactment.

In all four cases, it was possible to form plausible links between the enactments of expertise during the technical phase, and the way it was used during the political phase. This lends extra credence to the idea that expertise is enacted, given that Mol (2002) used the term to denote the performing or carrying out of practices, as well as to denote what results from them. Linking expertise and its uses therefore allows for the enactment of expertise to be thought of as extending across both the technical and the political phases. On this understanding, recognizing the multiplicity of contributory expertises that can exist during the technical phase would, in principle, allow sociological discrimination to be used to inform decisions during the political phase about where to seek expertise on a particular issue. This is especially important during controversies over technological decision-making, as the broad framing of an issue - as required by elective modernism (Collins et al. 2010) - may make a number of propositional questions relevant.

This is a positive potential outcome of recognizing the multiplicity of contributory expertise that can exist during a technical phase. Recognizing this multiplicity also allows for an analytical appreciation of cases where decisions about where to source expertise can have negative consequences. A closer examination of the political phase of the crypto wars showed that some actors felt that their expertise was being excluded, and that drawing on the secret expertise of CESC was inappropriate. It was argued that the expertise enacted by CESC was of dubious quality, and that they were attempting to force the adoption of technologies that would serve their own particular ends. In response, individuals from the Cambridge group (and others) engaged in a form of scientist activism. They mounted a successful opposition to the government's proposals through active engagement with the formal and informal processes of the political phase.

The events of the political phase of the crypto wars can therefore be used to highlight some of the contradictions in the current formulation of elective modernism. In particular, it was shown how decisions about where to source expertise threatened to collapse the distinction between the technical and political phase, and gave reason to question the notion that the political phase should consider as much of the available expertise from the technical phase as possible. As a way of creating an extra analytical foothold on these issues, I argued that, in the case of the crypto wars, it is possible to characterize a particular element of the relationship between the technical and political phases - namely, how the political phase came to know about the expertise enacted during the technical. I argued that it was possible to distinguish between expertise that was: not transferred; commissioned; delivered; and transferred invisibly. Under this scheme, as it was absent from the political phase, the expertise enacted within the Data Security Group was not transferred. As the expertise within the Information Security Group was, in part, enacted through a DTI-funded project, it was commissioned. As the expertise enacted within the Security Group was initially ignored, but was eventually integrated into the proposals during the consultation procedures, it was delivered. Finally, as the way in which the expertise of CESC was transferred was not made clear, it was transferred invisibly.

I have argued that it is possible to use this extra layer of data about the transfer process - together with the ideas expressed in the answers to the other research questions - to suggest two additional elective modernist ideas: the Minimum Transfer Requirement; and the problem of expert discrimination. The Minimum Transfer Requirement is an attempt to provide a means of determining whether it can be claimed by decision-makers that they have sufficiently considered the expertise enacted during the technical phase of a controversy. It states that, in light of the problems that resulted from the ignoring of particular enactments of expertise during the crypto wars, for it to be claimed that the available expertise has been sufficiently considered, decisions cannot be made solely on the basis of either commissioned expertise, or expertise that was transferred invisibly. The problem of expert discrimination recognizes that, in some cases - as with the secret expertise enacted by CESC - it may not be possible to use non-transmuted meta-expertise to assess its quality. In extreme cases, even the processes typically associated with domain-specific discrimination may not be available. In such cases, I have suggested, in order to preserve the distinction between the formative intentions of the

technical and political phases, the case can be made for prohibiting the transfer of this expertise from the technical to the political phase.

10.3 Reflections, Limitations, and Further Work

There are, of course, some limitations to the arguments that I have made and the methods I have used to arrive at them. In this final section, I will reflect on these limitations, and in some cases, use them to suggest ideas for further work. In many cases, it is clear that more case studies using the third wave are needed to probe both the ideas expressed in this thesis, and the ideas that underpin elective modernism more generally.

I will begin by reflecting on my methodological approach. Although, as was discussed in the chapter 3, some concerns over the generalizability of findings that result from case studies are perhaps exaggerated (Flyvbjerg 2006), they cannot be entirely dismissed. It is by no means clear that each of my claims would be equally visible in case studies based on, say, different fields of scientific research, or different controversies over technological decision-making. This thesis does not, and cannot, show that multiple research practices always produce multiple contributory expertises, or that multiple contributory expertises are always used in different ways and for different purposes during controversies over technological decision making. Nor does it claim that the types of contributory expertise that emerge from the descriptions are comprehensive. The purpose of this thesis was to work towards the development of new third wave ideas grounded in empirical data, rather than to confirm a pre-existing theory or relationship.

The main barrier to the generalizability of this case study could be seen as the emphasis on secrecy. The enacting of secret expertise was at the heart of what made the crypto wars distinctive, and the problem of expert discrimination aims to identify the problems associated with its use during controversies. However, it is possible to question the extent to which the enacting of expertise in secret is a pressing concern, especially given that Merton (1973) identified ‘communalism’ as one of the norms that guides science. Though Merton was right to stress the importance of communalism, this does not mean that secrecy is incompatible with science. In identifying the existence of a set of counter-norms, Mirtoff (1974) argued that secrecy is important to academic science because it allows scientists to

carry out a programme of research without having to worry about those working on something similar. This is also broadly true of commercial science - where keeping research secret can be used to protect a potential source of revenue - and of defence science - where keeping research secret can protect military power. A strand of literature within STS is emerging that aims to better understand these aspects of the relationship between secrecy and science. Given that secrecy guides much scientific activity, I believe that it is appropriate for the third wave to address the consequences of the use of secret expertise during controversies over technological decision-making.

Remaining concerns over generalizability and the importance of secrecy can be addressed by conducting more case studies. If this is done, there does not seem to be any good reason to limit case studies of the nature of contributory expertise to any particular scientific, technological, or medical field. Future case studies that probe multiple contributory expertises may even be able produce robust categories that could be incorporated into the periodic table of expertise. The more general categories that make up the research compass card could provide a useful basis for this. In terms of testing the existence of these categories of contributory expertise, it may be possible to use experimental methods similar to those used to test for interactional expertise (Collins et al. 2006).

Turning now to the specific methods used during this case study, in chapter 3 I outlined the way in which the fieldwork and analysis stages would be structured (see Table 3.2). Although this basic structure was broadly adhered to, there were some deviations. For example, the line between the first and second stage was blurred by the fact that some interviewees were involved in both the technical and political phases. As both were covered in the same interview, in some cases, data relating to the second stage was gathered before the first stage was completed. Also, within the first and second stages, a process similar to snowballing occurred whereby information gathered during the interview process pointed towards documents that had not yet been examined. This blurred the separation of the documentary analysis and semi-structured interview sections of each stage. This kind of pattern can be potentially problematic, because it is easy to see how such a snowballing process might serve to reinforce certain themes and conceal others. However, although it is difficult to be certain, it would appear that to not explore new avenues of investigation after they have been revealed would have

prevented me from claiming that a saturation point had been reached in terms of gathering more data.

Attempting to examine the relationship between a technical phase and a political phase of a controversy, to a certain extent requires the use of historical methods. It would be difficult to justify studying a technical phase in real time, without knowing whether a political phase would emerge later. One consequence of using historical methods is that they produce data that emphasises macro-social, institutional and historical research practices, rather than micro-social research practices associated with experimentation. Though I have argued that institutional and historical research practices exert a powerful influence on how expertise is enacted, not being able to observe research practices from the technical phase undoubtedly resulted in a lack of highly detailed descriptions of particular processes. On a similar note, there were also issues related to asking interviewees direct questions about vague and slightly abstract concepts such as ‘expertise’ and ‘research practices’. One interviewee, when asked to describe their research practices, responded by stating that “they were just normal, really”. Therefore, questions about expertise and research practices often had to be approached from a different angle. Furthermore, though I was aware from the methodological literature that interviewees often struggle to remember events from the past, with the benefit of hindsight, I do not think I was quite prepared for the extent to which this would be the case. Some themes, about which I would have liked to know more, simply could not be developed further due to a lack of specific details.

Turning now the substantive claims made in this thesis, I have described the research practices used to enact contributory cryptology expertise at four research sites during the technical phase of the crypto wars. However, the historical and sociological information contained in this thesis cannot be considered a complete historical or sociological study of modern cryptology research. There is evidence to suggest that cryptology research was undertaken at a number of other sites from 1970 to 2000. For example, the industrial cryptology research carried out by Zergo (and later Baltimore) would have been a good addition. Furthermore, it is known that early UK government-funded computer security research was carried out at the Royal Radar Establishment in Malvern in the 1980s, and that this may have included research into cryptology. However, in both cases, there was insufficient available data to be able to investigate this research in detail. Similarly, though this thesis necessarily focussed on collective research practices, the work

of lone scientists carrying out research into cryptology may be of relevance to future understandings of the crypto wars. Also, due to the fact that the third wave defines the technical phase in line with the formative intentions of science, the contributory expertise of lawyers and policymakers, and their relationship to the crypto wars, was not discussed. In the case of the crypto wars, this was a significant omission, as much of the activity from groups such as FIPR, that was used to construct an opposition to the government's proposals, was informed by non-scientific expertise. Future work on the third wave will surely have to address this shortcoming, as these types of expertise (and others) are likely to be pivotal during many controversies over technological decision-making.

Returning to the study of cryptology research, more can also be said about other social and historical facets. To take just one example, this study has not addressed the users of these technologies. Of particular interest may be the way in which the designers and developers of these technologies 'configured' their users' ideas about security (Woolgar 1991). Ideas about how cryptology use is configured may also be able to shed light on why some chose not to use these technologies, given that it has been acknowledged that non-users can also shape their development (Wyatt 2003).

In terms of data collection for the political phase, there were some barriers that, in general, did not exist during the data collection for the technical. Documents pertaining to the technical phase were usually publicly accessible in archives, and scientists and other interviewees were mostly happy to give up their time to talk to me. For the political phase, documents were often stored in archives that were not publicly accessible. FoI requests to see these documents were occasionally refused, and those that were successful often returned documents that had been heavily redacted. Potential interviewees, including politicians, were either too busy or otherwise unavailable to talk. Therefore, the description of the political phase provided was based on publicly accessible information. Consequently, for the political phase, it was not possible to carry out and define a full process of sociological discrimination that mirrored that which was carried out for the technical phase, given the lack of symmetry in the data. Though it was possible to provide a detailed description of the public face of the political phase, a lack of information about internal processes meant that an attempt to describe the evidence that existed for certain 'principles' in a thematic conceptual matrix would have been comprised by the possibility that evidence of this principle had not been made public. One

suspected reason for this difficulty is that, despite reaching a legislative conclusion around 2000, the political phase of the crypto wars, has, in a sense, continued. The tensions at the heart of the crypto wars did not disappear after 2000, and in many ways, they are still with us today. Debates over the appropriate use of cryptology technologies continue. The recent publication of classified intelligence documents leaked by Edward Snowden, detailing as they did the post-2000 attempts by GCHQ and the NSA to covertly undermine cryptographic technologies, demonstrates this clearly. Some cryptography advocates see the controversy that I have described as the first of many crypto wars that are sure to take place in the future (e.g. Assange 2012). As a result, information pertaining to the political phase of the controversy retains both sensitivity and contemporary relevance. On this point, in carrying out this study, I have learnt a lot about the impact that symmetry and availability of data can have on the conclusions that one is able to reach. It is now much clearer to me why some aspects of society are better understood than others. When designing future studies, I will be sure to place more emphasis on initial investigations that aim to reveal the likely availability of relevant data.

The answer I gave to the second research question was a qualified ‘yes’. The reason for this qualification is that it should be acknowledged that some aspects of the ontological framework discussed in chapter 2 do not align well with third wave ideas about expertise. For instance, in describing how multiple enactments of atherosclerosis interacted with one other, Mol (2002) described a process of co-ordination during which a ‘superior’ enactment came to determine a course of treatment. There appears to be a superficial resemblance between this process and the closing of a controversy. However, Mol’s description of co-ordination appears to hinge on there actually being an identifiable and indisputable enactment that is superior, and does not describe how such an identification may be achieved. Of the enactments that were produced during the technical phase of the crypto wars, it would be very difficult to pick one out as being superior to the others. Indeed, in identifying multiplicity, fundamental differences were described that implied a certain degree of incommensurability between enactments. Similarly, the ‘different worlds’ argument (Law & Mol 2011) did not align well with the descriptions of the work of the technical phase of the crypto wars. Whilst it would have been possible to delineate enactments of contributory expertise along the lines of ‘materials’ and ‘qualities’, it did not seem possible to identify a convincing distinction between their ‘spatial relations’ and their ‘staging of time’. As a result, these

aspects of multiplicity and the ontological framework could not be satisfactorily integrated into this case study. It was for this reason, and others, that I chose to use what Michael Lynch (2013) has termed ‘ontography’. Ontography refers to the adoption of the vocabulary of the ontological framework, rather than a complete acceptance of its heavily criticised philosophical consequences. However, in shedding the philosophical aspects of the ontological framework, it could be argued that it loses its *raison d’être*. After carrying out this study, I agree with Lynch that there is value in ontography due to the fact that it can be used to reorient descriptions. Though the crypto wars could have been studied using the epistemological framework, it is difficult to see how it could have been used to arrive at a similar appreciation of the fundamental ways in which contributory expertise differed, and the consequences that this had. I argue that it would have made little sense to claim that those researching cryptology were essentially carrying out investigations along the same lines, but each from a different point of view. That being said, it is also clear that more ontographical case studies need to be carried out before its legitimacy can be properly established.

Finally, the consequences of multiple contributory expertises during controversies over technological decision-making should be examined in more detail. Though the Minimum Transfer Requirement and the problem of expert discrimination were suggested as possible inclusions to elective modernism, they require further elaboration. In their current formulation, though I would argue that their application would be have been beneficial for the quality of debate during the crypto wars, it is likely that, in other scenarios, they would not have been. Similarly, though I have argued that recognising multiplicity could be used to guide decisions about where to source expertise during controversies, more work is needed to properly investigate and analyse the viability of this idea. Once again, these issues can be partially addressed through more case studies that attempt to use the third wave to understand and improve controversies over technological decision-making.

Appendix A

Information Sheet and Consent Form

Appendix B

Archival Sources

B.1 British Library

B.1.1 Collection: General Reference

- NPL Annual Reports and Accounts 1990/1991 - ZK.9.b.5764
- NPL Annual Reports and Accounts 1991/1992 - ZK.9.b.5764
- NPL Annual Reports and Accounts 1992/1993 - ZK.9.b.5764

B.1.2 Collection: Science, Business, and Technology

- NPL Data Security Group Bulletin 1991 - (P)PM580-E(13)
- NPL Data Security Group Bulletin 1992 - (P)PM580-E(13)
- NPL Data Security Group Bulletin 1993 - (P)PM580-E(13)
- NPL Data Security Group Bulletin 1994 - (P)PM580-E(13)
- NPL Data Security Group Bulletin 1995 - (P)PM580-E(13)
- NPL Data Security Group Bulletin 1996 - (P)PM580-E(13)
- Racal Review 1970 (6 issues) - P.523/183
- Racal Review 1971 (6 issues) - P.523/183
- Racal Review 1972 (6 issues) - P.523/183
- Racal Review 1973 (6 issues) - P.523/183
- Racal Review 1974 (6 issues) - P.523/183
- Racal Review 1975 (6 issues) - P.523/183

- Racal Review 1976 (6 issues) - P.523/183
- Racal Review 1977 (6 issues) - P.523/183
- Racal Review 1978 (6 issues) - P.523/183
- Racal Review 1979 (6 issues) - P.523/183
- Racal Review 1980 (6 issues) - P.523/183
- Racal Review 1981 (6 issues) - P.523/183
- Racal Review 1982 (6 issues) - P.523/183
- Racal Electronics for Defence 1982 - P.625/463
- Racal Electronics for Defence 1983 - P.625/463
- Racal Electronics for Defence 1984 - P.625/463
- Racal Electronics for Defence 1985 - P.625/463
- Racal Electronics for Defence 1986 - P.625/463
- Racal Electronics for Defence 1987 - P.625/463
- Racal Electronics for Defence 1988 - P.625/463
- Miscellaneous Racal Publications Not Catalogued Separately - ZA.9.d.252
- The Racal Handbook 1956-1975 - YK.1994.b.13305

B.1.3 Collection: Trade Literature

- NPL News 1975 - BS.38n/99
- NPL News 1976 - BS.38n/99
- NPL News 1977 - BS.38n/99
- NPL News 1978 - BS.38n/99
- NPL News 1979 - BS.38n/99
- NPL News 1980 - BS.38n/99
- NPL News 1981 - BS.38n/99
- NPL News 1982 - BS.38n/99
- NPL News 1983 - BS.38n/99
- NPL News 1984 - BS.38n/99
- NPL News 1985 - BS.38n/99
- NPL News 1986 - BS.38n/99
- NPL News 1987 - BS.38n/99
- NPL News 1988 - BS.38n/99
- NPL News 1989 - BS.38n/99
- NPL News 1990 - BS.38n/99

- NPL News 1991 - BS.38n/99
- NPL News 1992 - BS.38n/99

B.2 Cambridge University Archives

B.2.1 Collection: Archives of the Mathematical Laboratory and its successor, the Computer Laboratory

- Report on the functions of the Mathematical Laboratory, circulated to Faculty Board of Mathematics, with covering note - COMP 1/1
- Report of Mathematical Laboratory Committee to the Faculty Board of Mathematics on present conditions, proposals for improvements and plan of development - COMP 1/2
- Minutes and papers of the Laboratory Coordination Committee - COMP 1/5
- Minutes of Phoenix meetings - COMP 1/6
- Computing Service planning documents - COMP 1/7
- Minutes and papers of the Curriculum Committee - COMP 2/1/1-2
- Minutes and papers of the Teaching Committee - COMP 2/2
- University prospectus: Computer Science at Cambridge - COMP 2/6
- Papers relating to the Science Research Council grant for research staff and equipment - COMP 3/1/1
- Papers relating to the equipment grant from Xerox - COMP 3/1/2
- System evaluation documents - COMP 4/4
- Software Review meetings minutes - COMP 4/5
- Mainframe Division meetings minutes and papers - COMP 4/6
- Design Progress meetings minutes and papers - COMP 4/8/1-2
- Papers relating to EDSAC jubilee celebrations - COMP 7/1
- Papers and photographs relating to 'EDSAC 99', a conference to mark the 50th anniversary of the first program run on EDSAC - COMP 7/2

B.3 Imperial College London Archives

B.3.1 Collection: Donald Davies' Papers

- Personal - GB 98 B
- Notes of Miscellaneous Scientific Work - GB 98 B
- Data Security MS Notes - GB 98 B
- Public Key Ciphers - GB 98 B

B.4 National Archives

B.4.1 Collection: Records of the Department of Scientific and Industrial Research, the National Physical Laboratory

- Minutes of the Review Committee 1973 - DSIR 10/470
- Minutes of the Review Committee 1974 - DSIR 10/471
- Minutes of the Review Committee 1975 - DSIR 10/472
- Minutes of the Review Committee 1976 - DSIR 10/473
- Minutes of the Review Committee 1977 - DSIR 10/474
- Minutes of the Review Committee 1978 - DSIR 10/475
- Minutes of the Review Committee 1980 - DSIR 10/476
- Green Paper: Industrial Research and Development in Government Laboratories - DSIR 10/486
- Green Paper: Industrial Research and Development in Government Laboratories - DSIR 10/487
- Green Paper: A Framework for Government Research and Development - DSIR 10/488
- Advisory Board: Minutes and Papers 1969-1971 - DSIR 72/5
- Supervisory Board: Setting up and operation of the Board including some minutes and papers 1986-1988 - DSIR 72/13 - [FoI](#)
- Supervisory Board: Minutes and Papers 1987-1988 - DSIR 72/14 - [FoI](#)
- Supervisory Board: Minutes and Papers 1988-1990 - DSIR 72/15 - [FoI](#)
- Steering Board: Minutes and Papers 1990 - DSIR 72/16 - [FoI](#)

- Steering Board: Minutes and Papers 1991 - DSIR 72/17 - [FoI](#)
- Steering Board: Minutes and Papers 1992 - DSIR 72/18 - [FoI](#)
- Steering Board: Minutes and Papers 1993 - DSIR 72/19 - [FoI](#)
- Superintendents Meetings: Minutes 1970-1977 - DSIR 72/25
- Superintendents Meetings: Papers 1970-1974 - DSIR 72/37
- Superintendents Meetings: Papers 1975-1976 - DSIR 72/38
- Superintendents Meetings: Papers 1977 - DSIR 72/39
- Superintendents Meetings: Minutes and Papers 1978 - DSIR 72/40
- Superintendents Meetings: Minutes and Papers 1979 - DSIR 72/41
- Superintendents Meetings: Minutes and Papers 1980 - DSIR 72/42
- Superintendents Meetings: Minutes and Papers 1981 - DSIR 72/43
- Superintendents Meetings: Minutes and Papers 1982 - DSIR 72/44
- Superintendents Meetings: Minutes and Papers 1983 - DSIR 72/45 - [FoI](#)
- Superintendents Meetings: Minutes and Papers 1984 - DSIR 72/46 - [FoI](#)
- Superintendents Meetings: Minutes and Papers 1985 - DSIR 72/47 - [FoI](#)
- Superintendents Meetings: Minutes and Papers 1986 - DSIR 72/48 - [FoI](#)
- Superintendents Meetings: Minutes and Papers 1987 - DSIR 72/49 - [FoI](#)
- Superintendents Meetings: Minutes and Papers 1988 - DSIR 72/50 - [FoI](#)
- Superintendents Meetings: Minutes and Papers 1989 - DSIR 72/51 - [FoI](#)
- Superintendents Meetings: Minutes and Papers 1990 - DSIR 72/52 - [FoI](#)
- Directorate Committee: Minutes 1970 - DSIR 72/75
- Directorate Committee: Minutes 1971 - DSIR 72/76
- Directorate Committee: Minutes 1972 - DSIR 72/77
- Directorate Committee: Minutes 1973 - DSIR 72/78
- Directorate Committee: Minutes 1974 - DSIR 72/79
- Directorate Committee: Minutes 1975 - DSIR 72/80
- Directorate Committee: Minutes 1976 - DSIR 72/81
- Directorate Committee: Minutes 1977 - DSIR 72/82
- Directorate Committee: Minutes 1978 - DSIR 72/83
- Directorate Committee: Minutes 1979 - DSIR 72/84
- Directorate Committee: Minutes 1980 - DSIR 72/85
- Directorate Committee: Minutes 1981 - DSIR 72/86
- Directorate Committee: Minutes 1982 - DSIR 72/87
- Directorate Committee: Minutes 1983 - DSIR 72/88 - [FoI](#)

- Directorate Committee: Minutes 1984 - DSIR 72/89 - [FoI](#)
- Directorate Committee: Minutes 1985 - DSIR 72/90 - [FoI](#)
- Standards Committee: Minutes and Papers 1976 - [FoI](#)
- Standards Committee: Minutes and Papers 1977-1979 - DSIR 72/135
- Standards Committee: Agendas, Minutes, Correspondence and Papers 1975-1976 - DSIR 72/143

B.4.2 Collection: Records of the Home Office, Ministry of Home Security, and Related Bodies

- National Computing Centre Seminar on the Protection of Data by Cryptography - HO 261/198

B.5 Parliamentary Archives

B.5.1 Collection: Records of the House of Commons

- Trade and Industry Committee: Seventh Report, 1998-1999: “Building Confidence in Electronic Commerce”: The Government’s Proposals (HC 187) - HC/CL/CO/CZ/1/19 - [FoI](#)
- Trade and Industry Committee: Tenth Report, 1998-1999: Electronic Commerce (HC 648) - HC/CL/CO/CZ/1/22 - [FoI](#)
- Trade and Industry Committee: Fourteenth Report, 1998-1999: The Draft Electronic Communications Bill (HC 862) - HC/CL/CO/CZ/1/25 - [FoI](#)

B.6 Royal Holloway Archives

B.6.1 Collection: Royal Holloway College Papers

- Minutes of the Meeting of the Academic Board 1970 - RHC/AL/100/10
- Minutes of the Meeting of the Academic Board 1971 - RHC/AL/100/10
- Minutes of the Meeting of the Academic Board 1972 - RHC/AL/100/10
- Minutes of the Meeting of the Academic Board 1973 - RHC/AL/100/10

- Minutes of the Meeting of the Academic Board 1974 - RHC/AL/100/10
- Minutes of the Meeting of the Academic Board 1975 - RHC/AL/100/11
- Minutes of the Meeting of the Academic Board 1976 - RHC/AL/100/11
- Minutes of the Meeting of the Academic Board 1977 - RHC/AL/100/11
- Minutes of the Meeting of the Academic Board 1978 - RHC/AL/100/11
- Minutes of the Meeting of the Academic Board 1979 - RHC/AL/100/11
- Minutes of the Meeting of the Academic Board 1980 - RHC/AL/100/13
- Minutes of the Meeting of the Academic Board 1981 - RHC/AL/100/13
- Minutes of the Meeting of the Academic Board 1982 - RHC/AL/100/13
- Minutes of the Meeting of the Academic Board 1983 - RHC/AL/100/13
- Minutes of the Meeting of the Academic Board 1984 - RHC/AL/100/13
- Minutes of the Meeting of the Academic Board 1985 - RHC/AL/100/14
- Minutes of the Meeting of the Academic Planning Committee 1970 - RHC/18/19/3/4
- Minutes of the Meeting of the Academic Planning Committee 1971 - RHC/18/19/3/4
- Minutes of the Meeting of the Academic Planning Committee 1972 - RHC/18/19/3/4
- Minutes of the Meeting of the Academic Planning Committee 1973 - RHC/18/19/3/4
- Minutes of the Meeting of the Academic Planning Committee 1974 - RHC/18/19/3/4
- Minutes of the Meeting of the Academic Planning Committee 1975 - RHC/18/19/3/4
- Minutes of the Meeting of the Academic Planning Committee 1976 - RHC/18/19/3/4
- Minutes of the Meeting of the Academic Planning Committee 1977 - RHC/18/19/3/4
- Minutes of the Meeting of the Academic Planning Committee 1978 - RHC/18/19/3/4
- Minutes of the Meeting of the Academic Planning Committee 1979 - RHC/18/19/3/4
- Minutes of the Meeting of the Academic Planning Committee 1980 - RHC/18/19/3/4
- Minutes of the Meeting of the Academic Planning Committee 1981 - RHC/18/19/3/4
- Minutes of the Meeting of the Academic Planning Committee 1982 - RHC/18/19/3/4
- Minutes of the Meeting of the Faculty of Science 1970 - RHC/AL/170/4
- Minutes of the Meeting of the Faculty of Science 1971 - RHC/AL/170/4
- Minutes of the Meeting of the Faculty of Science 1972 - RHC/AL/170/4
- Minutes of the Meeting of the Faculty of Science 1973 - RHC/AL/170/4
- Minutes of the Meeting of the Faculty of Science 1974 - RHC/AL/170/4
- Minutes of the Meeting of the Faculty of Science 1975 - RHC/AL/170/4
- Minutes of the Meeting of the Faculty of Science 1976 - RHC/AL/170/4
- Minutes of the Meeting of the Faculty of Science 1977 - RHC/AL/170/4
- Minutes of the Meeting of the Faculty of Science 1978 - RHC/AL/170/4

- Minutes of the Meeting of the Faculty of Science 1979 - RHC/AL/170/4
- Minutes of the Meeting of the Faculty of Science 1980 - RHC/AL/170/4
- Minutes of the Meeting of the Faculty of Science 1981 - RHC/AL/170/4
- Minutes of the Meeting of the Faculty of Science 1982 - RHC/AL/170/4
- Minutes of the Meeting of the Faculty of Science 1983 - RHC/AL/170/4
- Minutes of the Meeting of the Faculty of Science 1984 - RHC/AL/170/4

B.6.2 Collection: Royal Holloway and Bedford New College Papers

- Minutes of the Meeting of the Academic Board 1986 - HB/CP/2/2/1/1
- Minutes of the Meeting of the Academic Board 1987 - HB/CP/2/2/1/1
- Minutes of the Meeting of the Academic Board 1988 - HB/CP/2/2/1/1
- Minutes of the Meeting of the Academic Board 1989 - HB/CP/2/2/1/2
- Minutes of the Meeting of the Academic Board 1990 - HB/CP/2/2/1/2
- Minutes of the Meeting of the Academic Board 1991 - HB/CP/2/2/1/2
- Minutes of the Meeting of the Joint Academic Planning Committee 1983 - HB/CP/1/1/4
- Minutes of the Meeting of the Joint Academic Planning Committee 1984 - HB/CP/1/1/4
- Minutes of the Meeting of the Joint Academic Planning Committee 1985 - HB/CP/1/1/4
- Minutes of the Meeting of the Academic Planning Committee 1986 - Uncatalogued
- Minutes of the Meeting of the Academic Planning Committee 1987 - Uncatalogued
- Minutes of the Meeting of the Academic Planning Committee 1988 - Uncatalogued
- Minutes of the Meeting of the Academic Planning Committee 1989 - Uncatalogued
- Minutes of the Meeting of the Academic Planning Committee 1990 - Uncatalogued
- Minutes of the Meeting of the Academic Planning Committee 1991 - Uncatalogued
- Minutes of the Meeting of the Faculty of Science 1985 - Uncatalogued

- Minutes of the Meeting of the Faculty of Science 1986 - Uncatalogued
- Minutes of the Meeting of the Faculty of Science 1987 - Uncatalogued
- Minutes of the Meeting of the Faculty of Science 1988 - Uncatalogued
- Minutes of the Meeting of the Faculty of Science 1989 - Uncatalogued
- Minutes of the Meeting of the Faculty of Science 1990 - Uncatalogued
- Minutes of the Meeting of the Faculty of Science 1991 - Uncatalogued

B.7 Wayback Machine

- Website of the National Physical Laboratory
www.npl.co.uk
- Website of the Cambridge Computer Laboratory
www.cl.cam.ac.uk
- Website of the Communications-Electronics Security Group
www.cesg.gov.uk
- Archives of the UK Cryptography Policy Discussion Group 1997
www.chiark.greenend.org.uk/mailman/listinfo/ukcrypto
- Archives of the UK Cryptography Policy Discussion Group 1998
www.chiark.greenend.org.uk/mailman/listinfo/ukcrypto
- Archives of the UK Cryptography Policy Discussion Group 1999
www.chiark.greenend.org.uk/mailman/listinfo/ukcrypto
- Archives of the UK Cryptography Policy Discussion Group 2000
www.chiark.greenend.org.uk/mailman/listinfo/ukcrypto
- Archives of the UK Cryptography Policy Discussion Group 2001
www.chiark.greenend.org.uk/mailman/listinfo/ukcrypto

Bibliography

- Abbate, J. (1999), *Inventing the Internet*, The MIT Press, Cambridge.
- Abelson, H., Anderson, R., Bellovin, S. M., Benaloh, J., Blaze, M., Diffie, W., Gilmore, J., Neumann, P. G., Rivest, R. L., Schiller, J. I. & Schneier, B. (1997), 'The Risks of Key Recovery, Key Escrow, and Trusted Third-Party Encryption', *World Wide Web Journal* **2**(3), 241–257.
- Abrams, P. (1982), *Historical Sociology*, Open Books, Shepton Mallet.
- Ahmed, H. (2013), *Cambridge Computing: The First 75 Years*, Third Millennium, London.
- Aldrich, R. J. (2001), 'GCHQ and Sigint in Early Cold War, 1945-70', *Intelligence and National Security* **16**(1), 67–96.
- Aldrich, R. J. (2010), *GCHQ: The Uncensored Story of Britain's Most Secret Intelligence Agency*, Harper, London.
- Altheide, D. L. (2004), Ethnographic Content Analysis, in M. S. Lewis-Beck, A. E. Bryman & T. F. Liao, eds, 'The SAGE Encyclopedia of Social Science Research Methods', SAGE, Thousand Oaks.
- Amsterdamska, O. (1990), 'Surely You Are Joking, Monsieur Latour!', *Science, Technology and Human Values* **15**(4), 495–504.
- Anderson, R. (1994), 'Why Cryptosystems Fail', *Communications of the ACM* **37**(11), 32–40.
- Anderson, R. (1997a), 'Problems with the NHS Cryptography Strategy', University of Cambridge.
- Anderson, R. (1997b), 'Response to the DTI TTP Proposals', University of Cambridge.

- Anderson, R. (1998), 'Letter to Mr Peter Mandelson, MP, Secretary of State for Trade and Industry'.
- Anderson, R. (2008), *Security Engineering: A Guide to Building Dependable Distributed Systems (Second Edition)*, Wiley, Indianapolis.
- Anderson, R. (2012), 'Ross Anderson's Home Page'.
URL: <http://www.cl.cam.ac.uk/~rja14/>
- Anderson, R. & Biham, E. (1996a), 'Tiger: A Fast New Hash Function', *Lecture Notes in Computer Science* **1039**, 89–97.
- Anderson, R. & Biham, E. (1996b), 'Two Practical and Provably Secure Block Ciphers: BEAR and LION', *Lecture Notes in Computer Science* **1039**, 113–120.
- Anderson, R. & Bond, M. (2004), Protocol Analysis, Composability and Computation, in A. Herbert & K. Spärck Jones, eds, 'Computer Systems: Theory Technology and Applications', Springer, New York, pp. 15–19.
- Anderson, R., Crispo, B., Lee, J.-H., Manifavas, C., Matyas, V. & Petitcolas, F. (1999), *The Global Internet Trust Register: 1999 Edition*, The MIT Press, Cambridge.
- Anderson, R. & Manifavas, C. (1997), 'Chameleon - A New Kind of Stream Cipher', *Lecture Notes in Computer Science* **1267**, 107–113.
- Anderson, R. & Needham, R. M. (1995), 'Programming Satan's Computer', *Computer Science Today* **1000**, 426–440.
- Anderson, R. & Roe, M. (1997), The GCHQ Protocol and its Problems, in W. Fumy, ed., 'Proceedings of the International Conference on the Theory and Application of Cryptographic Techniques, 11-15 May 1997, Konstanz', Springer, New York, pp. 134–148.
- Ashmore, M. (2005), 'The Left Inside/The Left-Hand Side', *Social Studies of Science* **35**(5), 827–830.
- Assange, J. (2012), *Cypherpunks: Freedom and the Future of the Internet*, OR, New York.
- Atkinson, P. & Coffey, A. (2011), Analysing Documentary Realities, in D. Silverman, ed., 'Qualitative Research (Third Edition)', SAGE, London, pp. 77–92.

- Back, A. (1998), 'the spooks _do_ want your keys', ukcrypto Mailing List.
- Balmer, B. (2012), *Secrecy and Science: A Historical Sociology of Biological and Chemical Warfare*, Ashgate, Farnham.
- Barlow, J. P. (1996), 'A Declaration of the Independence of Cyberspace', Electronic Frontier Foundation.
- Bazerman, C. (1988), *Shaping Written Knowledge: The Genre and Activity of the Experimental Article in Science*, University of Wisconsin Press, Madison.
- Beaulieu, A., Scharnhorst, A. & Wouters, P. (2007), 'Not Another Case Study: A Middle-Range Interrogation of Ethnographic Case Studies in the Exploration of E-Science', *Science, Technology and Human Values* **32**(6), 672–692.
- Becher, T. & Trowler, P. R. (2001), *Academic Tribes and Territories: Intellectual Enquiry and the Culture of Disciplines (Second Edition)*, Open University Press, Buckingham.
- Beck, U. (1992), *Risk Society: Towards A New Modernity*, SAGE, London.
- Beker, H. J. & Piper, F. C. (1985), *Secure Speech Communications*, Academic Press, London.
- Beker, H. & Piper, F. (1982), *Cipher Systems: The Protection of Communications*, Northwood, London.
- Bell, D. A. & Olding, S. E. (1978), 'An Annotated Bibliography of Cryptography', National Physical Laboratory.
- Bennett, G. (2002), 'Declassification and Release Policies of the UK's Intelligence Agencies', *Intelligence and National Security* **17**(1), 21–32.
- Bergmark, A. & Lundstrom, T. (2002), 'Education, Practice and Research: Knowledge and Attitudes to Knowledge of Swedish Social Workers', *Social Work Education* **21**(3), 359–373.
- Bingham, C. (1987), *The History of Royal Holloway College 1886-1986*, Constable, London.
- Blackburn, S. (2005), *The Oxford Dictionary of Philosophy (Second Edition)*, Oxford University Press, Oxford.

- Blackstone, T. & Plowden, W. (1988), *Inside the Think Tank: Advising the Cabinet 1971-1983*, Heinemann, London.
- Blanchette, J.-F. (2012), *Burdens of Proof: Cryptographic Culture and Evidence Law in the Age of Electronic Documents*, The MIT Press, Cambridge.
- Blaze, M. (1994), Protocol Failure in the Escrowed Encryption Standard, in D. Denning, ed., 'Proceedings of the 2nd ACM Conference on Computer and Communications Security, 2-4 November 1994, Fairfax', Association for Computing Machinery, New York, pp. 59–67.
- Blaze, M. (2011), Key Escrow from a Safe Distance: Looking Back at the Clipper Chip, in R. Hobbes, ed., 'Proceedings of the 27th Annual Computer Security Applications Conference, 5-9 December 2011, Orlando', Association for Computing Machinery, New York, pp. 317–321.
- Bloor, D. (1976), *Knowledge and Social Imagery*, Routledge, London.
- Boden, R., Cox, D., Georghiou, L. & Barker, K. (2001), Administrative Reform of United Kingdom Government Research Establishments: Case Studies of New Organisational Forms, in D. Cox, P. Gummett & K. Barker, eds, 'Government Laboratories: Transition and Transformation', IOS Press, Amsterdam, pp. 77–97.
- Boden, R., Cox, D. & Nedeva, M. (2006), 'The Appliance of Science? New Public Management and Strategic Change', *Technology Analysis & Strategic Management* **18**(2), 125–141.
- Boden, R., Cox, D., Nedeva, M. & Barker, K. (2004), *Scrutinising Science: The Changing UK Government of Science*, Palgrave Macmillan, Basingstoke.
- Boden, R., Gummett, P., Cox, D. & Barker, K. (1998), 'Men in White Coats... Men in Grey Suits: New Public Management and the Funding of Science and Technology Services to the UK Government', *Accounting, Auditing & Accountability Journal* **11**(3), 267–291.
- Bosch, O. (2005), Technology Transfer and Arms Control: The Impact of Controls on Encryption, PhD Thesis, University of Reading.
- Bowden, C. & Akdeniz, Y. (1999), Cryptography and Democracy: Dilemmas of Freedom, in Liberty, ed., 'Liberating Cyberspace: Civil Liberties, Human Rights, and the Internet', Pluto, London, pp. 81–125.

- Brooks, H. (1964), *The Scientific Adviser*, in R. Gilpin & C. Wright, eds, 'Scientists and National Policy-Making', Columbia University Press, New York, pp. 72–96.
- Brown, K. J. (2009), 'Freedom of Information as a Research Tool: Realising its Potential', *The Howard Journal* **48**(1), 88–91.
- Bryant, J. M. (1994), 'Evidence and Explanation in History and Sociology: Critical Reflections on Goldthorpe's Critique of Historical Sociology', *British Journal of Sociology* **45**(1), 3–19.
- Bryman, A. (2008), *Social Research Methods (Third Edition)*, Oxford University Press, Oxford.
- Burningham, K. (1998), 'A Noisy Road or Noisy Resident? A Demonstration of the Utility of Social Constructivism for Analysing Environmental Problems', *Sociological Review* **46**(3), 536–563.
- Burns, E. (2010), 'Developing Email Interview Practices in Qualitative Research', *Sociological Research Online* **15**(4).
- Burrows, M., Abadi, M. & Needham, R. M. (1990), 'A Logic of Authentication', *ACM Transactions on Computer Systems* **8**(1), 18–36.
- Butler, L. (1981), 'Letter to Professor Randolph Quirk, CBE, FBA, Vice-Chancellor of the University of London'.
- Cabinet Office (2013), 'Government Security Classifications', Cabinet Office.
- Cachia, M. & Millward, L. (2011), 'The Telephone Medium and Semi-Structured Interviews: A Complementary Fit', *Qualitative Research in Organizations and Management* **6**(3), 265–277.
- Calhoun, C. (2006), 'The University and the Public Good', *Thesis Eleven* **84**(7), 7–43.
- Callon, M. (1986), Some Elements of a Sociology of Translation: Domestication of the Scallops and the Fishermen of St. Briec Bay, in J. Law, ed., 'Power, Action and Belief: A New Sociology of Knowledge?', Routledge, London, pp. 196–234.
- Campbell-Kelly, M. (1981), 'Programming the Pilot ACE: Early Programming Activity at the National Physical Laboratory', *IEEE Annals of the History of Computing* **3**(2), 133–162.

- Campbell-Kelly, M. (1986), 'Interview with Donald Davies'.
URL: <http://www.cbi.umn.edu/oh/pdf.phtml?id=312>
- Campbell-Kelly, M. (1988), 'Data Communications at the National Physical Laboratory (1965-1975)', *IEEE Annals of the History of Computing* **9**, 221–247.
- Campbell-Kelly, M. (1992), 'The Airy Tape: An Early Chapter in the History of Debugging', *IEEE Annals of the History of Computing* **14**(4), 16–26.
- Campbell-Kelly, M. (2003), *From Airline Reservations to Sonic the Hedgehog: A History of the Software Industry*, The MIT Press, Cambridge.
- Campbell-Kelly, M. (2008), 'Pioneer Profiles - Donald Davies', *Computer Resurrection* **44**.
- Campbell-Kelly, M. & Aspray, W. (2004), *Computer: A History of the Information Machine (Second Edition)*, Westview Press, Boulder.
- Campbell-Kelly, M., Aspray, W., Ensmenger, N. & Yost, J. R. (2013), *Computer: A History of the Information Machine (Third Edition)*, Westview Press, Boulder.
- Caplan, N. (1979), 'The Two Communities Theory and Knowledge Utilization', *American Behavioral Scientist* **22**(3), 459–470.
- Carr, E. H. (1961), *What is History?*, Cambridge University Press, Cambridge.
- Ceruzzi, P. E. (2003), *A History of Modern Computing (Second Edition)*, The MIT Press, Cambridge.
- Cocks, C. (1973), 'A Note on 'Non-Secret Encryption'', Communications-Electronics Security Group.
- Cocks, C. (1997), Split Knowledge Generation of RSA Parameters, in 'IMA Conference on Cryptography and Coding'.
- Cocks, C. (2001), 'An Identity Based Encryption Scheme Based on Quadratic Residues', *Lecture Notes in Computer Science* **2260**, 360–363.
- Collins, H. (1975), 'The Seven Sexes: A Study in the Sociology of a Phenomenon, or the Replication of Experiments in Physics', *Sociology* **9**(2), 205–224.

- Collins, H. (1981a), 'Son of Seven Sexes: The Social Destruction of a Physical Phenomena', *Social Studies of Science* **11**(1), 33–62.
- Collins, H. (1981b), 'Stages in the Empirical Programme of Relativism', *Social Studies of Science* **11**(1), 3–10.
- Collins, H. (1985), *Changing Order: Replication and Induction in Scientific Practice*, SAGE, London.
- Collins, H. (1992), *Changing Order: Replication and Induction in Scientific Practice*, University of Chicago Press, Chicago.
- Collins, H. (1998), 'The Meaning of Data: Open and Closed Evidential Cultures in the Search for Gravitational Waves', *American Journal of Sociology* **104**, 293–338.
- Collins, H. (2004), *Gravity's Shadow: The Search for Gravitational Waves*, University of Chicago Press, Chicago.
- Collins, H. (2007a), 'Bicycling on the Moon: Collective Tacit Knowledge and Somatic-Limit Tacit Knowledge', *Organization Studies* **28**(2), 257–262.
- Collins, H. (2007b), 'Case Studies of Expertise and Experience', *Studies in History and Philosophy of Science Part A* **38**(4), 615–760.
- Collins, H. (2010), *Tacit and Explicit Knowledge*, University of Chicago Press, Chicago.
- Collins, H. (2014a), *Are We All Scientific Experts Now?*, Polity Press, Cambridge.
- Collins, H. (2014b), 'Code of Practice for Interviews'.
URL: <http://www.cardiff.ac.uk/socsi/contactsandpeople/harrycollins/code-of-practise.html>
- Collins, H. & Evans, R. (2002), 'The Third Wave of Science Studies: Studies of Expertise and Experience', *Social Studies of Science* **32**(2), 235–296.
- Collins, H. & Evans, R. (2003), 'King Canute Meets the Beach Boys: Responses to The Third Wave', *Social Studies of Science* **33**(3), 435–452.
- Collins, H. & Evans, R. (2007), *Rethinking Expertise*, University of Chicago Press, Chicago.

Collins, H., Evans, R., Ribeiro, R. & Hall, M. (2006), 'Experiments with Interactional Expertise', *Studies in History and Philosophy of Science Part A* **37**(4), 656–674.

Collins, H. & Pinch, T. (1993), *The Golem: What Everyone Should Know About Science*, Cambridge University Press, Cambridge.

Collins, H. & Weinel, M. (2011), 'Transmuted Expertise: How Technical Non-Experts Can Assess Experts and Expertise', *Argumentation* **25**(3), 401–413.

Collins, H., Weinel, M. & Evans, R. (2010), 'The Politics and Policy of the Third Wave: New Technologies and Society', *Critical Policy Studies* **4**(2), 185–201.

Collins, H., Weinel, M. & Evans, R. (2011), 'Object and Shadow: Responses to the CPS Critiques of Collins, Weinel and Evans', 'Politics and policy of the Third Wave', *Critical Policy Studies* **5**(3), 340–348.

Communications-Electronics Security Group (1996), 'Securing Electronic Mail within HMG - Part I: Infrastructure and Protocol (Draft C)', Communications-Electronics Security Group.

Communications-Electronics Security Group (1998a), 'CESG Cryptographic Services'.

URL: http://www.cesg.gov.uk/cesghtml/crypt_ce.htm

Communications-Electronics Security Group (1998b), 'The history of Non-Secret Encryption'.

URL: <https://web.archive.org/web/19980415070829/http://www.cesg.gov.uk/storynse.htm>

Communications-Electronics Security Group (1998c), 'UK ITSEC Scheme Certification Report No 98/84: CASM CryptServe', Communications-Electronics Security Group.

Communications-Electronics Security Group (2012), 'History of CESG'.

URL: <http://www.cesg.gov.uk/AboutUs/Pages/history-CESG.aspx>

Computer Security Group (1998a), 'Computer Security Group - Introduction'.

URL: <http://web.archive.org/web/19980219210836/http://www.cl.cam.ac.uk/Research/Security/>

Computer Security Group (1998b), 'NetCard Project'.

URL: <http://web.archive.org/web/19980219215551/http://www.cl.cam.ac.uk/users/cm213/P>

- Conley, T. M. (1990), *Rhetoric in the European Tradition*, University of Chicago Press, Chicago.
- Coppersmith, D. (1994), 'The Data Encryption Standard (DES) and its Strength Against Attacks', *IBM Journal of Research and Development* **38**(3), 243–250.
- Course, M. (2010), 'Of Words and Fog: Linguistic Relativity and Amerindian Ontology', *Anthropological Theory* **10**(3), 247–263.
- Court, J. & Young, J. (2003), 'Bridging Research and Policy: Insights from 50 Case Studies', Overseas Development Institute.
- Creswell, J. W. (2007), *Qualitative Inquiry and Research Design: Choosing Among Five Approaches (Second Edition)*, SAGE, Thousand Oaks.
- Croarken, M. G. (1992), 'The Emergence of Computing Science Research and Teaching at Cambridge, 1936-1949', *IEEE Annals of the History of Computing* **14**(4), 10–15.
- Data Security Group (1991*a*), 'An Implementation of MAA from a VDM Specification', National Physical Laboratory.
- Data Security Group (1991*b*), 'Data Security Group Bulletin', National Physical Laboratory.
- Data Security Group (1993), 'Standards and Conformance Testing in Data Security: Results of Initial Programme of Investigation', National Physical Laboratory.
- Davies, D. W. (1990), 'Evolution of the ACE', *Computer Resurrection* **2**.
- Davies, D. W. (1993), 'Early Computer Development at NPL', *Computer Resurrection* **8**.
- Davies, D. W. & Clayden, D. O. (1983), 'A Message Authenticator Algorithm Suitable for a Main Frame Computer', National Physical Laboratory.
- Davies, D. W. & Clayden, D. O. (1988), 'The Message Authenticator Algorithm (MAA) and its Implementation', National Physical Laboratory.
- Davies, D. W. & Price, W. L. (1980*a*), 'An Annotated Bibliography of Recent Publications on Data Security and Cryptography', National Physical Laboratory.

- Davies, D. W. & Price, W. L. (1980*b*), 'Selected Papers in Cryptography and Data Security', National Physical Laboratory.
- Davies, D. W. & Price, W. L. (1984), *Security for Computer Networks*, John Wiley and Sons, New York.
- de Laet, M. & Mol, A. (2000), 'The Zimbabwe Bush Pump: Mechanics of a Fluid Technology', *Social Studies of Science* **30**(2), 225–263.
- Delanty, G. & Isin, E. F. (2003), Introduction: Reorienting Historical Sociology, in 'Handbook of Historical Sociology', SAGE, London, pp. 1–9.
- Department of Trade and Industry (1976), 'Report of the British Bureau of Standards Working Party', Department of Trade and Industry.
- Department of Trade and Industry (1996*a*), 'Consultation on Strategic Export Controls', Department of Trade and Industry.
- Department of Trade and Industry (1996*b*), 'Paper On Regulatory Intent Concerning Use Of Encryption On Public Networks', Department of Trade and Industry.
- Department of Trade and Industry (1997*a*), 'Licensing Of Trusted Third Parties For The Provision Of Encryption Services: Public Consultation Paper on Detailed Proposals for Legislation', Department of Trade and Industry.
- Department of Trade and Industry (1997*b*), 'Strategic Review of UK Export Licensing Controls'.
URL: <http://webarchive.nationalarchives.gov.uk/+http://www.dti.gov.uk/europeandtrade/strategic-review-of-uk-export-control/licensing-policy/page16442.html>
- Department of Trade and Industry (1998*a*), 'Secure Electronic Commerce Statement', Department of Trade and Industry.
- Department of Trade and Industry (1998*b*), 'White Paper on 'Strategic Export Controls'', Department of Trade and Industry.
- Department of Trade and Industry (1999*a*), 'Building Confidence in Electronic Commerce - A Consultation Document', Department of Trade and Industry.
- Department of Trade and Industry (1999*b*), 'Electronic Commerce Bill', Department of Trade and Industry.

- Department of Trade and Industry (2001), 'Consultation on Draft Legislation: the Export Control and Non-Proliferation Bill', Department of Trade and Industry.
- Diffie, W. & Hellman, M. (1976), 'New Directions in Cryptography', *IEEE Transactions on Information Theory* **22**, 644–654.
- Diffie, W. & Landau, S. (2007), *Privacy on the Line: The Politics of Wiretapping and Encryption (Updated and Expanded Edition)*, The MIT Press, Cambridge.
- Dobson, M. & Ziemann, B. (2009), Introduction, in M. Dobson & B. Ziemann, eds, 'Reading Primary Sources: The Interpretation of Texts from Nineteenth- and Twentieth-Century History', Routledge, Abingdon, pp. 1–18.
- Donaldson, A., Lowe, P. & Ward, N. (2002), 'Virus-Crisis-Institutional Change: The Foot and Mouth Actor Network and the Governance of Rural Affairs in the UK', *Sociologica Ruralis* **42**(3), 201–214.
- Dreyfus, H. L. & Dreyfus, S. E. (1986), *Mind Over Machine: The Power of Human Intuition and Expertise in the Era of the Computer*, The Free Press, New York.
- Duke, K. (2002), 'Getting Beyond the 'Official Line': Reflections on Dilemmas of Access, Knowledge and Power in Researching Policy Networks', *Journal of Social Policy* **31**(1), 39–59.
- Dunleavy, P. & Hood, C. (1994), 'From Old Public Administration to New Public Management', *Public Money & Management* **14**(3), 9–16.
- Easter, D. (2008), 'GCHQ and British External Policy in the 1960s', *Intelligence and National Security* **23**(5), 681–706.
- Edgerton, D. (1993), 'Tilting at Paper Tigers', *The British Journal for the History of Science* **26**(1), 67–75.
- Edgerton, D. (2006), *Warfare State: Britain, 1920-1970*, Cambridge University Press, Cambridge.
- Edgerton, D. (2009), 'The 'Haldane Principle' and Other Invented Traditions in Science Policy', *History & Policy*.
- Edgerton, D. (2011), *Britain's War Machine: Weapons, Resources and Experts in the Second World War*, Allen Lane, London.

- Ellis, J. H. (1970), 'The Possibility of Secure Non-Secret Digital Encryption', Communications-Electronics Security Group.
- Ellis, J. H. (1987), 'The Story of Non-Secret Encryption', Communications-Electronics Security Group.
- Elton, G. R. (1967), *The Practice of History*, Methuen, London.
- Elzen, B. (1986), 'Two Ultracentrifuges: A Comparative Study of the Social Construction of Artifacts', *Social Studies of Science* **16**, 621–662.
- Engelhardt, H. T. & Caplan, A. L. (1987), Patterns of Controversy and Closure: The Interplay of Knowledge, Values, and Political Forces, in H. T. Engelhardt & A. L. Caplan, eds, 'Scientific Controversies: Case Studies in the Resolution and Closure of Disputes in Science and Technology', Cambridge University Press, New York.
- Engineering and Physical Sciences Research Council (2013), 'Grants on the web'.
URL: <http://gow.epsrc.ac.uk>
- Ensmenger, N. L. (2010), *The Computer Boys Take Over: Computers, Programmers, and the Politics of Technical Expertise*, The MIT Press, Cambridge.
- Epstein, S. (1995), 'The Construction of Lay Expertise: AIDS, Activism and the Forging of Credibility in the Reform of Clinical Trials', *Science, Technology and Human Values* **20**(4), 408–437.
- Epstein, S. (1996), *Impure Science: AIDS, Activism, and the Politics of Knowledge*, University of California Press, Berkeley.
- Epstein, S. (2011), 'Misguided Boundary Work in Studies of Expertise: Time to Return to the Evidence', *Critical Policy Studies* **5**(3), 323–328.
- Etzkowitz, H. & Leydesdorff, L. (2000), 'The Dynamics of Innovation: From National Systems and "Mode 2" to a Triple Helix of University–Industry–Government Relations', *Research Policy* **29**(2), 109–123.
- European Commission (1997), 'Towards A European Framework for Digital Signatures And Encryption', European Commission.
- Evans, R. & Collins, H. (2008), Expertise: From Attribute to Attribution and Back Again?, in E. J. Hackett, O. Amsterdamska, M. Lynch & J. Wajcman,

- eds, 'The Handbook of Science and Technology Studies (Third Edition)', The MIT Press, Cambridge, pp. 609–631.
- Evans, R. J. (1999), *In Defence of History*, W. W. Norton, New York.
- Evans, R. & Plows, A. (2007), 'Listening Without Prejudice?: Re-discovering the Value of the Disinterested Citizen', *Social Studies of Science* **37**(6), 827–853.
- Evans, S. (2009), *Technological Ambiguity and the Wassenaar Arrangement*, PhD Thesis, University of Oxford.
- Ezrahi, Y. (1990), *The Descent of Icarus: Science and the Transformation of Contemporary Democracy*, Harvard University Press, Cambridge.
- Fischer, F. (2011), 'The 'Policy Turn' in the Third Wave: Return to the Fact-Value Dichotomy?', *Critical Policy Studies* **5**(3), 311–316.
- Flyvbjerg, B. (2006), 'Five Misunderstandings About Case-Study Research', *Qualitative Inquiry* **12**(2), 219–246.
- Forsyth, T. (2011), 'Expertise Needs Transparency Not Blind Trust: A Deliberative Approach to Integrating Science and Social Participation', *Critical Policy Studies* **5**(3), 317–322.
- Foundation for Information Policy Research (1998a), 'FIPR Launch Press Release'.
URL: <http://www.fipr.org/pr.html>
- Foundation for Information Policy Research (1998b), 'Strategic Export Controls: The Impact on Cryptography', Foundation for Information Policy Research.
- Foundation for Information Policy Research (2003), 'Academic Freedom'.
URL: <http://web.archive.org/web/20030605151219/http://www.fipr.org/academic.html>
- Foundation for Information Policy Research (2005), 'The Crypto Wars Are Over!'.
URL: <http://www.fipr.org/press/050525crypto.html>
- Franken, O. I. (1993), 'Babbage and Cryptography. Or, the Mystery of Admiral Beaufort's Cipher', *Mathematics and Computers in Simulation* **35**, 327–367.
- Frickel, S. (2004), *Chemical Consequences: Environmental Mutagens, Scientist Activism, and the Rise of Genetic Toxicology*, Rutgers University Press, Piscataway.

- Frisch, M. (1998), Oral History and Hard Times: A Review Essay, *in* R. Perks & A. Thompson, eds, 'The Oral History Reader', Routledge, London, pp. 29–37.
- Fujimura, J. (1992), Crafting Science: Standardized Packages, Boundary Objects and Translation, *in* A. Pickering, ed., 'Science as Practice and Culture', University of Chicago Press, Chicago, pp. 168–214.
- Garforth, L. (2012), 'In/Visibilities of Research: Seeing and Knowing in STS', *Science, Technology and Human Values* **37**(2), 264–285.
- Gibbons, M., Limoges, C., Nowotny, H., Schwartzman, S., Scott, P. & Trow, M. (1994), *The New Production of Knowledge: The Dynamics of Science and Research in Contemporary Society*, SAGE, London.
- Giere, R. N. (1987), Controversies Involving Science and Technology: A Theoretical Perspective, *in* H. T. Engelhardt & A. L. Caplan, eds, 'Scientific Controversies: Case Studies in the Resolution and Closure of Disputes in Science and Technology', Cambridge University Press, New York.
- Gilbert, G. N. & Mulkay, M. (1984), *Opening Pandora's Box: A Sociological Analysis of Scientists' Discourse*, Cambridge University Press, Cambridge.
- Gladman, B. (1996a), 'DTI Announces HMS Clipper?', talk.politics.crypto.
- Gladman, B. (1996b), 'UK Cryptography Policy Discussion Group', ukcrypto Mailing List.
- Gladman, B. (1997), 'US government admission', ukcrypto Mailing List.
- Glasziou, P. & Haines, B. (2005), 'The Paths from Research to Improve Health Outcomes', *ACP Journal Club* **142**(2), A8–A10.
- Goldsmith, J. & Wu, T. (2006), *Who Controls the Internet?: Illusions of a Borderless World*, Oxford University Press, Oxford.
- Goldthorpe, J. H. (1991), 'The Uses of History in Sociology: Reflections on Some Recent Tendencies', *British Journal of Sociology* **42**(2), 211–230.
- Golinski, J. (2005), *Making Natural Knowledge: Constructivism and the History of Science*, University of Chicago Press, Chicago.
- Goodwin, M. (2014), 'Political Science? Does Scientific Training Predict UK MPs Voting Behaviour', *Parliamentary Affairs* **Online First**.

- Gummett, P. (1980), *Scientists in Whitehall*, University of Manchester Press, Manchester.
- Hanna, P. (2012), 'Using Internet Technologies (Such as Skype) as a Research Medium: A Research Note', *Qualitative Research* **12**(2), 239–242.
- Hart, N. (1994), 'John Goldthorpe and the Relics of Sociology', *British Journal of Sociology* **45**(1), 21–30.
- Harte, N. (1986), *The University of London 1836-1986: An Illustrated History*, Athlone Press, London.
- HC (1994), 'Deb 21 April 1994, vol 241, col 603'.
- Henare, A., Holbraad, M. & Wastell, S., eds (2007), *Thinking Through Things: Theorising Artefacts Ethnographically*, Routledge, Oxford.
- Henderson, K. (1991), 'Flexible Sketches and Inflexible Data Bases: Visual Communication, Conscripted Devices, and Boundary Objects in Design Engineering', *Science, Technology and Human Values* **16**(4), 448–473.
- Henderson, M. (2011), *The Geek Manifesto: Why Science Matters*, Corgi, London.
- Henderson, S. & Segal, E. H. (2013), Visualizing Qualitative Data in Evaluation Research, in T. Azzam & S. Evergreen, eds, 'Data Visualization, Part 1: New Directions for Evaluation', pp. 53–71.
- Herman, M. (1996), *Intelligence Power in Peace and War*, Cambridge University Press, Cambridge.
- Hess, D. J. (2001), Ethnography and the Development of Science and Technology Studies, in P. Atkinson, A. Coffey, S. Delamont, J. Lofland & L. Lofland, eds, 'SAGE Handbook of Ethnography', SAGE, Thousand Oaks, pp. 234–245.
- Heywood, P. (2010), 'Anthropology and What There Is: Reflections on "Ontology"', *Cambridge Anthropology* **30**(1), 143–151.
- Hill, M. R. (1993), *Archival Strategies and Techniques*, SAGE, Newbury Park.
- Hodges, A. (1983), *Alan Turing: The Enigma*, Burnett Books, London.
- Holt, A. (2010), 'Using the Telephone for Narrative Interviewing: A Research Note', *Qualitative Research* **10**(1), 113–121.

- Hood, C. (1991), 'A Public Management for All Seasons?', *Public Administration* **69**(1), 3–19.
- Hosein, G. (1997), 'Consultation and Contemplation - What Has Gone Before'.
URL: <http://www.fipr.org/publications/consult.html>
- Hosein, I. R. (2003), *Regulating the Technological Actor: How Governments Tried to Transform the Technology and the Market for Cryptography and Cryptographic Services and the Implications for the Regulation of Information and Communications Technologies*, PhD Thesis, London School of Economics and Political Science.
- Hughes, T. P. (1994), Technological Momentum, in M. R. Smith & L. Marx, eds, 'Does Technology Drive History?: The Dilemma of Technological Determinism', The MIT Press, Cambridge, pp. 101–114.
- Hyman, A. (1982), *Charles Babbage: Pioneer of the Computer*, Oxford University Press, Oxford.
- Import, Export and Customs Powers (Defence Act) (1939), The Stationary Office, London.
- Information Security Group (2008), 'The Information Security Group: A Brief History', Royal Holloway, University of London.
- Intel (1997), 'Response to the UK Government Consultation Paper on the Licensing of Trusted Third Parties for the Provision of Encryption Services', Intel Corporation Ltd.
- Interception of Communications Act (1985), The Stationary Office, London.
- Irvine, A. (2011), 'Duration, Dominance and Depth in Telephone and Face-to-Face Interviews: A Comparative Exploration', *International Journal of Qualitative Methods* **10**(3), 202–220.
- Irwin, A. (1995), *Citizen Science: A Study of People, Expertise and Sustainable Development*, Routledge, London.
- Jansen, R. (1990), *The RACAL Handbook: A Review of Racal Communication Equipment, 1956-1975*, G C Arnold, Broadstone.
- Jasanoff, S. (1996), 'Beyond Epistemology: Relativism and Engagement in the Politics of Science', *Social Studies of Science* **26**(3), 393–418.

- Jasanoff, S. (2003), 'Breaking the Waves in Science Studies: Comment on H.M. Collins and Robert Evans, 'The Third Wave of Science Studies'', *Social Studies of Science* **33**(3), 389–400.
- Jasanoff, S. (2005), *Designs on Nature: Science and Democracy in Europe and the United States*, Princeton University Press, Princeton.
- Jefferies, N., Mitchell, C. & Walker, M. (1995), A Proposed Architecture for Trusted Third Party Services, in E. P. Dawson & J. Golic, eds, 'Proceedings of Cryptography Policy and Algorithms Conference, 3-5 July 1995, Brisbane', Springer, New York, pp. 67–81.
- Kahn, D. (1991), *Seizing the Enigma: The Race to Break the German U-Boat Codes, 1939-1943*, Houghton Mifflin, Boston.
- Kahn, D. (1997), *The Codebreakers: The Comprehensive History of Secret Communication From Ancient Times to the Internet (Revised and Updated Edition)*, Scribner's and Sons, New York.
- Kelly, K. (2011), *What Technology Wants*, Penguin, New York.
- Kelty, C. M. (2008), *Two Bits: The Cultural Significance of Free Software*, Duke University Press, Durham.
- Knorr Cetina, K. (1981), *The Manufacture of Knowledge: An Essay on the Constructivist and Contextual Nature of Science*, Pergamon, Oxford.
- Knorr Cetina, K. (1999), *Epistemic Cultures: How the Sciences Made Knowledge*, Harvard University Press, Cambridge.
- Knott, J. & Wildavsky, A. (1980), 'If Dissemination is the Solution, What is the Problem?', *Knowledge: Creation, Diffusion, Utilization* **1**(4), 537–578.
- Koops, B.-J. (1998), *The Crypto Controversy: A Key Conflict in the Information Society*, Kluwer Law International, The Hague.
- Koops, B.-J. (2013), 'Crypto Law Survey'.
URL: <http://www.cryptolaw.org/>
- Kuhn, T. S. (1962), *The Structure of Scientific Revolutions*, University of Chicago Press, Chicago.

- Labour Party (1997), 'New Labour Because Britain Deserves Better', Labour Party.
- Lachmann, R. (2013), *What is Historical Sociology?*, Polity Press, Cambridge.
- Laidlaw, J. (2012), 'Ontologically Challenged', *Anthropology of This Century* 4.
- Landry, R., Amara, N. & Lamari, M. (2001), 'Utilization of Social Science Research Knowledge in Canada', *Research Policy* 30(2), 333–349.
- Larédo, P. & Mustar, P. (2000), 'Laboratory Activity Profiles: An Exploratory Approach', *Scientometrics* 47(3), 515–539.
- Larédo, P., Mustar, P., Callon, M., Birac, A. M. & Fourest, B. (1992), Defining the Strategic Profile of Research Labs: The Research Compass Card Method, in A. F. J. van Raan, R. E. de Bruin, H. F. Moed, A. J. Nederhof & R. W. J. Tijssen, eds, 'Science and Technology in a Policy Context', DSWO Press, Leiden.
- Latour, B. (1987), *Science in Action*, Harvard University Press, Cambridge.
- Latour, B. (1996), *Aramis, or, The Love of Technology*, Harvard University Press, Cambridge.
- Latour, B. (2005), *Reassembling the Social: An Introduction to Actor-Network-Theory*, Oxford University Press, Oxford.
- Latour, B. & Woolgar, S. (1979), *Laboratory Life: The Social Construction of Scientific Facts*, SAGE, Beverly Hills.
- Laudel, G. & Gläser, J. (2007), 'Interviewing Scientists', *Science, Technology and Innovation Studies* 3(2), 91–111.
- Lavington, S. (1980), *Early British Computers: The Story of Vintage Computers and the People Who Built Them.*, Manchester University Press, Manchester.
- Law, J. (1986), On the Methods of Long Distance Control: Vessels, Navigation, and the Portuguese Route to India, in 'Power, Action and Belief: A New Sociology of Knowledge?', Routledge, Henley, pp. 234–263.
- Law, J. (1992), 'Notes on the Theory of the Actor-Network: Ordering, Strategy and Heterogeneity', *Systems Practice* 5(4), 379–393.

- Law, J. & Lien, M. E. (2013), 'Slippery: Field Notes in Empirical Ontology', *Social Studies of Science* **43**(3), 363–378.
- Law, J. & Mol, A. (2011), 'Veterinary Realities: What is Foot and Mouth Disease?', *Sociologica Ruralis* **51**(1), 1–16.
- Law, J. & Singleton, V. (2005), 'Object Lessons', *Organization* **12**(3), 331–355.
- Lessig, L. (1999), *Code and Other Laws of Cyberspace*, Basic Books, New York.
- Levy, S. (2001), *Crypto: Secrecy and Privacy in the New Code War*, Allen Lane, London.
- Lindsey, C. (1997), 'Critique of the DTI Proposals', ukcrypto Mailing List.
- Lynch, M. (1985), *Art and Artifact in Laboratory Science: A Study of Shop Work and Shop Talk in a Research Laboratory*, Routledge & Kegan Paul, London.
- Lynch, M. (2013), 'Ontography: Investigating the Production of Things, Deflating Ontology', *Social Studies of Science* **43**(3), 444–462.
- Lynch, M. & Cole, S. (2005), 'Science and Technology Studies on Trial: Dilemmas of Expertise', *Social Studies of Science* **35**(2), 269–311.
- MacKenzie, D. (1990), *Inventing Accuracy: A Historical Sociology of Nuclear Missile Guidance*, The MIT Press, Cambridge.
- Magnello, E. (2000), *A Century of Measurement: An Illustrated History of the National Physical Laboratory*, Canopus, Bath.
- Mann, M. (1994), 'In Praise of Macro-Sociology: A Reply to Goldthorpe', *British Journal of Sociology* **45**(1), 37–54.
- Marginson, S. (2011), 'Higher Education and Public Good', *Higher Education Quarterly* **65**(4), 411–433.
- Marie, J. (2008), 'For Science, Love and Money: The Social Worlds of Poultry and Rabbit Breeding in Britain, 1900-1940', *Social Studies of Science* **38**(6), 919–936.
- Marres, N. (2013), 'Why Political Ontology Must Be Experimentalized: On Eco-Show Homes as Devices of Participation', *Social Studies of Science* **43**(3), 417–443.

- Meho, L. I. (2006), 'E-mail Interviewing in Qualitative Research: A Methodological Discussion', *Journal of the American Society for Information Science and Technology* **57**(10), 1284–1295.
- Merton, R. K. (1973), The Normative Structure of Science, in 'The Sociology of Science: Theoretical and Empirical Investigations', Chicago University Press, Chicago, pp. 267–281.
- Mirtoff, I. I. (1974), 'Norms and Counter-Norms in a Select Group of the Apollo Moon Scientists: A Case Study of the Ambivalence of Scientists', *American Sociological Review* **39**(4), 579–595.
- Mitchell, C., Murphy, S., Piper, F. & Wild, P. (1996), 'Red Pike - An Assessment', Codes & Ciphers Ltd.
- Mol, A. (1998), Missing Links, Making Links: The Performance of Some Atheroscleroses, in A. Mol & M. Berg, eds, 'Differences in Medicine: Unravelling Practises, Techniques and Bodies', Duke University Press, Durham, pp. 144–165.
- Mol, A. (2002), *The Body Multiple: Ontology in Medical Practice*, Duke University Press, Durham.
- Mol, A. (2013), 'Mind Your Plate! The Ontonorms of Dutch Dieting', *Social Studies of Science* **43**(3), 379–396.
- Mol, A. & Berg, M. (1994), 'Principles and Practices of Medicine: The Co-Existence of Various Anaemias', *Culture, Medicine and Psychiatry* **18**(2), 247–265.
- Mol, A. & Elsmann, B. (1996), 'Detecting Disease and Designing Treatment: Duplex and the Diagnosis of Diseased Leg Vessels', *Sociology of Health and Illness* **18**(5), 609–631.
- Mol, A. & Law, J. (1994), 'Regions, Networks and Fluids: Anaemia and Social Topology', *Social Studies of Science* **24**(4), 641–671.
- Mol, A. & Law, J. (2004), 'Embodied Action, Enacted Bodies: The Example of Hypoglycaemia', *The Body and Society* **10**(2), 43–62.
- Moran, C. (2013), *Classified: Secrecy and the State in Modern Britain*, Cambridge University Press, Cambridge.

- Morison, R. S. (1978), 'Misgivings About Life-Extending Technologies', *Daedalus* **107**(2), 211–227.
- Moseley, R. (1978), 'The Origins and Early Years of the National Physical Laboratory: A Chapter in the Pre-History of British Science Policy', *Minerva* **16**(2), 222–250.
- Mouzelis, N. (1994), 'In Defence of 'Grand' Historical Sociology', *British Journal of Sociology* **45**(1), 31–36.
- Mulkay, M. (1997), *The Embryo Research Debate: Science and the Politics of Reproduction*, Cambridge University Press, New York.
- Murphy, M. (2006), *Sick Building Syndrome and the Problem of Uncertainty: Environmental Politics, Technoscience, and Women Workers*, Duke University Press, Durham.
- Murray, L. (1972), 'Governance of the University of London: Report of the Committee of Enquiry into the Governance of the University of London', University of London.
- National Archives (2012), 'History of the Public Records Acts'.
URL: <https://www.nationalarchives.gov.uk/information-management/legislation/history-of-pra.htm>
- National Computing Centre (1977), National Computing Centre Seminar on the Protection of Data by Cryptography, Cumberland Hotel, London.
- National Physical Laboratory (1971), 'Summary of Comments on the Rothschild Report', National Physical Laboratory.
- National Physical Laboratory (1973), 'Director's Review', National Physical Laboratory.
- National Physical Laboratory (1974), 'Minutes of Meeting of the National Physical Laboratory Review Committee', National Physical Laboratory.
- National Physical Laboratory (1978a), 'First Thoughts on Strategy', National Physical Laboratory.
- National Physical Laboratory (1978b), 'Requirements Boards Customer View and Future Trends', National Physical Laboratory.

- Needham, R. (2003), 'The Clifford Paterson Lecture, 2002: Computer Security?', *Philosophical Transactions of the Royal Society A* **361**, 1549–1555.
- Needham, R. M. (1992), 'Later Developments at Cambridge: Titan, CAP, and the Cambridge Ring', *IEEE Annals of the History of Computing* **14**(4), 57–58.
- Needham, R. M. & Schroeder, M. D. (1978), 'Using Encryption for Authentication in Large Networks of Computers', *Communications of the ACM* **21**(12), 993–999.
- Nelkin, D. (1971), *Nuclear Power and Its Critics: The Cayuga Lake Controversy*, Cornell University Press, Ithaca.
- Nelkin, D. (1979), *Controversy: Politics of Technical Decisions*, SAGE, Beverly Hills.
- Nelkin, D. (1984), *Controversy: Politics of Technical Decisions (Second Edition)*, SAGE, Beverly Hills.
- Nelkin, D. (1993a), Science, Technology, and Political Conflict: Analyzing the Issues, in D. Nelkin, ed., 'Controversy: Politics of Technical Decisions (Third Edition)', SAGE, Beverly Hills.
- Nelkin, D. (1995), Science Controversies: The Dynamics of Public Disputes in the United States, in S. Jasanoff, G. E. Markle, J. C. Peterson & T. Pinch, eds, 'Handbook of Science and Technology Studies', SAGE, Thousand Oaks, pp. 444–456.
- Nelkin, D., ed. (1993b), *Controversy: Politics of Technical Decisions (Third Edition)*, SAGE, Beverly Hills.
- Nelkin, D. & Jasper, J. M. (1993), The Animal Rights Controversy, in D. Nelkin, ed., 'Controversy: Politics of Technical Decisions (Third Edition)', SAGE, Beverly Hills.
- Noble, D. F. (1977), *America by Design: Science, Technology and the Rise of Corporate Capitalism*, Oxford University Press, Oxford.
- Noble, D. F. (1984), *Forces of Production: A Social History of Industrial Automation*, Alfred A. Knopf, New York.
- Nutley, S. M., Walter, I. & Davies, H. T. O. (2007), *Using Evidence: How Research Can Inform Public Services*, Policy Press, Bristol.

- Organisation for Economic Co-operation and Development (1997), 'OECD Guidelines for Cryptography Policy', Organisation for Economic Co-operation and Development.
- Owens, S. (2011), 'Three Thoughts on the Third Wave', *Critical Policy Studies* 5(3), 329–333.
- Parviainen, S.-P. (2000), Cryptographic Software Export Controls in the EU, PhD Thesis, University of Helsinki.
- Performance and Innovation Unit (1999), 'Encryption and Law Enforcement', Performance and Innovation Unit.
- Perlroth, N., Larson, J. & Shane, S. (2013), 'N.S.A. Able to Foil Basic Safeguards of Privacy on Web'.
URL: <http://www.nytimes.com/2013/09/06/us/nsa-foils-much-internet-encryption.html>
- Pettigrew, A. (1985), *The Awakening Giant: Continuity and Change in Imperial Chemical Industries*, Blackwell, Oxford.
- Pickering, A. (1981), 'Constraints on Controversy: The Case of the Magnetic Monopole', *Social Studies of Science* 11(1), 63–93.
- Pickering, A. (1984), *Constructing Quarks: A Sociological History of Particle Physics*, University of Chicago Press, Chicago.
- Pinch, T. J. & Bijker, W. E. (1989), The Social Construction of Facts and Artifacts: Or How the Sociology of Science and the Sociology of Technology Might Benefit Each Other, in W. E. Bijker, T. P. Hughes & T. J. Pinch, eds, 'The Social Construction of Technological Systems: New Directions in the Sociology and History of Technology', The MIT Press, Cambridge, pp. 17–50.
- Piper, F. & Murphy, S. (2002), *Cryptography: A Very Short Introduction*, Oxford University Press, Oxford.
- Polanyi, M. (1958), *Personal Knowledge: Towards a Post-Critical Philosophy*, University of Chicago Press, Chicago.
- Polanyi, M. (1962), 'The Republic of Science: Its Political and Economic Theory', *Minerva* 1(1), 54–74.

- Porter, B. (2009), 'Other People's Mail', *London Review of Books* **31**(22), 15–17.
- Preneel, B., Rumen, V. & van Oorschot, P. C. (1997), 'Security Analysis of the Message Authenticator Algorithm (MAA)', *European Transaction on Telecommunications* **8**(5), 455–470.
- Price, W. L. (1979), 'A Further Annotated Bibliography of Cryptography as Applied to Data Protection in Computing', National Physical Laboratory.
- Price, W. L. (1980), 'A Fourth Annotated Bibliography of Recent Publications on Data Security and Cryptography', National Physical Laboratory.
- Price, W. L. (1982), 'A Fifth Annotated Bibliography of Recent Publications on Data Security and Cryptography', National Physical Laboratory.
- Price, W. L. (1983), 'A Sixth Annotated Bibliography of Recent Publications on Data Security and Cryptography', National Physical Laboratory.
- Pyatt, E. (1983), *The National Physical Laboratory: A History*, Adam Hilger, Bristol.
- Racal (1975), 'The Communications Specialists', *Racal Review* **4**(10), 3–9.
- Racal (1977), 'Security is the Key', *Racal Review* **5**(3), 12.
- Rappert, B. (2009), *Experimental Secrets: International Security, Codes, and the Future of Research*, University Press of America, Lanham.
- Rappert, B. (2010), Making Silence Matter: The Place of the Absences in Ethnography, in 'Ethnography Practice in Industry Conference Proceedings', American Anthropological Association, pp. 260–273.
- Rashid, R. (2004), Introduction: Roger Needham, in A. Herbert & K. Spärck Jones, eds, 'Computer Systems: Theory Technology and Applications', Springer, New York, pp. 1–8.
- Rip, A. (2003), 'Constructing Expertise: In a Third Wave of Science Studies?', *Social Studies of Science* **33**(3), 419–434.
- Rivest, R., Shamir, A. & Adleman, L. (1978), 'A Method for Obtaining Digital Signatures and Public-Key Cryptosystems', *Communications of the ACM* **22**(2), 120–126.

- Robbins, C. (1963), *Higher Education: Report of the Committee Appointed by the Prime Minister Under the Chairmanship of Lord Robbins*, The Stationary Office, London.
- Rogers, A. (1997), *Secrecy and Power in the British State: A History of the Official Secrets Act*, Pluto, London.
- Rogers, R. & Walters, R. (2006), *How Parliament Works (Sixth Edition)*, Routledge, London.
- Rothschild, V. (1971), *A Framework for Government Research and Development*, The Stationary Office, London.
- Royal Anniversary Trust (2013), 'Previous Prize-winners'.
URL: <http://www.royalanniversarytrust.org.uk/the-prizes/previous-prize-winners>
- Royal Holloway (1978), 'Minutes of the Meeting of the Academic Board', Royal Holloway, University of London.
- Royal Holloway (1984), 'Minutes of the Meeting of the Academic Board', Royal Holloway, University of London.
- Royal Holloway (1986a), 'Minutes of the Meeting of the Academic Board', Royal Holloway, University of London.
- Royal Holloway (1986b), 'Minutes of the Meeting of the Academic Board', Royal Holloway, University of London.
- Royal Holloway (1986c), 'Minutes of the Meeting of the Faculty of Science', Royal Holloway, University of London.
- Royal Holloway (1991), 'Proposal Document - M.Sc. in Information Security', Royal Holloway, University of London.
- Rubin, H. S. & Rubin, I. S. (1995), *Qualitative Interviewing: The Art of Hearing Data*, SAGE, London.
- Rudner, M. (2004), 'Britain Betwixt and Between: UK SIGINT Alliance Strategy's Transatlantic and European Connections', *Intelligence and National Security* **19**(4), 571–609.

- Santos, B. d. S. (2004), Toward a Counter Hegemonic Globalization, *in* J. Sen, A. Anand, A. Escobar & P. Waterman, eds, 'World Social Forum: Challenging Empires', The Viveka Foundation, New Delhi, pp. 235–245.
- Saunders, A. (1997), 'CESG Recommendations for Secure Electronic Mail', Communications-Electronics Security Group.
- Schaffer, S. (1994), 'Babbage's Intelligence: Calculating Engines and the Factory System', *Critical Inquiry* **21**, 203–227.
- Schneier, B. (1996), *Applied Cryptography: Protocols, Algorithms and Source Code in C (Second Edition)*, John Wiley and Sons, New York.
- Schneier, B. & Banisar, D., eds (1997), *The Electronic Privacy Papers: Documents on the Battle for Privacy in the Age of Surveillance*, John Wiley and Sons, New York.
- Schrader, A. (2010), 'Responding to Pfiesteria Piscicida (the Fish Killer): Phantomatic Ontologies, Indeterminacy, and Responsibility in Toxic Microbiology', *Social Studies of Science* **40**(2), 275–306.
- Scott, J. (1990), *A Matter of Record: Documentary Sources in Social Research*, Polity, Cambridge.
- Scott, R. (1996), *Report of the Inquiry into the Export of Defence Equipment and Dual-Use Goods to Iraq and Related Prosecutions*, The Stationary Office, London.
- Sebag-Montefiore, S. (2001), *Enigma: The Battle for the Code*, Orion, London.
- Select Committee (1998), 'Second Report: Strategic Export Controls', Commons Select Committee on Trade and Industry.
- Select Committee (1999), '"Building confidence in Electronic Commerce": The Government's Proposals', Select Committee on Trade and Industry.
- Shapin, S. (1975), 'Phrenological Knowledge and the Social Structure of Early Nineteenth-Century Edinburgh', *Annals of Science* **32**(3), 219–243.
- Shapin, S. & Schaffer, S. (1985), *Leviathan and the Air Pump: Hobbes, Boyle, and the Experimental Life*, Princeton University Press, Princeton.

- Shepherd, S. J. (1997), 'Responses to Ross's Paper on the GCHQ Protocol', ukcrypto Mailing List.
- Shopes, L. (2011), Oral History, in N. K. Denzin & Y. S. Lincoln, eds, 'The SAGE Handbook of Qualitative Research', SAGE, Thousand Oaks, pp. 451–465.
- Singh, S. (1999), *The Code Book: The Secret History of Codes and Code-Breaking*, Random House, New York.
- Slaughter, S. & Leslie, L. L. (1997), *Academic Capitalism: Politics, Policies, and the Entrepreneurial University*, Johns Hopkins University Press, Baltimore.
- Slaughter, S. & Rhoades, G. (2004), *Academic Capitalism and the New Economy: Markets, State and Higher Education*, Johns Hopkins University Press, Baltimore.
- Spärck Jones, K. (1999), 'A Brief Informal History of the Computer Laboratory', University of Cambridge.
- Stake, R. E. (1995), *The Art of Case Study Research*, SAGE, Thousand Oaks.
- Staples, W. G. (1998), 'The Electronic Keyhole', *American Scientist* **86**(5), 487.
- 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.
- Stephens, N. (2007), 'Collecting Data from Elites and Ultra Elites: Telephone and Face-to-Face Interviews with Macroeconomists', *Qualitative Research* **7**(2), 203–216.
- Suryanarayanan, S. & Kleinman, D. L. (2013), 'Be(e)coming Experts: The Controversy Over Insecticides in the Honey Bee Colony Collapse Disorder', *Social Studies of Science* **43**(2), 215–240.
- Thompson, C. (2007), *Making Parents: The Ontological Choreography of Reproductive Technologies*, The MIT Press, Cambridge.
- Thompson, F. M. L. (1990), Introduction, in 'The University of London and the World of Learning, 1836-1986', Hambledon Press, London, pp. ix–xxii.
- UK Parliament (2013), 'Passage of a Bill'.
URL: <http://www.parliament.uk/about/how/laws/passage-bill/>

- van Heur, B., Leydesdorff, L. & Wyatt, S. (2013), 'Turning to Ontology in STS? Turning to STS Through 'Ontology'', *Social Studies of Science* **43**(3), 341–362.
- van Rooij, A. (2011), 'Knowledge, Money and Data: An Integrated Account of the Evolution of Eight Types of Laboratory', *The British Journal for the History of Science* **44**(3), 427–448.
- van Rooij, A. (2013), 'Gaps and Plugs: TNO, and the Problems of Getting Knowledge Out of Laboratories', *Minerva Online First*.
- Vaughan, D. (1996), *The Challenger Launch Decision: Risky Technology, Culture, and Deviance at NASA*, University of Chicago Press, Chicago.
- Vincent, D. (1997), *The Culture of Secrecy: Britain 1832-1998*, Oxford University Press, Oxford.
- Wallard, A. (2001), Successful Contractorisation - The Experience of the National Physical Laboratory, in D. Cox, P. Gummett & K. Barker, eds, 'Government Laboratories: Transition and Transformation', IOS Press, Amsterdam, pp. 202–209.
- Wassenaar Arrangement (2013), 'Introduction'.
URL: <http://www.wassenaar.org/introduction/index.html>
- Webber, D. J. (1986), 'Explaining Policymakers' Use of Policy Information', *Knowledge: Creation, Diffusion, Utilization* **7**(3), 249–290.
- Weinel, M. (2010), Technological Decision-Making Under Scientific Uncertainty: Preventing Mother-to-Child Transmission of HIV in South Africa, PhD Thesis, Cardiff University.
- Wheeler, D. & Needham, R. (1994), 'TEA, a Tiny Encryption Algorithm', *Lecture Notes in Computer Science* **1008**, 363–366.
- Wheeler, J. M. (1992), 'Applications of the EDSAC', *IEEE Annals of the History of Computing* **14**(4), 27–33.
- Whelan, R. C. (2000), 'Management of Scientific Institutions NPL 1995-98: The Transition from Agency to Government-Owned Contractor Operated (GOCO)', *R&D Management* **30**(4), 312–322.
- Wilkes, M. V. (1985), *Memoirs of a Computer Pioneer*, The MIT Press, Cambridge.

- Wilkes, M. V. (1992), 'EDSAC 2', *IEEE Annals of the History of Computing* **14**(4), 49–56.
- Wilkie, T. (1991), *British Science and Politics Since 1945*, Blackwell, Oxford.
- Williamson, M. J. (1974), 'Non-Secret Encryption Using a Finite Field', Communications-Electronics Security Group.
- Williamson, M. J. (1976), 'Thoughts on Cheaper Non-Secret Encryption', Communications-Electronics Security Group.
- Wilsdon, J. & Willis, R. (2004), 'See-Through Science: Why Public Engagement Needs to Move Upstream', Demos.
- Wilson, J. P. (1980), 'Growing with Racal', *Radio and Electronic Engineer* **50**(10), 499–502.
- Winner, L. (1977), *Autonomous Technology: Technics-out-of-Control as a Theme in Political Thought*, The MIT Press, Cambridge.
- Winner, L. (1980), 'Do Artifacts Have Politics?', *Daedalus* **109**(1), 121–136.
- Wittgenstein, L. (1953), *Philosophical Investigations*, Macmillan, New York.
- Woods, A. (2004), *A Manufactured Plague: The History of Foot and Mouth Disease in Britain*, Earthscan, London.
- Woolgar, S. (1991), Configuring the User: The Case of Usability Trials, in J. Law, ed., 'A Sociology of Monsters: Essays on Power, Technology and Domination', Routledge, London.
- Woolgar, S. & Lezaun, J. (2013), 'The Wrong Bin Bag: A Turn to Ontology in Science and Technology Studies?', *Social Studies of Science* **43**(3), 321–340.
- Wray, K. B. (2005), 'Rethinking Scientific Specialization', *Social Studies of Science* **35**(1), 151–164.
- Wyatt, S. (2003), Non-Users Also Matter: The Construction of Users and Non-Users of the Internet, in N. Oudshoorn & T. Pinch, eds, 'How Users Matter: The Co-Construction of Users and Technology', The MIT Press, Cambridge, pp. 67–79.

- Wynne, B. (1982), *Rationality and Ritual: The Windscale Inquiry and Nuclear Decisions in Britain*, British Society for the History of Science, Oxford.
- Wynne, B. (1992), 'Misunderstood Misunderstanding: Social Identities and Public Uptake of Science', *Public Understanding of Science* **1**(3), 281–304.
- Wynne, B. (1996), May the Sheep Safely Graze? A Reflexive View of the Expert-Lay Knowledge Divide, in S. Lash, B. Szerszynski & B. Wynne, eds, 'Risk, Environment & Modernity: Towards a New Ecology', SAGE, London, pp. 44–83.
- Wynne, B. (2003), 'Seasick on the Third Wave? Subverting the Hegemony of Propositionalism: Response to Collins & Evans (2002)', *Social Studies of Science* **33**(3), 401–417.
- Wynne, B. (2007), 'Public Participation in Science and Technology: Performing and Obscuring a Political-Conceptual Category Mistake', *East Asian Science, Technology and Society* **1**, 99–110.
- Yates, D. M. (1997), *Turing's Legacy: A History of Computing at the National Physical Laboratory 1945-1995*, Science Museum, London.
- Yearley, S. (2005), *Making Sense of Science: Understanding the Social Study of Science*, SAGE, London.
- Yin, R. K. (2009), *Case Study Research: Design and Methods (Fourth Edition)*, SAGE, Thousand Oaks.
- Zergo (1996), 'The Use of Encryption and Related Services with the NHSnet: A Report for the NHS Executive by Zergo Limited', NHS Executive.