



**UCL**

**Measuring and Understanding  
Security Behaviours**

Ingolf Becker

A thesis submitted for the degree of

*Doctor of Philosophy*

at

UCL

2019



I, Ingolf Becker confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the thesis.



## Abstract

Information security embodies the complex interaction between security policies, user perceptions of these policies, productive activity and the security culture in general. The vast majority of organisations consist not solely of data and technology, but have human actors involved in the productive activity, and are thus socio-technical systems. The aim of this thesis is to understand how individuals perceive, understand and react to information security policies, and how they fit into productive tasks, while investigating the viability of measuring each of these aspects.

An analytical evaluation and empirical user study in three countries of banking policies evidences difficulties in understanding policies. A second study quantifies actual user characteristics and shows that the assumptions on user behaviour in the policies are unrealistic. Advice attempting to explain security aspects to the general public fail to improve user understanding, and security awareness is promoted without measuring the impact of the interventions. Better understanding and measurements of security culture are needed.

This demand is pursued in the remainder of the thesis: in two companies, the results of context-aware surveys that elicit responses to typical scenarios of non-compliant behaviours are evaluated. The responses are used to define the security culture of the company, and to re-frame the notion of Security Champions based on the observed security cultures.

Finally, the impact of a change in password policy in a university with over 100,000 users for 17 months is studied. Virtually all users respond positively to the policy change, adopting a more secure password over time in response to a longer password lifetime. This work gives evidence for the benefit of involving users in security decisions.

The metrics developed in this thesis allow security to be grounded in the actual circumstances of the organisation and its human actors and security to be evaluated objectively. By involving and empowering individuals, security can become workable and sustainable.



## Impact Statement

My academic publications are actively being disseminated, with a total of 34 citations to date. Due to the Productive Security project, we have received two small grants for further work, both part of this thesis: the security construct work ([Chapter 3](#)) and the password project ([Chapter 7](#)). Wherever possible, the research in this thesis has been conducted in a reproducible manner. I have released datasets and analysis code under a permissive license, allowing future researchers and practitioners to easily understand and build upon my results.

Based on our research on banking terms and conditions we submitted evidence to the European Banking authority’s consultation on strong customer authentication in the Payment Service Directive 2.

Organisations can use the methodologies presented in this thesis to measure their security culture. This enables them to tailor interventions to their organisation, and measure the impact of interventions. At Companies A & B we presented the findings to their respective boards, and we are continuing to work with UCL to integrate our feedback into their password policy. To disseminate my research to the wider public, a number of blog posts have been released our Information Security Group’s Blog<sup>1</sup> and the blog of the Research Institute in Science of Cyber Security<sup>2</sup>, each with over 2500 unique readers per month.

---

<sup>1</sup>Bentham’s Gaze

<sup>2</sup>RISCS





## Acknowledgements

I would like to thank my supervisors, Angela Sasse and Sebastian Riedel, for their support and advice throughout my PhD. Their enthusiasm and drive was an essential motivation for this thesis and will continue to inspire me for years to come. In addition, Simon Parkin, my frequent collaborator, was invaluable as a constant source of feedback, inspiration and complicated holistic analyses that put my research into perspective. I was supported by an EPSRC grant through the Doctoral Training Centre in Security and Crime Science. I would like to thank the staff in the Department of Security and Crime Science as well as those in the Department of Computer Science for their kind assistance in all administrative matters. I am also grateful for the anonymous companies A, B and C who generously donated their employees' time to this research.

I enjoyed my time as a research student at UCL, and this is primarily thanks to a great research group, and the many current and past students and post-docs of 6.07 and 6.22. This environment has led to many interesting and fruitful discussions and collaborations.

I am also grateful for the comments from my examiners, Emiliano De Cristofaro and Kami Vaniea. I would also like to thank my co-authors, especially those of the papers on which this thesis is based: Ruba Abu-Salma, Ross Anderson, Adam Beutement, Nicholas Bohm, Albesë Demjaha, Alice Hutchings, Kat Krol, Steven Murdoch, Simon Parkin, Angela Sasse, Jono Spring and Gianluca Stringhini.

I appreciate the support my friends, new and old, and family have given me throughout this journey, and especially my parents who have provided nothing but useful advice and guidance. Last but not least, I thank Louise for her love and encouragement, and for sharing this journey with me.



# Contents

<b>Abstract</b>	<b>iii</b>
<b>Impact Statement</b>	<b>v</b>
<b>Acknowledgements</b>	<b>vii</b>
<b>Contents</b>	<b>ix</b>
<b>List of Figures</b>	<b>xv</b>
<b>List of Tables</b>	<b>xvii</b>
<b>1 Introduction</b>	<b>1</b>
1.1 Security scope . . . . .	2
1.2 Security manifestation . . . . .	3
1.3 Security awareness . . . . .	3
1.4 Security decision making . . . . .	4
1.5 Research questions . . . . .	4
1.6 Overview of thesis . . . . .	5
1.6.1 Motivational research . . . . .	6
1.6.2 Productive Security . . . . .	9
1.6.3 Password project . . . . .	12
1.6.4 Other research conducted during my PhD . . . . .	14
1.7 My original contributions of this thesis . . . . .	14
<b>2 Literature review</b>	<b>17</b>
2.1 Measuring security . . . . .	17
2.2 Factors influencing compliance . . . . .	20
2.3 Perceived security fatigue . . . . .	25
2.3.1 Sources of fatigue . . . . .	26
2.3.2 Responses to fatigue . . . . .	26

2.3.3	Types of compliance . . . . .	26
2.4	Coping strategies and workarounds . . . . .	27
<b>3</b>	<b>Security constructs survey</b>	<b>29</b>
3.1	Introduction . . . . .	30
3.2	Methodology . . . . .	31
3.2.1	Construct validation . . . . .	33
3.2.2	Construct grouping . . . . .	34
3.3	The database . . . . .	35
3.3.1	Publication pages . . . . .	35
3.3.2	Construct pages . . . . .	38
3.3.3	Category pages . . . . .	40
3.4	Analysis . . . . .	41
3.5	A case study of construct origins . . . . .	43
3.6	Limitations . . . . .	46
3.7	Conclusion . . . . .	47
3.7.1	Recommendations for researchers and practitioners . . . . .	48
<b>4</b>	<b>Bank Terms and Conditions: perceptions and reality</b>	<b>49</b>
4.1	Introduction . . . . .	49
4.1.1	My contribution in this chapter . . . . .	51
4.2	Related literature . . . . .	52
4.2.1	Legal and regulatory context . . . . .	54
4.3	Review of banking Terms and Conditions internationally . . . . .	55
4.3.1	Methodology . . . . .	55
4.3.2	Results . . . . .	58
4.3.3	Discussion . . . . .	68
4.3.4	Limitations . . . . .	69
4.4	Survey of payment card PIN usage . . . . .	69
4.4.1	Questionnaire setup . . . . .	70
4.4.2	Results . . . . .	70
4.4.3	Discussion . . . . .	73
4.5	Survey of understanding and interpretation of banking T&Cs . . . . .	74
4.5.1	Related literature . . . . .	75
4.5.2	Survey design . . . . .	76
4.5.3	The scenarios . . . . .	77
4.5.4	The Terms and Conditions . . . . .	79
4.5.5	Demographics . . . . .	79

---

4.5.6	Payment demographics . . . . .	81
4.5.7	Scenario overview . . . . .	83
4.5.8	Scenario 1: card loss . . . . .	84
4.5.9	Scenario 2: phone scam . . . . .	87
4.5.10	Understanding of Terms and Conditions . . . . .	90
4.6	Discussion . . . . .	93
4.6.1	Data availability . . . . .	95
<b>5</b>	<b>Studies on perceptions of security</b>	<b>97</b>
5.1	Security metaphors . . . . .	97
5.1.1	Metaphors considered . . . . .	99
5.1.2	Survey testing of metaphors . . . . .	100
5.1.3	Reliability and validity . . . . .	101
5.1.4	Findings . . . . .	102
5.1.5	Discussion . . . . .	104
5.2	SANS analysis . . . . .	106
5.3	Conclusion . . . . .	108
<b>6</b>	<b>Productive Security</b>	<b>109</b>
6.1	Introduction . . . . .	110
6.2	Related literature . . . . .	111
6.3	Methodology . . . . .	114
6.4	Surveys . . . . .	116
6.4.1	Online scenario-based survey . . . . .	116
6.4.2	Scenario commonality . . . . .	117
6.4.3	Survey design decisions . . . . .	118
6.4.4	Company A . . . . .	119
6.4.5	Company B . . . . .	120
6.4.6	Survey tasks scoring . . . . .	123
6.4.7	Selecting interventions . . . . .	123
6.4.8	Research ethics and data handling . . . . .	124
6.5	Survey results . . . . .	124
6.5.1	Company A . . . . .	125
6.5.2	Company B . . . . .	128
6.6	Discussion of survey results . . . . .	134
6.6.1	Company A . . . . .	135
6.6.2	Company B . . . . .	135
6.6.3	Comparisons between Companies A and B . . . . .	136

6.6.4	Limitations . . . . .	136
6.7	Survey conclusion . . . . .	137
6.8	Security Champions . . . . .	137
6.8.1	Free-text survey responses . . . . .	139
6.8.2	Source data . . . . .	139
6.9	Culture analysis . . . . .	140
6.9.1	Sales & Services division . . . . .	142
6.9.2	Operations division . . . . .	144
6.9.3	Business division . . . . .	145
6.9.4	Finance & Professional Services division . . . . .	147
6.9.5	Discussion and limitations . . . . .	148
6.9.6	Recommendations for practitioners and researchers . . . . .	150
6.9.7	Conclusion . . . . .	150
6.10	Validation of Surveys . . . . .	151
6.10.1	Employee types . . . . .	152
6.10.2	Survey design . . . . .	154
6.10.3	Methodology . . . . .	154
6.10.4	Results . . . . .	157
6.10.5	Conclusions . . . . .	161
6.11	Discussion . . . . .	161
6.12	Conclusion . . . . .	162
<b>7</b>	<b>Password Project</b>	<b>165</b>
7.1	Introduction . . . . .	166
7.2	Related literature . . . . .	167
7.2.1	Password strength estimation . . . . .	168
7.2.2	The role of users in password security . . . . .	169
7.2.3	Studying passwords in the wild . . . . .	169
7.2.4	Password policy . . . . .	171
7.3	Methodology . . . . .	171
7.3.1	The interface . . . . .	171
7.3.2	Empirical evaluation . . . . .	174
7.3.3	The dataset . . . . .	175
7.3.4	Calculation of entropy . . . . .	176
7.3.5	Uses of a password . . . . .	176
7.3.6	UCL's threat model . . . . .	177
7.3.7	Perceived value of a password . . . . .	177
7.3.8	User interviews . . . . .	178

---

7.4	Results . . . . .	178
7.4.1	Noteworthy events during the study . . . . .	180
7.4.2	Password change behaviour . . . . .	181
7.4.3	Time dependence of subsequent changes/resets on prior life-times . . . . .	184
7.4.4	Password change time series . . . . .	186
7.4.5	Password change time series by school . . . . .	188
7.4.6	Password change time series by relationship . . . . .	189
7.4.7	User feedback . . . . .	190
7.5	Discussion . . . . .	192
7.5.1	Limitations . . . . .	194
7.6	Conclusion . . . . .	194
7.6.1	Policy interventions . . . . .	195
<b>8</b>	<b>Conclusions</b>	<b>197</b>
8.1	Recommendations for researchers . . . . .	200
8.2	Recommendations for practitioners . . . . .	201
8.3	Limitations . . . . .	201
8.4	Future work . . . . .	202
	<b>Bibliography</b>	<b>203</b>
<b>A</b>	<b>Productive Security</b>	<b>225</b>
A.1	Company A additional tables . . . . .	226
A.2	Company B additional tables . . . . .	228
A.3	Company A Behaviour and Attitude scenarios . . . . .	230
A.3.1	Scenario A (Behaviour): File Sharing . . . . .	230
A.3.2	Scenario B (Behaviour): Managing Permissions . . . . .	231
A.3.3	Scenario C (Behaviour): USB Stick Usage . . . . .	232
A.3.4	Scenario D (Attitude): Tailgating . . . . .	232
A.3.5	Scenario E (Attitude): Document Control . . . . .	233
A.3.6	Scenario F (Attitude): Information Disposal . . . . .	234
A.3.7	Scenario G (Behaviour): Backing Up Information . . . . .	235
A.3.8	Scenario H (Attitude): External Threats . . . . .	235
A.3.9	Scenario I (Behaviour): Information Requests . . . . .	236
A.3.10	Scenario J (Attitude): Working Practises . . . . .	237
A.4	Company B Behaviour and Attitude scenarios . . . . .	238
A.4.1	Scenario A (Attitude): ID Badges . . . . .	238
A.4.2	Scenario B (Attitude): Clear Desk Policy . . . . .	238

A.4.3	Scenario C (Behaviour): Password Manager . . . . .	239
A.4.4	Scenario D (Behaviour): VPN . . . . .	239
A.4.5	Scenario E (Attitude): Tailgating . . . . .	240
A.4.6	Scenario F (Behaviour): File Storage . . . . .	240
A.4.7	Scenario G (Attitude): Secure Disposal . . . . .	241
A.4.8	Scenario H (Behaviour): Credit Check . . . . .	241
A.5	Maturity model . . . . .	242
<b>B</b>	<b>SANS analysis</b>	<b>245</b>
B.1	Analysis, 2016 . . . . .	245
B.1.1	Conclusions . . . . .	251
B.2	Analysis, 2017 . . . . .	253
B.2.1	Memorable quotes . . . . .	254
B.2.2	Conclusions . . . . .	256



# List of Figures

3.1	The home page of the constructDB . . . . .	35
3.2	ConstructDB Screenshot for Sohrabi Safa, Von Solms and Furnell (2016)	36
3.3	ConstructDB Screenshot for Ifinedo (2014) . . . . .	37
3.4	ConstructDB Screenshot for construct <u>Attitude towards compliance</u> .	38
3.5	ConstructDB Screenshot for category <i>Social bond</i> . . . . .	40
4.1	Histogram of our participants' age . . . . .	80
4.2	Participants' number of payment cards . . . . .	81
4.3	Participants' number of bank accounts . . . . .	82
6.1	Overview of processes in our methodology . . . . .	114
6.2	Kiviat diagram for the entire organisation . . . . .	140
6.3	Kiviat diagram for the the Sales & Service division . . . . .	143
6.4	Kiviat diagram for the Operations division . . . . .	144
6.5	Kiviat diagram for the Business division . . . . .	146
6.6	Kiviat diagram for the Finance & Prof. Services division . . . . .	147
7.1	Screenshot of the myaccount home page . . . . .	172
7.2	Screenshot of the login page . . . . .	172
7.3	Screenshot of the change password website . . . . .	173
7.4	Screenshot of the register your phone website . . . . .	174
7.5	Screenshot of the reset your password website . . . . .	174
7.6	Empirical analysis of the password strength validation function. . . .	174
7.7	Normalised frequency of password lifetime . . . . .	179
7.8	Number of password changes and resets over time . . . . .	180
7.9	Distribution of the number of changes and resets the users in the dataset have made. . . . .	181
7.10	Average password lifetime of unexpired passwords by number of password resets. . . . .	182

LIST OF FIGURES

---

7.11	Distribution of the change in the password lifetime after the password change/reset . . . . .	183
7.12	The frequency of password changes by the number of days relative to password expiration (day 0). . . . .	184
7.13	The distribution of the time between consecutive password changes. . . . .	185
7.14	31-day moving average of the mean password strength of all users and new users. . . . .	186
7.15	31-day moving average password expiration for selected schools over time. . . . .	188
7.16	31-day moving average password expiration for various relationships with the university over time. . . . .	189

## List of Tables

3.1	Distribution of construct questions given by type of the publication. .	41
3.2	Distribution of types of content validation of a construct by type of the publication. . . . .	42
3.3	Distribution of types of piloting or pretesting for a construct by type of the publication. . . . .	42
3.4	Distribution of answer option types for constructs . . . . .	43
4.1	Banking documents included in survey. . . . .	58
4.2	Description of coding categories used. . . . .	59
4.3	Summary of banks' T&Cs related to PIN security . . . . .	60
4.4	Distribution of the number of 4-, 5- and 6-digit PINs the participants have. . . . .	70
4.5	Distribution of how frequently our participants use each of their PINs.	71
4.6	Source of participants' 4-digit PINs . . . . .	71
4.7	Location of written down PINs by participants . . . . .	72
4.8	A variety of locations where participants' PINs are re-used . . . . .	72
4.9	Sharing of PINs by participants . . . . .	73
4.10	Gender of our participants . . . . .	79
4.11	Employment demographics of our participants . . . . .	80
4.12	Educational demographics of our participants . . . . .	81
4.13	Frequency of use of any of our participants' payment cards . . . . .	82
4.14	Frequency of Fraud . . . . .	82
4.15	Thematic analysis of the description of fraud experienced by participants	83
4.16	Percentage of participants that say that the money should be returned in each of the scenarios . . . . .	84
4.17	Thematic analysis of the answers in support of reimbursement in Scenario 1: Card Loss. . . . .	84
4.18	Thematic analysis of the answers not in support of reimbursement in Scenario 1: Card Loss. . . . .	85

LIST OF TABLES

---

4.19	Thematic analysis of the answers in support of reimbursement in Scenario 1: Card Loss, after the participants have seen the T&Cs. . . . .	86
4.20	Thematic analysis of the answers not in support of reimbursement in Scenario 1: Card Loss, after the participants have seen the T&Cs. . .	86
4.21	Thematic analysis of the answers in support of reimbursement in Scenario 2: Phone Scam. . . . .	88
4.22	Thematic analysis of the answers not in support of reimbursement in Scenario 2: Phone Scam. . . . .	88
4.23	Thematic analysis of the answers in support of reimbursement in Scenario 2: Phone Scam, after the participants have seen the T&Cs. . . . .	89
4.24	Thematic analysis of the answers not in support of reimbursement in Scenario 2: Phone Scam, after the participants have seen the T&Cs. .	89
4.25	Thematic analysis of the answers to the comprehension question: ‘ <i>When are you liable for an unauthorised transaction?</i> ’ . . . . .	91
4.26	Thematic analysis of the answers to the question: ‘ <i>What is gross negligence?</i> ’ . . . . .	91
4.27	Thematic analysis of the answers to the question: ‘ <i>What can you do to remember your PIN?</i> ’ . . . . .	92
4.28	Responses to the question ‘ <i>How confident are you that you have understood the T&amp;Cs?</i> ’ . . . . .	93
4.29	Thematic analysis of understanding issues of the T&Cs of the participants. . . . .	93
5.1	Distribution of the changes in the participants’ responses . . . . .	103
5.2	Comparison between 2016 and 2017 for question: ‘ <i>How would you classify the maturity of your organisation’s security awareness program?</i> ’ . . .	106
5.3	Frequency of answers to question 16 in 2016 and 2017 . . . . .	107
5.4	Frequency analysis of 2016’s Q14 vs Q16 . . . . .	107
6.1	Demographics for Company A . . . . .	125
6.2	Factor Analysis in Company A . . . . .	126
6.3	Participant demographics in Company B . . . . .	129
6.4	Factor Analysis in Company B . . . . .	130
6.5	Absolute percentages of Behaviour Types compared with Attitude Levels	141
6.6	Response rates to free-text response questions by Behaviour Type and Attitude Levels . . . . .	142
6.7	The dimensions by which survey responses are measured in Company A & B . . . . .	153

---

6.8	Inter-coder confusion matrix . . . . .	156
6.9	Coder-mapping confusion matrix . . . . .	156
6.10	Krippendorff's $\alpha$ measures for compA with 95% confidence intervals.	158
6.11	Krippendorff's $\alpha$ measures for compB with 95% confidence intervals.	158
6.12	Appropriateness scores for each attitude question. . . . .	160
6.13	Acceptability of failing to complete the task and severity scores analysis.	160
A.1	Additional Factor Analysis in Company A . . . . .	226
A.2	Additional Factor Analysis in Company B . . . . .	228
B.1	Frequency of coded responses to 2016's Q14 . . . . .	245
B.2	Frequency of coded responses to 2016's Q15 . . . . .	246
B.3	Frequency analysis of 2016's Q14 vs Q15 . . . . .	247
B.4	Frequency of coded responses to 2016's Q16 . . . . .	248
B.5	Frequency analysis of 2016's Q16 vs Q15 . . . . .	249
B.6	Frequency of answers to 2016's Q17 . . . . .	249
B.7	Frequency analysis of 2016's Q14 vs Q17 . . . . .	250
B.8	Frequency analysis of 2016's Q17 vs Q15 . . . . .	251
B.9	Frequency of answers to 2017's Q14 . . . . .	253
B.10	Frequency of answers to 2017's Q16 . . . . .	253
B.11	Frequency of answers to 2017's Q17 . . . . .	254
B.12	Frequency analysis of 2017's Q14 vs Q16 . . . . .	255
B.13	Frequency analysis of 2017's Q14 vs Q17 . . . . .	255



# Introduction

Every aspect of human life today is influenced by security considerations, security decisions and security design. Our home internet router is a cyber security device that requires configuration, we shop online, we bank online, we date online, and almost every job has a digital element to it that invariably requires some security considerations. Yet many recent security and data breaches have put security in the limelight and have caused the general public to start questioning whether security is really achieves all that it claims to do.

The common element to all security considerations are humans. They are the ones who design and build the protections, they are the ones that are to be protected, they manage the protection and reap (or suffer) the consequences of these security considerations. This thesis studies their view of security. I build and validate methods and tools to gather the perceptions of individuals when interacting with security, both as employees in organisations and as individuals interacting with organisations and businesses. I measure the impact of security on individuals' behaviours and understanding of systems, and provide evidence as to what constitutes *reasonable*. My contribution is a picture of security from the perspective of those living it, rather than from the view of those that designed it.

In the following **Sections 1.1 to 1.4** I scope the actors and threats considered in this thesis. The threat model here is circular: security specifications are treated both as the solution and the threat to the system. External events and actors are obviously motivators and drivers for some of these specifications, but this thesis does not consider the reactions to these explicitly. Instead, I focus on the externalities of these responses (i.e. the impact on individuals and their behaviours) and build

tools and methodologies that are holistically applicable.

I state my research questions in [Section 1.5](#) and then subsequently outline the contributions of the remaining chapters and their relevance to them.

### 1.1 Security scope

Security is in its essence defined as a freedom from harm. This emphasises that security should be supportive for individuals and organisations alike. It should support individuals to live their lives and enable productivity, as without this security the productive task may not exist or not produce fruitful results. By this measure security should not be in conflict with the productive tasks. However, as it is often the case, the manifestation of security is not perfect. This results in inefficiencies, and failures in the capacity to protect: this is the cost of security. The understanding of threats and consequences differs between security managers, security architects and individuals; these differences are often sources of conflicts. Further, the decisions around security are (and will) always be based on incomplete information of the end users' context. My work, particularly through the Productive Security and perception work in [Chapters 5 and 6](#), provides an overview of the existing security culture, and how this end user/employee culture sits with the intended security culture of the organisation. Consider the following example: an organisation decides to implement a clear desk policy in order to protect confidential information, but at the same time, many business processes are paper based. If employees follow the policy rigorously, an additional 20 minutes are spent each day on storage and retrieval of these physical documents from the secure lockers. Some employees adapt their working practices by batching tasks, which causes delays in other departments. Others decide to hide the documents under their desks overnight. Only by repeatedly re-assessing this disparity can improvements to the inefficiencies of security objectively be made. The research in this thesis contributes to the methodology of identifying and measuring these conflicts and misalignments in a real-world setting, and investigates if user needs are considered or represented in security decisions, and how this might explain the resultant user response.

The security I consider here is essentially a solution for managing the risk of non-intended consequences and events. One may manage the risk by reducing the likelihood or the impact of the consequences, by transferring the risk to another party who is better suited to managing the risk, or by mitigating it entirely. In the environments that I study as part of this thesis, this management of security is described through a statement of intent, the security policy. The policy can be driven by external pressures (mainly regulatory, or even as re-affirmations of laws)



or through requirements that enable productive activity (for example, statements designed to protect confidential information such as intellectual property). In many cases, aspects of the policy evolve over time in response to events. The rigorous and repeatable methodologies presented in this thesis are sufficiently granular and targeted to capture these complexities. In principle, the policy should be applicable to those and workable for those that should adhere to it. This aspect is fundamental to this thesis, and I investigate this claim in organisations and other environments.

## 1.2 Security manifestation

The security policy is implemented through either soft or hard measures. Soft measures are guidelines, *dos* and *do not dos* that rely on the individuals goodwill (or threat of sanctions) for execution. Hard security measures are reliant on systems to enforce the policy such as physical access control measures, firewalls, internet gateway proxy servers or password expiration. Individuals can work around these measures or develop coping strategies (see [Section 2.4](#)) to manage these security measures. A significant part of this thesis focusses on understanding the reasons behind individuals' strategies when faced with security measures.

Often new security measures are added in response to new threats being discovered. These new security measures are regularly implemented through restrictions to existing procedures and modifications to protocols. In essence, the new security is bolted on. This reactive security behaviour causes a number of frictions: productivity is harmed through non-optimal procedures. Further, the new policies may not actually be *workable*, causing those delivering the procedure to have to be non-compliant in order to get their work done (see [Section 2.2](#)). Essentially, modifications to existing procedures and protocols require the cooperation of the individuals who execute them (Guo et al. 2011), but also need to acknowledge that employee populations are not homogeneous. As part of my research, I have focussed on the impact of modifications to policy on security behaviours (see [Section 1.6.3](#)). There are many aspects that influence whether these modifications are successfully applied, and this thesis contributes to the existing knowledge base.

## 1.3 Security awareness

In order to encourage the uptake of good security practices, security awareness professionals promote the policy and attempt to improve security behaviours in the organisation. They often act as a mediator between policy and practice, as

individuals actively promoting the policy are more approachable than the authors of the policy who may have few actual interactions with the rest of the organisation.

Security Awareness is often supported by knowledgeable individuals in teams, who are often called Security Champions as they act as promoters of security in their local teams, even though it is not their official role. As part of this thesis, I study which type of security champions organisations should promote for their organisational circumstances. I find that there is a role in the development of effective and secure organisation security policies for employees who not only follow and promote policy, but also those who question the adequacy of policy, or challenge it through finding alternative solutions; socialise security solutions through engagement with peers; and those who would expect security to justify itself by being a critical part of their productive work. These findings demonstrate that organisations can find a range of security champions if they engage with distinct viewpoints on policy. In [Section 6.8](#) we find that investing solely in Security Champions who rigidly follow policy misses opportunities to involve the wider organisation in the shaping of effective and workable security.

### 1.4 Security decision making

Security decisions dictate the organisations security policy. This thesis focusses on the effect of the policies. The understanding of productive tasks and feedback from the actual executing individuals can only benefit the quality of the decision making process.

Through the methodologies presented in this thesis, I contribute to building a thorough, measurable understanding of the security behaviours in an organisation, with the knowledge of existing frictions between organisational levels. I have provided means to rigorously and repetitively measure the impact, or indeed co-existence, of security decisions upon end users and employees and their productive tasks. This understanding can then be used to articulate a workable security policy that is *good enough* for the organisation.

### 1.5 Research questions

The studies presented in this thesis focus on understanding user perceptions and behaviour when faced with security, in particular when security changes. This requires us to be able to measure user intentions, behaviour and understanding in a comparable way. I address this challenge through the work on the banking project (overview in [Section 1.6.1.1](#), more detail in [Chapter 4](#)), on the Productive Secur-

ity project (overview in [Section 1.6.2](#), and more detail in [Chapter 6](#)) and on the password project (overview in [Section 1.6.3](#), more detail in [Chapter 7](#)).

In particular, I consider the following research questions (RQ):

- RQ1** what is the capacity of data collection tools to identify conflicts between security and other organisational drivers (such as the organisation’s security behaviours)?
- RQ2** How can this data be used to allow organisations to target remediations?
- RQ3** How are user needs considered or represented in user-facing policies, and how might this explain the resultant user response?
- RQ4** Does involving users in security decision making affect the organisation’s security behaviours?
- RQ5** Can security compliance be measured through individual responses to situations that are in the individual’s context of use?

I discuss these questions in more detail throughout the remainder of the introduction; in particular when giving an overview of all projects presented in this Thesis in [Section 1.6](#). I discuss the natural extension to my research in the final section of my thesis ([Chapter 8](#)), where I summarise my findings and discuss if one can combine user expectations and needs with security requirements to devise security that is *good enough* for the organisation and its members/customers/users as a whole.

## 1.6 Overview of thesis

In this section I give an overview of the research projects that form parts of this thesis as well as outline the structure of the content. The current, introductory chapter serves as motivation and summary. The literature review in [Chapter 2](#) focusses on the related literature that is relevant to the thesis as a whole, while individual chapters relating to specific studies have their own short sections on relevant related literature (i.e. [Sections 4.2](#), [6.2](#) and [7.2](#)). The main literature review chapter is followed by a short chapter where I specifically examine survey constructs used in measuring security behaviours ([Chapter 3](#)).

The main research studies in this thesis commence by understanding the complexities of policies, which in turn motivate the subsequent chapters. In the banking studies, I combine a qualitative analysis of banking Terms and Conditions with two survey studies (a short summary is given below in [Section 1.6.1.1](#), the full research is described in [Chapter 4](#)), highlighting the divergence between the expected behaviour assumed by the policy and actual behaviour by consumers, as well as the difficulties in comprehending the actual Terms and Conditions.

This is followed by two studies on how security aspects are perceived by end users (Section 5.1) and security professionals (Section 5.2), highlighting the demand for better understanding and measuring security behaviours.

As I will discuss in Chapter 3, much of existing research uses survey instruments to deduce what the likely behaviour might be, rather than explicitly going out and observing peoples behaviours. In contrast, the survey methodology developed in Productive Security (Section 1.6.2 and Chapter 6) attempts to approach the inferred behaviour problem by putting the questions into realistic circumstances for the security behaviour. I also develop an approach for statistically evaluating the validity of this approach.

Based on the measurements obtained through the Productive Security methodology, I enhance the notion of Security Champions to become a useful tool to adapt security interventions to an organisation's security behaviours (Section 6.8).

The security behaviours in an organisation are particularly apparent when challenged by a change in policy. In the password study (Section 1.6.3 and Chapter 7), a policy change involves users, allowing them to determine the expiration of their password lifetime through the complexity of their passwords. I measure how users respond to this policy change, allowing the organisation to fine tune their intervention.

I will now discuss these topics in more detail.

### 1.6.1 Motivational research

The research described in this section serves as motivation for later work, and directly describes where policies, user behaviour and understanding diverge.

A common approach to understanding security behaviours is to build behavioural models using surveys (see Chapter 3). These surveys utilise a number of survey constructs that infer some human characteristic through a set of abstract questions, and make predictions on actual behaviour. I think that understanding human behaviour is essential to improving security. But rather than building an explanatory model of human actions (which leads researchers and practitioners to believe that they can fix the problem by altering humans so that the model predicts a more favourable outcome), I think that studying human perceptions of security and the interactions between security and primary tasks is more fruitful. It focuses our future research more towards making security actually *workable* for everyone and the business, and it allows us to understand the motivations behind users' behaviours that were unintended by the engineers.

The following publications underpin this section:

**Are Payment Card Contracts Unfair?** by Steven J. Murdoch, **Ingolf Becker**, Ruba Abu-Salma, Ross Anderson, Nicholas Bohm, Alice Hutchings, M. Angela Sasse and Gianluca Stringhini, *Financial Cryptography and Data Security*, 2016 (Murdoch et al. 2016);

**International Comparison of Bank Fraud Reimbursement: Customer Perceptions and Contractual Terms** by **Ingolf Becker**, Alice Hutchings, Ruba Abu-Salma, Ross Anderson, Nicholas Bohm, Steven J. Murdoch, M. Angela Sasse and Gianluca Stringhini, first appeared in WEIS 2016 (Becker et al. 2016), and then in an extended form in the *Journal of Cybersecurity* 2017 (Becker et al. 2017);

**Metaphors considered harmful? An exploratory study of the effectiveness of functional metaphors for end-to-end encryption** by Albesë Demjaha, Jonathan M. Spring, **Ingolf Becker**, Simon Parkin and M. Angela Sasse, in USEC 2018 (Demjaha et al. 2018);

**Awareness Is Hard: A Tale of Two Challenges & It's Time to Communicate** Two SANS Security Awareness reports for which I analysed the data (SANS Securing The Human 2016, 2017).

The first two items from the list are summarised below in [Section 1.6.1.1](#), and are discussed in more detail in [Chapter 4](#). The last two items are described in [Chapter 5](#) with summaries given below in [Sections 1.6.1.2](#) and [1.6.1.3](#).

### 1.6.1.1 Banking research

The banking studies investigate to what extent bank customers understand the Terms and Conditions (T&Cs) they have signed up to. If many customers are not able to understand T&Cs and the behaviours they are expected to comply with, they risk not being compensated when their accounts are breached.

We conduct an expert analysis of 30 bank contracts across 25 countries and find that most contract terms are too vague for customers to infer required behaviour. In some cases the rules vary for different products, meaning the advice can be contradictory at worst. While many banks allow customers to write Personal Identification Numbers (PINs) down (as long as they are disguised and not kept with the card), 20% of banks categorically forbid writing PINs down, and a handful stipulate that the customer have a unique PIN for each account.

We test our findings with two surveys. First, we build an understanding of customers' PIN behaviour by conducting a survey of 241 British residents. We find that while only a third of PINs are ever changed, almost half of bank customers write at least one PIN down. We also find bank conditions that are too vague to

test, or even contradictory on whether PINs could be shared across cards. Yet, some hazardous practices are not forbidden by many banks: of the 22.9% who re-use PINs across devices, half also use their bank PINs on their mobile phones.

The second survey compares the participant’s perceptions and understanding of the T&Cs themselves. We recruit 151 participants in Germany, the US and UK. Their responses mostly agree with the outcomes of our expert evaluation of the contracts: only 35% fully understand the T&Cs, and 28% find important sections are unclear. There are strong regional variations: Germans find their T&Cs particularly hard to understand, and US bank customers assume some of their behaviours contravened the T&Cs, but are reassured when they actually read them.

We conclude that many bank contracts fail a simple test of fairness, and ‘strong authentication’, as required by the Payment Services Directive II, should include usability testing.

These three studies neatly highlight the gaps between actual security policy, security policy understanding by customers, and actual security behaviour.

### 1.6.1.2 Metaphors research

Research has shown that users do not use encryption and fail to understand the security properties that encryption provides. We hypothesise that one contributing factor to failed user understanding is poor explanations of security properties, as the technical descriptions used to explain encryption focus on structural mental models. We methodically generate and analytically evaluate metaphors for end-to-end (E2E) encryption that cue functional models and test the effect of these metaphors and other commonly used explanations on users’ understanding of E2E-encryption through a survey with 211 participants.

While the analytical evaluation showed promising results, none of the descriptions tested in the survey improve understanding; descriptions frequently cue users in a way that undoes their previously correct understanding. Metaphors developed from user language are better than existing industry descriptions, in that ours cause less harm.

We conclude that creating explanatory metaphors for encryption technologies is hard. Short statements that attempt to cue mental models do not improve participants’ understanding. Better solutions should build on our methodology to test a variety of potential metaphors, to understand both the improvement and harm that metaphors may elicit.

My main contribution to this research (Demjaha et al. 2018) is the design, execution and analysis of the survey with 211 respondents. I also took part in the

metaphor coding exercise and performed the statistical analysis. This research is relevant to this thesis because it highlights that oversimplifications and attempts to explain complex issues through analogies is rarely successful. Individuals need to have some level of understanding of the technical details in order to be able to understand the consequences of their actions and interactions with the system.

### 1.6.1.3 SANS Securing the Human research

The SANS Securing The Human group conducts an annual questionnaire on the state of security awareness in the industry. The questions capture professional demographics, and attempt to discover broad trends throughout the community. In [Section 5.2](#), I have included a short extract from the survey as it demonstrates the lack of scientific approaches to improving organisational security behaviours and therefore serves as motivation and impact of my research: security awareness professionals are desperate to better their interventions.

I have been fortunate enough to gain access to the raw survey data, and I have been particularly interested in the metrics employed by professionals in the security awareness field as well as the effectiveness of security awareness in general. The SANS team creates a well-designed report on the quantitative results each year, and my analysis has been part of this report in 2016 and 2017.

The survey highlights the necessity of my research: awareness interventions need to be supportive of the organisational environment that they are deployed in. My research in [Chapter 6](#) identifies the organisational structures that should be targeted as well as the security behaviours, through Security Champions, that will aid the success of the intervention.

The interventions need to be harmonised with other advice so that individuals do not have to reconcile opposing arguments (such as the conflicts I have identified in [Chapter 4](#)), and the advice has to be understandable (unlike advice I have studied in [Chapter 4](#) and [Section 5.1](#)).

Further, when interventions are deployed, the response should be measured, for example as my research on the password project demonstrates ([Chapter 7](#)).

## 1.6.2 Productive Security

Productive Security was one of three projects of the EPSRC funded Research Institute in the Science of Cyber Security (RISC). I was not funded through this project, however my thesis topic was chosen to support the work of this project.

The work presented here responds to the findings from the security perceptions work above by developing and executing a methodology that explains the

organisational security habits and allows interventions to be tailored to the existing behaviours.

From this project, the following publications form the basis of the work described in this thesis:

**Productive Security: A Scalable Methodology for Analysing Employee Security Behaviours** by Adam Beutement, **Ingolf Becker**, Simon Parkin, Kat Krol and M. Angela Sasse, in SOUPS 2016 (Beutement et al. 2016);

**Applying Cognitive Control Modes to Identify Security Fatigue Hotspots** by Simon Parkin, Kat Krol, **Ingolf Becker** and M. Angela Sasse, in SOUPS 2016 Workshop on Security Fatigue (Parkin et al. 2016);

**Finding Security Champions in Blends of Organisational Culture** by **Ingolf Becker**, Simon Parkin and M. Angela Sasse, in EuroUSEC 2017 (Becker, Parkin and Sasse 2017a);

**Measuring the Success of Context-Aware Security Behaviour Surveys** by **Ingolf Becker**, Simon Parkin and M. Angela Sasse, in LASER 2017 (Becker, Parkin and Sasse 2017b).

My research on this project focuses on two surveys delivered to the employees at two large, multi-national organisations (referred to Company A & B). We received 1486 responses at Company A and 641 responses at Company B.

This project is summarised in the following sections. The studies are discussed in detail in [Chapter 6](#).

### 1.6.2.1 Background

Organisational security policies are often written without sufficiently taking into account the goals and capabilities of the employees who must follow them. Effective security management requires that security managers are able to assess the effectiveness of their policies, including their impact on employee behaviour. Security managers define policies and procedures to express how employees should behave to ‘do their bit’ for information security. They assume these policies are compatible with the business processes and individual employees’ tasks as they know them. Security managers usually rely on the ‘official’ description of how those processes are run; the day-to-day reality is different, and this is where security policies can cause friction. Organisations need employees to participate in the construction of workable security, by identifying where policies causes friction, are ambiguous, or just do not apply.

Yet security tasks can burden the individual, to the extent that security fatigue promotes habits that undermine security.



### 1.6.2.2 Aim

Effective security management requires that security managers are able to assess the effectiveness of their policies, including their impact on employee behaviour. Current efforts to involve employees in security act to identify employees who can be local representatives of policy—as with the currently popular idea of ‘Security Champions’—rather than as a representative of employee security needs.

Additionally, as a reflective task, we aim to establish the usefulness of framing survey questions around active security controls and problems experienced by employees, by assessing the validity of the clustering. We introduce measures for the appropriateness of the survey scenarios for each company and the quality of candidate answer options. We use these scores to articulate the methodological improvements between the two surveys.

### 1.6.2.3 Method

We present a methodology for gathering large scale data sets on employee behaviour and attitudes via scenario-based surveys. The survey questions are grounded in rich data drawn from interviews, and probe perceptions of security measures and their impact. We analyse the responses by clustering the data on various demographic information to establish distinguishing features in the organisation, and attempt to identify a structural analysis towards helping organisations ‘close the loop’ and get input from employees.

We develop a methodology to verify the clustering of participants, where 516 (Company A) and 195 (B) free-text responses are coded by two annotators. Inter-annotator metrics are adopted to identify agreement. Further, we analyse 5196 (A) and 1824 (B) appropriateness and severity scores to measure the appropriateness and quality of the questions.

### 1.6.2.4 Results

We demonstrate that our approach is capable of determining important differences between various population groups. The post-hoc validation raises questions on a number of clusterings, although Krippendorff’s  $\alpha$  improves for reliable questions by 0.15 from A to B. We find that the scenarios presented in B are more recognisable to the participants, suggesting that the survey design has indeed improved.

The analysis of 608 responses finds that attitude to policy and behaviour types—the prevailing security cultures—vary greatly in Company B and across four business divisions examined in further detail. There is a role in contributing to the effective-

ness of security policies not only for those who follow policy, but also for those who question policy, socialise solutions, or expect security to justify itself as a critical part of their productive work.

By examining routine security behaviours, we identify perceived contributors and consequences of security fatigue, and the strategies that a person may adopt when feeling overburdened by security. Behaviours and strategies are framed according to a model of cognitive control modes, to explore the role of human performance and error in producing security fatigue. Security tasks are then considered in terms of modes such as unconscious routines and knowledge-based ad-hoc approaches.

### 1.6.2.5 Conclusions

Our work has been used to set policy within the partner companies, illustrating the real-world impact of our research. Security Champions cannot be uniform across the organisation. Organisations should re-think the role of security champions as diverse ‘bottom-up’ agents to change policy for the better, rather than communicators of existing ‘top-down’ policies.

However, to be able to draw valid conclusions from survey responses, the train of analysis needs to be well documented. Our approach allows us to further validate the clustering of responses by utilising free-text responses. Further, we establish the relevance and appropriateness of the scenarios for individual organisations. While much prior research draws on survey instruments from research before it, this is then often applied in a different context; in these cases adding metrics of appropriateness and severity to the survey design can ensure that results relate to the security experiences of employees.

### 1.6.3 Password project

The password project was a RISCS small grants funded project where I studied the effect of a policy change by measuring the impact on user behaviour through log events.

It supports this thesis by evaluating a security policy change and measuring the response of users.

From this project, the following publications form the basis of the work described in this thesis:

**Quantifying the impact of password policy change** by Simon Parkin, **Ingolf Becker**, Albesë Demjaha, Julianne Park, Nissy Sombatruang and M. Angela Sasse, unpublished report (Parkin et al. 2017);

**Rewarding Users for Stronger Passwords: Linking Password Lifetime to Strength** by **Ingolf Becker**, Simon Parkin and M. Angela Sasse, in USENIX Security 2018 (Becker, Parkin and Sasse 2018).

This research is a case study where we observe a change in UCL’s password policy. We analyse the password change behaviour of 100,000 staff and students over a period of 14 months. The following sections outline the research, and the full research summary can be found in [Chapter 7](#).

### 1.6.3.1 Aim

The goal of the IT staff who conceived the policy was to encourage stronger passwords by varying password lifetime according to password strength. Strength was measured through Shannon entropy (acknowledged to be a poor measure of password strength by the academic community, but still widely used in practice). When users change their password, a password meter informs them of the lifetime of their new password, which may vary from 100 days (50 bits of entropy) to 350 days (120 bits of entropy).

### 1.6.3.2 Method

We analysed data of nearly 200,000 password changes and 115,000 resets of passwords that were forgotten/expired over a period of 14 months.

### 1.6.3.3 Results

The new policy took over 100 days to gain traction, but after that, average entropy rose steadily. After another 12 months, the average password lifetime increased from 146 days (63 bits) to 170 days (70 bits). Maths and Physical Sciences schools demonstrate the strongest passwords with 179 days average lifetime (72 bits); research and teaching staff choose one order of magnitude stronger passwords than postgraduate students.

We also found that passwords with more than 300 days of lifetime are 4 times as likely to be reset as passwords of 100 days of lifetime. Users who reset their password more than once per year (27% of users) choose passwords with over 10 days fewer lifetime, and while they also respond to the policy, maintain this deficit.

### 1.6.3.4 Conclusions

Where password expiration has been seen in related studies to have no positive impact on the strength of passwords, it appears that here the inconvenience of

changing a main password pushes users to choose stronger passwords than the minimum requirement at this institution. We conclude that linking password lifetime to strength at the point of password creation is a viable strategy for encouraging users to choose stronger passwords (at least when measured by Shannon entropy).

### 1.6.4 Other research conducted during my PhD

I also contributed to the following publications, however they are outside the scope of this thesis:

**No Good Reason to Remove Features: Expert Users Value Useful Apps over Secure Ones** by Steve Dodier-Lazaro, **Ingolf Becker**, Jens Krinke and M. Angela Sasse, in International Conference on Human Aspects of Information Security, Privacy, and Trust, 2017 (Dodier-Lazaro, Becker et al. 2017);

**From Paternalistic to User-Centred Security: Putting Users First with Value-Sensitive Design** by Steve Dodier-Lazaro, Ruba Abu-Salma, **Ingolf Becker** and M. Angela Sasse, in CHI 2017 Workshop on Values in Computing (Dodier-Lazaro, Abu-Salma et al. 2017);

**Light-touch Interventions to Improve Software Development Security** by Charles Weir, Lynne Blair, James Noble, **Ingolf Becker**, M. Angela Sasse, in IEEE SecDev, 2018 (Weir et al. 2018).

## 1.7 My original contributions of this thesis

This thesis focusses on methodologies to measure interactions with security. Methodologies put policies in contrast with perceptions of them and measures of actual use. Security behaviours are reliably analysed in realistic contexts, and cross-sectional studies and longitudinal studies demonstrate variations of approaches and habits of security.

All of my research has been conducted in collaboration with other researchers. In writing this thesis, I have made use of manuscripts that were written by myself and my co-authors. While I was the first and main author of the majority of publications, and in particular those that I mainly focus on in this thesis, it is likely that I have re-used sections that were not phrased by myself. Throughout this thesis I indicate which original publication the work is based on.

Despite this, the following contributions I consider my own, in order of appearance in this thesis:

1. A systematic review of survey constructs used for understanding security behaviours ([Chapter 3](#)) found that much existing research focusses on identifying

- psychological metrics that explain human factors. The focus in this research is on establishing internal validity rather than on organisational applicability.
2. I reviewed the Terms and Conditions of three banks as part of [Section 4.3](#).
  3. I designed, executed and analysed a survey study to understand customers' payment card behaviour and perceptions ([Section 4.4](#)), finding that participants have many different payment cards and PINs, which cause most customers to violate their banks' Terms and Conditions to some extent.
  4. I designed a survey methodology for contrasting users expectations of policies with their understanding of the policy, taking care not to bias participants' responses [Section 4.5.2](#).
  5. I executed the above methodology in a cross-cultural survey study between the UK, US and Germany to understand customers' perceptions and understanding of banking Terms and Conditions ([Section 4.5](#)), finding that only a minority of individuals understand the T&Cs.
  6. I designed a survey for measuring the understanding of the functionality provided by E2E-encrypted messaging apps using a test-retest survey. I executed and analysed this survey on 211 respondents ([Section 5.1](#)).
  7. I analysed the survey responses of the SANS Security Awareness in 2016 and 2017 ([Section 5.2](#)), highlighting the desire for actual metrics on employee security behaviour.
  8. I created a methodology for continuous evaluation of an organisation's security behaviours through context-aware surveys ([Section 6.3](#)).
  9. I conducted an analysis of 1486 survey responses at Company A and 641 survey responses at Company B ([Section 6.4](#)), finding strong regional and divisional differences in security maturity, attitude and behaviour types that explain responses to security conflicts.
  10. I created a methodology for validating context-based surveys ([Section 6.4](#)) by manually coding free-text responses to questions that align with the mapping of the context-based surveys. I executed it on the two surveys ([Section 6.10](#)) with two annotators each, finding that a number of mappings are not sufficiently supported by the validation.
  11. I identified profiles of Security Champions in organisations ([Section 6.8](#)), arguing that the security culture of individual organisational units warrants a specific security mark-up of a security champion in order to promote culture change efficiently.
  12. I measured the effect of a policy intervention at an institution with over 100,000 users for over a year ([Chapter 7](#)), finding that throughout the institution users will choose stronger passwords when rewarded with longer password lifetime.



## Literature review

This chapter puts my research on measuring and understanding security behaviours into the perspective of existing research. I establish the state of the art and identify the gaps in existing knowledge that my contribution in the following chapters aims to fill.

The review begins with an overview of existing methodologies and studies on measuring security (Section 2.1), which is then followed by a review of literature of user-centred studies on explanations and consequences of unworkable security.

### 2.1 Measuring security

Engaging users is important in the development of meaningful, effective security behaviour surveys. If studies are conducted out of context, reproduction of results is difficult (Open Science Collaboration 2015). Yet much of security awareness research examines individuals' abilities to internalise and enact knowledge of security risks and controls in an abstract setting. Efforts to measure security behaviour frequently assess individuals' competency in general security skills. Much of this research ignores the bounded effort of the individual (Beautement, Sasse and Wonham 2008; Herley 2009, 2014), and that employees in organisations have other responsibilities (Ashenden and Lawrence 2016).

Egelman, Harbach and Peer developed the Security Behavior Intentions Scale (SeBIS) to predict security behaviours for common controls (awareness<sup>1</sup>, passwords, updating, and securement) (2016). The SeBIS survey comprises of 16 items on a

---

<sup>1</sup>For clarity, constructs/abstract ideas that are measured are underlined throughout this thesis.

5-point Likert scale. SeBIS was deployed on several occasions through Mechanical Turk (and in one case, PhoneLab). The goal of the work was to determine if self-reported, *intended*, behaviours translated into actual behaviour. To this end, tasks were set relating to each behaviour category (such as identifying fake login pages). The authors accepted that the designed tasks were targeted and narrow in scope, but with a focus on exploring SeBIS' predictive capabilities in this limited setting. Here we use scenarios and options based in real organisational settings to establish an individual's behaviour type and attitude toward the security apparatus around them; the focus is not on predicting behaviour, but rather to capture a snapshot of how effectively security provisions are perceived to be supporting the business.

Parsons *et al.* sought to validate a survey tool for measuring information security awareness and awareness initiatives, the Human Aspects of Information Security Questionnaire (HAIS-Q) (2017). The HAIS-Q is a 63-item questionnaire exploring a number of areas: password management, email use, internet use, social media use, mobile devices, information handling, and incident reporting. The HAIS-Q has been deployed across a number of groups of working professionals. Two studies were conducted: in one study, participants completed the HAIS-Q and were tested for security skills (in this case, identifying potential phishing links amongst a range of fabricated emails); in the second study, engagement of participants in the survey was examined by establishing the level of non-responsivity. Here, we similarly seek to determine whether the scenarios and response options in our surveys resonate with participants, through examination of internal measures within our situated surveys. By doing so we identify repeatable measures for measuring engagement. Parsons *et al.* suggested applying the HAIS-Q to measure the effectiveness of security awareness initiatives, where here we measure the capacity of a security behaviour survey tool to measure employees' relationship with the security apparatus—be it provisioned or appropriated—that is available to them in the organisation.

Rajivan *et al.* propose a questionnaire for capturing users' level of security expertise, presented as being a critical factor in how well an individual can assess risk and use available security controls (2017). The questionnaire seeks to separate respondents across the dimensions of skills, rules, and knowledge, toward understanding how individuals apply these in different situations. Here we discuss our survey methodology as a means to not only determine how employees use the tools available to them as individuals and groups, but also how they respond to specific risks which can potentially arise in their working environment. Rajivan *et al.* also included free-text questions to capture additional comments from participants, where we use a similar internal mechanism in our situated scenarios so that participants can further describe security experiences from their own perspective (further



informing the picture of security on the ground).

Karlsson, Karlsson and Åström posit that in organisations, information security compliance must be evaluated relative to employees' work tasks (and with this, competing goals and their related *values* such as productivity and efficiency) (2017). The authors speak of there being '*tensions and dilemmas*' where one option is preferable to others that are available.

The authors evaluate both a value-monistic (non-contextualised) measure of compliance (i.e. '*a traditional measure of information security compliance that asks respondents to report their compliance with information security regulations*') and a value-pluralistic (contextualised) measure of compliance ('*Reports respondents' tendencies to neglect information security regulations in situations when information security has come into conflict with other organisational imperatives*'). The authors results highlight the importance of contextualising security measures to individual's circumstances:

*'When a value-monistic measure of compliance was adopted as the dependent variable, the results suggested that information security compliance was a function of employees' intentions to comply, as well as their self-efficacy and awareness of information security policies. However, information security compliance was not at all related to the occurrence of conflicts between information security and other organisational value systems.*

*When the dependent variable was changed to a more value-pluralistic measure, the results suggested that employees' compliance behaviour was to a great extent a function of the occurrence of conflicts between information security and other organisational values. In a situation where information security policies come into conflict with the efficiency and effectiveness of workflow, as well as the personal integrity and well-being of individuals, employees are less likely to comply.'*

(Karlsson, Karlsson and Åström 2017, Section 6)

The measures used by Karlsson, Karlsson and Åström of *contextualised* compliance are still generic and not tailored to an organisation specifically, yet the difference in results are considerable. This leads us to believe that properly contextualised, scenario-based survey responses improve external validity even further.

We argue that surveys that are situated in scenarios that the participants can relate to will engage them and evoke genuine responses, which can inform efforts to improve the effectiveness of security solutions in an organisation. Some research has focused on basing scenario design in literature, where Blythe (2015) argues

that scenarios should *‘[avoid] unusual events and characters but nonetheless resonate with the respondent in a way that they are readily understood while presenting multiple solutions.’* Siponen and Vance argue that research needs to be practically relevant by ensuring contextual relevance (2014). Their five suggestions focus on studying information system policy violations, but are equally transferable to other behavioural research (and related to the principles derived by Krol et al. (2016) for studying usability in security and privacy). Many examples of security behaviour research using questionnaire instruments do not consider the role of task conflicts characterised by Karlsson, Karlsson and Åström (2017).

These works instead draw on existing questions from prior research (e.g. Sohrabi Safa, Von Solms and Furnell (2016), Ifinedo (2014), Tamjidyamcholo et al. (2014), Cheng et al. (2013) and Witherspoon et al. (2013) as described in Chapter 3). The importance of scenario-based surveys is underlined by Wash, Rader and Fennell’s findings that individuals do not self-report security accurately: it undermines much of the traditional self-reported constructs used for inferring personal security behaviour (2017).

I will revisit some of these construct based measurement studies in Chapter 3, where I focus on the ancestry and validation of survey constructs in a separate rigorous study.

## 2.2 Factors influencing compliance

Previous work has explored factors and activities within organisations which can influence individual compliance with security policies (e.g., Topa and Karyda (2015)), implying that the act of declaring a security policy does not in itself guarantee compliance. A survey of works in this area by Sommestad et al. implies that factors such as perceived behavioural control and the types of training delivered around policy are reliable predictors of compliance, summarising that for anticipating compliance constructs for values and norms are more effective than systems around sanctions and rewards (2014). Here we consider ways to engage with characteristics of organisational and security culture at scale, toward aligning individual security responsibility and the accessibility of provisioned security systems. In this regard, some works consider the role of ‘security champions’ as role models within distinct groups of employees, as examples that others can follow.

Connolly, Lang and Tygar explore the role of deterrent factors in organisational security and the minimisation of ‘human error’ (2015). The authors examine how organisational culture, national culture, and security countermeasures can influence employee security behaviours. An existing organisational culture framework guided

semi-structured interviews with employees from a range of companies in the USA and Ireland. Where employees are encouraged to make local decisions and voice their opinions within the organisation, they are more likely to comply with security policies and procedures. Excluding employees from the conversation around security encourages non-compliance. Compliant behaviours may also rest on secure working being an integral value of the organisation, and the visible presence of countermeasures such as policy and security training is implied as reducing insecure behaviour. In this thesis, I explore how different security cultures are measured and can be engaged for their strengths, to help organisations define workable behaviours.

Posey et al. consider the differences between information security professionals and other employees in organisations, in terms of their perceptions around security (2014). Security professionals consider workarounds by employees as a threat to the organisation, where here we consider whether such activities—essentially, deviation from what professionals expect—are an opportunity for the organisation to develop security which integrates naturally with business processes (Kirlappos, Parkin and Sasse 2014). They found that both security professionals and other employees were concerned about careless behaviour threatening the security of the organisation, yet also noted the potential systematic threats that stem from insecure, yet mandatory technologies. Security professionals underestimated the negative impact of security infrastructure upon employees, and overestimated employees' tendency to distance themselves from security. The authors note that future solutions could act to unify the differences in views between security professionals and employees, rather than determining which group is 'correct'. In their work, perceptions are framed around responses to security events, where here the survey used to drive analysis asks participants to choose amongst responses to scenarios; both the scenario and the range of responses are based on self-reports from employees. In some sense employees noted potential threats which involved peers around them, whereas a weak theme amongst security professionals was inadequate policies. Organisational culture was considered as a global theme, although different kinds of behaviour were reported as concerns to varying degrees, where here we explore the potential for culture to be made up of a range of behaviour types existing all at once.

Hsu et al. examine extra-role security behaviours: those not specified in the information security policy recognised by the organisation (2015). The influence of social controls, rather than formal controls alone, is also explored. The authors asserted that employee involvement in the development of information security policy is critical, given that 'involvement' is a foundational social control. The authors also found that maintaining a mix of formal and social controls can benefit an organisation. Here we explore whether variations of following and challenging of policy

can be blended to benefit an organisation. Driving compliance through formal controls only is seen to stifle extra-role security behaviour. Combinations of security cultures—security leaders—may be necessary to encourage both proactive security behaviour and allegiance to the fundamentals principles of security policy.

Johnston et al. deploy a scenario-based survey, to 242 individuals with experience of using computers and working in organisations with security procedures, to examine how personality traits and derived perceptions of situational factors (such as sanction certainty and threat severity) determine an individual's inclination toward policy violation (2016). The study focuses on specific dispositional factors, those being stability and plasticity, and how these meta-traits interact with the derived perceptions of sanctions, threat appraisal and coping appraisal. Rather than arguing in terms of dispositional factors, we use scenarios as an experimental construct, to explore participants' perceptions around security. Individuals exhibiting Stability may conform with rules, whereas those exhibiting Plasticity may take more risks but only when there is seen to be a benefit in doing so. The authors note that differing forms of persuasive engagement may be necessary to target different personality types, but that controls such as sanctions must be seen as equally fair to all.

Furnell and Rajendran develop a model of the influences on security behaviour, relating workplace-based influences and workplace-independent influences; colleague behaviour is an element of workplace interactions, for instance (2012). By considering *situational factors*, the authors note that factors such as fatigue and the perceived importance of the primary task can impact the enactment of security behaviours. It is further noted that intention to follow security policy can be separate from factors such as the usability of security controls. Based upon a focus group exercise with a number of professionals, disciplinary procedures are seen as a strong influential factor for compliance, but so is colleague behaviour.

Rather than attempting to shift all employee's behaviours to the characteristics of security champions, I explore whether employees and groups of employees, experiencing policy in practice, can actively contribute to the identification and removal of barriers which may otherwise undo any good intentions towards security.

Similar to Furnell and Rajendran's work above, Gabriel and Furnell explore the characteristics of a 'security champion', positing that awareness, motivation, and compliance are foundations for defining a security champion, and further that the personality of a security champion rests primarily on increased imagination and minimising any tendency toward immoderation (i.e., fixation on short-term gains rather than long-term consequences) (2011). This research highlights the limited view of security champions for promoting passive compliance rather than as an

agent for change. Other factors such as Emotionality, Anxiety, and Altruism were seen as positive elements of a security champion, where here we discuss behaviour types which frame characteristics like these in terms of security culture and the connection to the rest of the organisation. We consider with our security maturity levels that employees have to balance their security competence with a primary task, and that being a security champion may be about finding workable security that allows them to complete their primary, productive tasks in a secure way (for instance limiting the need to make short-term compromises by avoiding cumbersome security controls).

In contrast to the attributes prescribed to security champions above, Beris, Beautement and Sasse identified sixteen theoretical behaviour types based upon the analysis of semi-structured interviews with staff in large organisations (2015). The authors distinguish between risk understanding and affective security. Risk understanding is regarded as an individual's competence, and affective security the person's emotional response to security. One of the theoretical behaviour types identified by Beris, Beautement and Sasse is a *security champion*, someone who is motivated to engage with security while also understanding the risks relevant to their work. An awareness of relevant risks is seen to allow security champions to repurpose their skills to address situations not directly or explicitly addressed by policy; this suggests potential in engaging with staff to shape policy for the better.

Prior experiences of security matter. Users are less likely to install future security updates when they have a prior negative experience of updates (Vaniaea, Rader and Wash 2014). In my own work, we have studied the impact of sandboxing on usability (Dodier-Lazaro, Becker et al. 2017). Sandboxing makes applications more secure, but expert users reject this benefit as features need to be removed. We argue that security experts need to identify user values and deliver on them, rather than persuade users to pay attention to security and make secure choices (Dodier-Lazaro, Abu-Salma et al. 2017).

In early work, Smith studies the evolution of privacy policies in banks and insurance companies (1993). He distinguishes between explicit policies (written down), implicit policies (engrained) and practices (created as necessary). Through 105 semi-structured interviews as well as surveys and examinations of the policies themselves, he concludes that the majority of policies are either non-existing, or implicit. Only for deliberate errors and improper, computerised access there exist explicit policies. Further, he identifies a '*policy-practice gap*' (Smith, 1993, 118), with about half of respondents identifying an explicit mismatch between the statements in the policy and their organisation's behaviour. The explicit policy is only written and updated under external threats to the business model through legislation. Until

then, the burden of resolving the privacy conflicts within the organisation are left to the employees themselves, which is central to this thesis. In the 25 years since Smith's assessment of privacy policies, information systems have progressed to the stage where the policy is not just inadequate, it is essentially unworkable.

Bulgurcu, Cavusoglu and Benbasat posit that the quality and fairness of an information security policy, as perceived by employees, are factors in employee security compliance. 'Security hygiene', as defined by Pfleeger, Sasse and Furnham, is a combination of workable security habits that also delivers effective risk management to the level required by the organisation (2014). This can only be achieved if employees are involved in shaping policy that they can adhere to. Otherwise, employees may wonder why policies or specific rules are necessary, or abandon policy at the first sign of any friction or contradiction with other goals (Kirlappos, Beautelement and Sasse 2013).

Posey et al. (2014) found in their interview-based study that employees saw compliance with policy as adding excessive effort to the primary task, stifled the general working environment, caused frustration or seemed difficult—and potentially irrelevant—even with good intentions toward security. Security policies and employee work activities must then be considered as part of the same conversation, to understand how security fits with the existing environment from the 'ground up' rather than exclusively 'top down'.

At a high level, the work of Beris, Beautelement and Sasse defines 4 behaviour types which exhibit strong affective security, but vary in risk understanding. Of note is that all of these 4 behaviour types can be eager to play an active role in security, engaging with the organisation, but vary in how prepared they are to play that role *relative to policy*. There is then a relationship between the *quality of policy* and the the kind of security leader that is needed. Beris et al. note that a 'Security Champion' understands policy but is also able to adhere to the tenets of the policy even in situations where policy is not defined or under-defined. A 'Willing' team member may follow and champion the precise content of policy by rote, but not understand where it is lacking.

Within large organisations, an individual employee's approach to resolving friction with security may be determined by not only their own inclinations but also their interactions with others around them (Beautelement et al. (2016) and Beris, Beautelement and Sasse (2015) & Section 6.4). Not every employee needs to be a 'security champion' but can still be a useful resource for security, as for instance there may be individuals who have knowledge of the risks affecting the business, but who do not have the skill-set of a dedicated security expert (Beris, Beautelement and Sasse 2015).

Equally, there can be employees for whom security is not a natural part of the job, but a *levy on productivity*. This can inadvertently perpetuate the myth that there is a ‘tradeoff’ between security and productivity (Sasse et al. 2016). In the realm of safety, Dekker deprecates the oft-expressed desire of experts to ‘*get people engaged*’ (2014). He argues that the problem is not that ‘operational people’ need to engage with safety, but that ‘*safety experts are not engaged with operational people*’. Dekker further refers to the work of Pink in that motivation to be safe involves *autonomy, mastery, and purpose* (2011). These elements together allude to a larger purpose than safety, best served with an open dialogue around workable behaviours and meaningful interaction between experts and non-experts. Similar arguments may be applied to security.

Employees who rigidly follow the rules are powerless in situations where rules are unwritten or unclear. Organisations who want employees who can be a ‘hero’ who keeps the organisation secure in unanticipated circumstances need to enable them by supporting individual and collective awareness of the risks and an understanding of actions that could mitigate those effects (Reason 2008).

Where security is important to the business, it will enter the discourse of discussions between peers, who may approach security as a social responsibility. Those who are firmly *part of the organisation*, but not necessarily ‘on board’ with security, are nonetheless part of the ‘pulse’ of the organisation. A proactive security champion or network of champions *will* find gaps in policy and process, so the organisation must have the capability to approach shortcomings in policy. This can nonetheless be done in a way that aligns with the *intent* of the security policy, where this is most naturally achieved through alignment of security with the goals of the business.

## 2.3 Perceived security fatigue

Individuals perform a number of security tasks in their daily lives, both private and professional. These tasks support some primary activity, towards a personal goal such as purchasing an event ticket, or as part of their efforts to comply with an employer’s security policies. Security tasks can burden the individual, to the extent that security fatigue promotes habits that undermine security: they can drain individuals’ ‘*energy reservoir*’ (Pfleeger, Sasse and Furnham 2014) by demanding conscious, laboured cognitive effort.

In my work in collaboration with Simon Parkin, Kat Krol & Angela Sasse (Parkin et al. 2016), we review prior works to identify common occurrences of security fatigue and responses to it.

### 2.3.1 Sources of fatigue

In our work, we find that there are a number of sources of perceived security fatigue:

**Excessive cognitive load.** Recalling passwords, recalling the steps of a complicated security process.

**Excessive physical load and preparedness.** Remembering in advance to carry an authentication token, and coordinating its timely use as a part of a security process (Sasse et al. 2014; Krol et al. 2015).

**Distraction from time-sensitive tasks.** Work pressures can make security demands seem more burdensome (Beautement, Sasse and Wonham 2008).

**Blocking of tasks and missed opportunities.** Waiting for IT support to resolve issues, avoiding new business partnerships because setting up secure IT access seems too troublesome (Sasse et al. 2014).

**Potential embarrassment.** The possibility of being unable to access data in the presence of a client adds excessive stress to the process (Beautement, Sasse and Wonham 2008).

### 2.3.2 Responses to fatigue

Related studies also identified potential responses to security fatigue. These responses may inspire changes to habits should security effort cross a threshold and be perceived as too much to justify maintaining proscribed behaviour:

### 2.3.3 Types of compliance

**Continued goodwill.** This is the starting position, which in an organisation can be regarded as altruism for the greater good of all individuals in that organisation (Beautement, Sasse and Wonham 2008). An appropriate number of tasks to perform, that support *diligence*. The individual will do what they are asked to do in the name of security, by rote, because it is reasonable and achievable.

**Grudging compliance.** Continued, visible compliance with security expectations, but only because technology is constraining behaviour. Critically, this behaviour may be *indistinguishable* from security goodwill if the impact on productivity is not also visible (Beautement, Sasse and Wonham 2008) (an individual appears to be complying with policy, but their productive tasks are suffering an impact elsewhere).

**Shadow security.** Where policy or guidance is not visible—or is not descriptive enough to guide behaviour—individuals may create security solutions of their



own. These solutions leverage individuals' own existing knowledge of security, and may develop in response to immediate needs or persistently effortful security that is imposed upon them. Crucially, shadow security (Kirlappos, Parkin and Sasse 2014) happens because people want to behave securely, but take action on their own (or collectively in groups) to manage security fatigue and develop workable, repeatable solutions that can be called upon to address recurring security challenges. Rather than being an act of pushing back against security, shadow security is an attempt to match security effort to the task in the absence of external support.

**Sub-optimal compliance.** Examples focus on 'batching' of tasks such as in Sasse et al. (2014) and Steves et al. (2014). Productive tasks are performed at times which are sub-optimal, to reduce the need to carry out effortful security routines that are wrapped around any one access to a particular system. Technically compliant with security, but the value of security for the process is diminished. An example would be logging in to a system at the end of the day to perform multiple data entry tasks which could have been carried out at any time, meaning data enters the system later than it normally would.

**Learned resignation.** Individuals have too few tasks to perform. It may not be appropriate to limit individuals' involvement in their own security (Krol et al. 2015). Waiting for face recognition technology to decide whether a person is a human or not is one example (Krol, Parkin and Sasse 2016). This has the same fundamental dilemma as potential embarrassment, where here a person may experience stress at the possible outcome for lack of being able to do anything else to influence the end result.

## 2.4 Coping strategies and workarounds

Balancing the demands of primary, productive tasks and secondary tasks—such as security—introduces cost-benefit dilemmas in which individuals are forced to choose between security and productivity. In particular, security that over-burdens the user and is not aligned with their working practices can become less effective (Beautement, Sasse and Wonham 2008). Security is presented to employees as being for their own good, but can introduce externalities, burdening the individual with indirect costs (e.g., changing an increasing number of passwords at regular intervals) (Herley 2009). Individuals may, rationally, perceive the personal cost of compliance as greater than the security benefits gained. As a result, detecting instances where security and business processes are in conflict is critical for both an organisation's security and productivity.

Users develop coping strategies and workarounds to manage the security effort required from them. Users modify characteristics of the task to make repeated application of a routine that is known to support reaching the primary goal without risk of their work being blocked by security. Coping strategies mean that users reorganise their work and security task through for example batching logins. When batching logins, users do not log in to a system as they need it for their work but batch multiple activities that need to be done on a system to reduce the number of login procedures required. One type of a coping strategy can be disengagement where individuals abandon or avoid the use of technologies or services altogether (Steves et al. 2014). This can include workarounds, and indeed avoidance of the productive task that security was intended to support.

An example of a workaround is when users decide to switch to an alternative to the advocated technology. This can include other security mechanisms (e.g., transfer files on an encrypted USB stick instead of via a file-share), or ‘home-made’, *shadow security* solutions leveraging a personal approximation of security (Kirlappos, Parkin and Sasse (2014) and also discussed in [Chapter 6](#)) (such as a password-protected record of passwords, kept on a personal machine which never leaves the office). These may be workarounds, but also appropriations of other technologies advocated in policy. Banking customers may use an account less often unify all PINs of their payment cards ([Section 4.4](#)).

## Security constructs survey

There has been much prior research to measure aspects of security behaviours. In this chapter, I perform a meta-review of constructs used in security research with the specific aim of evaluating the validity of existing measures. It informs the choice of validation techniques used for my own research, as well as gives guidance to fellow researchers. The main output from this work is a database that will greatly simplify the identification and re-use of previously published constructs.

In particular, I focus on the questions that are used to study aspects of human behaviour through surveys. It is a scientific intervention to support the NCSC's People: The Strongest Link policy. Whilst that policy smashes the traditional view that *'people need fixing'* if they don't comply with security policy, existing practice largely relies on tools such as employee surveys to identify constructs associated with compliant or non-compliant security behaviour, and to measure the effectiveness of interventions to change behaviour. These constructs (and often associated tools for measuring them) have been borrowed from a range of disciplines, including psychology and organisational behaviour, and 'transmogrified' into security versions and/or 'mashed up' with other versions. This section investigates the constructs used and their origins, and provides a first step towards assessing how valid they are.

I have identified 688 publications that study human behaviour in information security or support the survey methodology in those studies. Between them, these publications use 984 constructs from 92 categories to measure various aspects of human behaviour. I present these constructs and publications alongside summary information on their validation in the form of a website, which can be found at

<https://verdi.cs.ucl.ac.uk/constructDB/>.

The majority (695/984) of constructs identified have been re-used from existing literature, although we have identified 217 newly created ‘security constructs’. Most authors do not identify the actual source of their construct, even if they reference prior literature for the validity of their instruments. We conclude that there is little methodological rigour is underpinning the instruments and results of most of the papers since only a minority of publications perform proper validation (43/984) and pre-tests (436/984).

I will also include some case studies of a few commonly used survey instruments. The full content of the review can be found on the web page.

In this thesis I will give describe the methodology for selecting publications and evaluating the reliability of the constructs in [Section 3.2](#). I also include a case study of a few commonly used survey instruments in [Section 3.5](#). The full content of the review can be found on the web page, which is described in [Section 3.3](#). We briefly discuss the findings and limitations in [Sections 3.4](#) and [3.6](#) and conclude in [Section 3.7](#).

## 3.1 Introduction

Security solutions are designed to protect specific data or systems; these are part of organisations that engage in some form of productive activity. The vast majority of organisations consist not solely of data and technology (at least not yet!) but have human actors involved in the productive activity, and are thus socio-technical systems. The traditional approach to information security has been to treat human actors as components whose behaviour can be specified and controlled in the same way as that of the technical elements: via architectural measures (security mechanisms) or formal rules (security policies) (as discussed in [Section 1.2](#)). We know from research into socio-technical systems, human factors and behavioural economics that this perspective is misguided (Sasse, Brostoff and Weirich 2001; Pallas 2009): human actors are principal agents whose behaviour is driven by primary economic incentives set (explicitly or implicitly) by the organisation and themselves. Thus, an organisation can achieve desired security behaviour by designing security that is aligned with its the primary business goals and processes and causes minimum friction or distraction. The NCSC’s People: The Strongest Link policy recognises this and advises engagement of employees and design, rather than the traditional sanctions, as the correct response to rule-breaking (National Cyber Security Centre (NCSC) 2016b).

Security expert's traditional view of rule-breaking in cyber security is that the human actors are at fault and *'need fixing.'* There is a wide-spread belief that *'the faults in employees that need fixing'* can be identified via employee surveys: that there are certain constructs that can be measured and that cause compliance or rule-breaking. The underlying assumption is that those employees who score highly on constructs associated with rule-breaking can be 'weeded out', or transformed through security awareness measures. The effectiveness of interventions such as security awareness campaigns is then also measured through such surveys (or not at all, as I will discuss in [Section 5.2](#)). The collection of data via surveys consumes a precious resource: employee time and effort, and many organisations report that employees suffer from 'survey fatigue'. It is thus reasonable to ask whether they deliver any useful insights towards improving security behaviour. If they do not, the organisation would be better served by using those resources for employee engagement and improving security hygiene (Pfleeger, Sasse and Furnham 2014).

Psychologists have utilised surveys to study human behaviour for decades before information security became a necessity, and many of the questions and techniques are borrowed from that discipline. The central element that I focus on in this section is the (survey) construct, which is an explanatory variable which is not directly observable. It is measured through one or many questions. However, it is often not apparent what has been used previously and what the sources of the questions and constructs are. This makes it difficult to establish prior results and previous findings of the constructs. Further, I have noticed that often researchers and practitioners alike will re-invent the wheel, as they are unaware of existing constructs. This resources will attempt to aid their tasks.

## 3.2 Methodology

In this study I systematically explore the constructs used in survey studies in information security. We are interested in any publication that

- Creates new constructs;
- Applies constructs in the form of user studies;
- Discusses results derived from constructs (i.e. meta analysis).

The exploration was approached through three methods:

1. As a first step a google scholar search for combinations of the terms 'information security', 'security', 'survey', 'questionnaire' and 'construct' was carried out. Unfortunately these terms are incredibly broad, and over 3 million relevant articles were returned by the search engine. The first 30 pages of search

results (i.e. 3000 articles) were analysed against the three criteria above and 124 relevant publications were identified.

2. For every article analysed, the articles that were cited for constructs were added to the analysis queue (going backwards in time).
3. For every article analysed, we used Google Scholar to identify citing publications (i.e. going forward in time). Here we limited us to the 30 most cited publications (some psychology publications had tens of thousands of citations) and added these to the analysis queue if they conformed to the selection criteria above.

Initially I found that steps 2 and 3 above increased the size of the analysis queue exponentially: for every paper analysed we would add 10 papers to the queue. However after analysing 400 publications the queue started to become of fixed length, i.e. for every paper analysed we added one more paper to the queue. I have not finished processing the queue of papers, but I decided after identifying over 1000 constructs (before merging) that a comprehensive view of constructs in security has been achieved.

For each publication, we collected the following data:

- A short description of the research (usually a snippet of the abstract);
- research type and sample size (for example ‘user study with 180 students’ / ‘meta review’ / ‘construct validation study’);
- the source PDF file;
- The constructs used or discussed, where for each construct we collected:
  - The exact sources referenced for the construct, or any comment by the authors if they created the construct themselves;
  - The type of the construct (usually the theory on which the construct is based);
  - whether the article lists the exact questions used;
  - whether the article gives the answer options to the questions (and if so, what type they are);
  - two measures of validation, as described below in [Section 3.2.1](#).

The data was collected through Zotero and multiple excel spreadsheets, with a number of custom scripts. The website is statically build using Flask<sup>1</sup> and Jinja2<sup>2</sup> for templating.

I have also explored the construct database of Muhlenberg College<sup>3</sup>, but it is not security specific and not well maintained.

---

<sup>1</sup><http://flask.pocoo.org/>

<sup>2</sup><http://jinja.pocoo.org/>

<sup>3</sup><http://www.muhlenberg.edu/depts/psychology/Measures.html>

Throughout this section, specific instances of constructs have been underlined.

### 3.2.1 Construct validation

While the collection of the constructs is the primary task of this section, I also have to consider whether they have been used in a valid manner. Construct validity is the extent to which an operationalisation measures the concepts that it purports to measure (Straub 1989). Convergent (is the construct giving the same results as other constructs that measure the same thing), discriminant (does the construct disagree with other constructs in a way that is supported by theory), and nomological validation (does the construct fit into the researchers' world view of constructs/the nomological network (Cronbach and Meehl 1955)) are all considered to be components of construct validity, as well as criterion-related validity and its sub-types, predictive (does it correlate to other variables as theory predicts) and concurrent validity (is the construct able to measure and distinguish between all the nuances that it is supposed to) (Cronbach 1949; Rogers 1995).

It should be noted that these criteria for construct validity are measured through tools that concern themselves with internal validity: they are relating to other constructs and theory and are calculated through a wide range of statistical methods, but not to observed behaviour.

Separate to construct validity there is content validity.

*'Generally speaking, an academic achievement test is considered content valid if and when (a) the curriculum universe has been defined (called the "content domain") and (b) the test adequately samples that universe.'* (Lawshe 1975, 565)

In other words, content validity concerns itself with the external validity of the research. Lawshe developed a methodology based on inter-annotator agreement of subjective expert methods to measure it.

In the context of the Management Information Systems (MIS) literature, Boudreau, Gefen and Straub reviews MIS positivist quantitative methodologies and their validity and reliability (Boudreau, Gefen and Straub 2001). The authors limit their analysis to high-level validation techniques:

- Pretest
- Pilot
- Previous Instrument Utilized
- Content Validity
- Construct Validity

They find that only between 25% and 60% of articles studied in their field perform these analysis techniques. Their research forms the basis of the two columns that I have coded for in this research: ‘Content Validity’ and ‘Pretest’.

Under ‘Content Validity’ I code for the various techniques that researchers can perform to establish the degree to which items in an instrument reflect the content universe to which the instrument will be generalised (Cronbach 1949; Rogers 1995). This aspect is particularly important in the case of surveys in information security, where often constructs are borrowed from other disciplines, and the content validation is never repeated. This validity is generally established through literature reviews and expert judges or panels. Lawshe describes a statistical approach to measuring content validity, although I have not seen it employed in this survey.

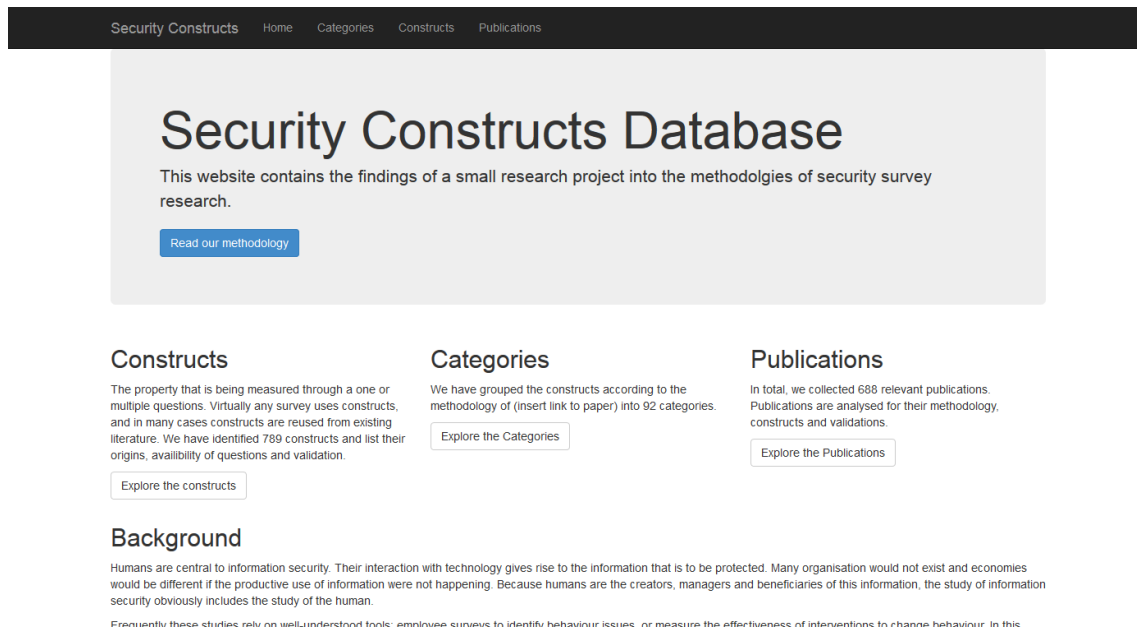
Under ‘Pretest’ I denote the measures the authors have taken to ensure that there are no unanticipated difficulties when executing the survey. Every survey should be pre-tested no matter how skilled the researcher (Fowler 2009). While a study should contain both pre-tests and pilot (the pilot being the dress-rehearsal), many authors use these terms interchangeably, and hence I code both parts here under the heading ‘pretest’.

#### 3.2.2 Construct grouping

There are many constructs that have near identical questions, but have been renamed to capture a specific context, for example Ifinedo (2014, [ConstructDB]) uses the construct Attitude toward ISSP compliance (Information System Security Policy), while Sohrabi Safa, Von Solms and Furnell (2016, [ConstructDB]) uses Attitude towards compliance with ISOP (Information Security Organisational Policies).

Sommestad et al. (2014, [ConstructDB]) conduct a systematic review of empirical studies to identify variables that influence compliance with information security policies of organizations and to identify how important these variables are. 29 studies fit their criteria and they find that none of the 60 factors identified strongly influences compliance. The authors find that even in their relatively small set of studies there is considerable duplication in survey questions and construct intentions. The authors describe a methodology for disambiguating these constructs, as well as the full, detailed disambiguation conducted. This methodology forms the basis for our grouping of constructs that overlap strongly in their questions and intentions. I expand on their table considerably: in our case, I initially identified 984 constructs that the grouping reduced to 789 constructs.





**Figure 3.1:** The home page of the constructDB at <https://verdi.cs.ucl.ac.uk/constructDB/>

## 3.3 The database

The main deliverable of this review is the website. A screenshot of the website can be found in [Figure 3.1](#). In this section I will describe the structure of the website's three main parts: the 688 [publications pages](#) (see [Section 3.3.1](#)), the 789 [construct pages](#) (see [Section 3.3.2](#)) and lastly the 92 [construct categories](#) (see [Section 3.3.3](#)).

### 3.3.1 Publication pages

To continue with the example from [Section 3.2.2](#), Sohrabi Safa, Von Solms and Furnell find that information security knowledge sharing, collaboration, intervention and experience all have a significant effect on employees' attitude towards compliance with organisational information security policies. Attachment does not have a significant effect. They write '*the lack of information security awareness, ignorance, negligence, apathy, mischief, and resistance are the root of users' mistakes*' (Sohrabi Safa, Von Solms and Furnell (2016)) based on 462 questionnaire responses from 4 companies.

Their questionnaire uses nine information security constructs from prior literature to model their interactions. The information gathered for this publication can be seen in a screenshot presented in [Figure 3.2](#). The page begins with a short description of their study as well a short description of their study size. The full, formatted citation is included as well as a the bibliography information in the form

### Sohrabi Safa et al., 2016: Information Security Policy Compliance Model in Organizations

**Topic:**

Information security knowledge sharing, collaboration, intervention and experience all have a significant effect on employees' attitude towards compliance with organizational information security policies. Attachment does not have a significant effect.

Survey, 462 responses, argue that "the lack of information security awareness, ignorance, negligence, apathy, mischief, and resistance are the root of users' mistakes"

**Constructs in this publication:**

Construct	Cites	Category	Questions given?	Content validity	Pretests	Response type	Notes
Information security knowledge sharing	Cheng et al., 2013, Ifinedo, 2014, Tamjidyamcholo et al., 2014, Witherspoon et al., 2013	Social bond/involvement	yes	no		pilot	
Information security collaboration	Cheng et al., 2013, Ifinedo, 2014, Tamjidyamcholo et al., 2014, Witherspoon et al., 2013	Social bond/involvement	yes	no		pilot	
Information security intervention	Cheng et al., 2013, Ifinedo, 2014, Tamjidyamcholo et al., 2014, Witherspoon et al., 2013	Social bond/involvement	yes	no		pilot	
Information security experience	Cheng et al., 2013, Ifinedo, 2014, Tamjidyamcholo et al., 2014, Witherspoon et al., 2013	Social bond/involvement	yes	no		pilot	
Attachment	Cheng et al., 2013, Ifinedo, 2014, Tamjidyamcholo et al., 2014, Witherspoon et al., 2013	Social bond	yes	no		pilot	
Commitment	Cheng et al., 2013, Ifinedo, 2014, Tamjidyamcholo et al., 2014, Witherspoon et al., 2013	Social bond	yes	no		pilot	
Personal norms	Cheng et al., 2013, Ifinedo, 2014, Tamjidyamcholo et al., 2014, Witherspoon et al., 2013	Social bond	yes	no		pilot	
Attitude towards compliance with ISOP	Cheng et al., 2013, Ifinedo, 2014, Tamjidyamcholo et al., 2014, Witherspoon et al., 2013		yes	no		pilot	
ISOP compliance behavioural intentions	Cheng et al., 2013, Ifinedo, 2014, Tamjidyamcholo et al., 2014, Witherspoon et al., 2013		yes	no		pilot	

**Citation:**

Nader Sohrabi Safa, Rossouw Von Solms, and Steven Furnell. Information security policy compliance model in organizations. *Computers & Security*, 56:70–82, February 2016. doi:10.1016/j.cose.2015.10.006.

**Bibtex**

```
@article{sohrabisafa_information_2016,
  abstract = {The Internet and information technology have influenced human life significantly. However, information security is still an important concern for both users and organizations. Technology cannot solely guarantee a secure environment for information; the human aspects of information security should be taken into consideration, besides the technological aspects. The lack of information security awareness, ignorance, negligence, apathy, mischief, and resistance are the root of users' mistakes. In this research, a novel model shows how complying with organizational information security policies changes and mitigates the n
```

**Figure 3.2: ConstructDB Screenshot for Sohrabi Safa, Von Solms and Furnell (2016) at <https://verdi.cs.ucl.ac.uk/constructDB/publications/sohrabisafa-information-2016.html>**

of a bibtex snippet. The paper's fulltext pdf is accessible to machines from inside UCL.

As can be seen from the website, the authors cite for each of their constructs four prior works as sources: Cheng et al. (2013, [ConstructDB]), Ifinedo (2014, [ConstructDB]), Tamjidyamcholo et al. (2014, [ConstructDB]), and Witherspoon et al. (2013, [ConstructDB]). Each of the constructs is linked to their respective construct pages (as described further in Section 3.3.2) and I link the construct's categories (as described further in Section 3.3.3), state if the publication lists the actual question text (yes/no/partially) and the answer format. I also include some short remarks regarding the paper's content validations and pretests applied to the constructs, as described in Section 3.2.1. The origins of the constructs are discussed in more detail in a case study in Section 3.5.

## Ifinedo, 2014: Information Systems Security Policy Compliance: An Empirical Study of the Effects of Socialisation, Influence, and Cognition

### Topic:

Social bonds that are formed at work largely influence attitudes towards compliance and subjective norms, with both constructs positively affecting employees' ISSP compliance.

Questionnaire, 68 and 124 responses

Constructs in this publication:

Construct	Cites	Category	Questions given?	Content validity	Pretests	Response type	Notes
Attachment	Lee et al., 2004	Social bond	yes	expert review	none	7-point Likert scale ranging from Strongly Disagree to Strongly Agree	
Commitment	Lee et al., 2004, Herath, 2009	Social bond	yes	expert review	none	7-point Likert scale ranging from Strongly Disagree to Strongly Agree	
Involvement	Lee et al., 2004	Social bond	yes	expert review	none	7-point Likert scale ranging from Strongly Disagree to Strongly Agree	
Personal norms	Li et al., 2010	Social bond	yes	expert review	none	7-point Likert scale ranging from Strongly Disagree to Strongly Agree	
Attitude toward ISSP compliance	Woon, 2007, Bulgurcu et al., 2010, Herath, 2009, Herath, 2009		yes	expert review	none	7-point Likert scale ranging from Strongly Disagree to Strongly Agree	
Subjective norms	Bulgurcu et al., 2010		yes	expert review	none	7-point Likert scale ranging from Strongly Disagree to Strongly Agree	
Locus of control	Workman et al., 2008	Social control	yes	expert review	none	7-point Likert scale ranging from Strongly Disagree to Strongly Agree	
Self-efficacy	Woon, 2007, Workman et al., 2008, Compeau, 1995	Social control	yes	expert review	none	7-point Likert scale ranging from Strongly Disagree to Strongly Agree	
ISSP compliance behavioral intentions	Woon, 2007, Herath, 2009, Herath, 2009		yes	expert review	none	7-point Likert scale ranging from Strongly Disagree to Strongly Agree	
Detection probability	Herath, 2009, Herath, 2009		yes	expert review	none	7-point Likert scale ranging from Strongly Disagree to Strongly Agree	
Sanction severity	Herath, 2009, Herath, 2009		yes	expert review	none	7-point Likert scale ranging from Strongly Disagree to Strongly Agree	

This publication is cited by the following publications:

- Sohrabi Safa et al., 2016: Information security knowledge sharing
- Sohrabi Safa et al., 2016: Information security collaboration
- Sohrabi Safa et al., 2016: Information security intervention
- Sohrabi Safa et al., 2016: Information security experience
- Sohrabi Safa et al., 2016: Attachment
- Sohrabi Safa et al., 2016: Commitment
- Sohrabi Safa et al., 2016: Personal norms
- Sohrabi Safa et al., 2016: Attitude towards compliance with ISOP
- Sohrabi Safa et al., 2016: ISOP compliance behavioural intentions

### Citation:

Princely Ifinedo. Information systems security policy compliance: An empirical study of the effects of socialisation, influence, and cognition. *Information & Management*, 51(1):69–79, January 2014. doi:10.1016/j.im.2013.10.001.

**Figure 3.3: ConstructDB Screenshot for Ifinedo (2014) at <https://verdi.cs.ucl.ac.uk/constructDB/publications/ifinedo-information-2014.html>**

Obviously, each of these publications have their own pages, similar to Sohrabi Safa, Von Solms and Furnell (2016, [ConstructDB]) in Figure 3.2; for example a screenshot of the page for Ifinedo (2014, [ConstructDB]) is shown in Figure 3.3. This page also contains the paper's dependent publications, i.e. publications that cite it for their constructs. In this case, Ifinedo (2014, [ConstructDB]) was only cited by Sohrabi Safa, Von Solms and Furnell (2016, [ConstructDB]), not for one specific construct, but for all of them.

As Sohrabi Safa, Von Solms and Furnell do not make the exact sources of each of their constructs clear, it is difficult to highlight the exact origins of the constructs, as the four cited articles in turn cite over 100 papers other unique articles as sources of their constructs.

### 3.3.2 Construct pages

#### Construct: Attitude towards compliance

Also quoted in literature as:

- attitude
- personal norms
- attitude towards issp compliance
- Attitude towards compliance with ISOP
- Attitude toward ISSP compliance
- attitude towards information security policy and procedures

This construct belongs to the following groups:

- Compliance
- theory of planned behavior
- Social bond
- HAIS-Q

This construct is used by:

- Sohrabi Safa et al., 2016: [Information Security Policy Compliance Model in Organizations](#)  
Questions given: yes
- Ifinedo, 2014: [Information Systems Security Policy Compliance: An Empirical Study of the Effects of Socialisation, Influence, and Cognition](#)  
Questions given: yes
- Parsons et al., 2015: [The Influence of Organizational Information Security Culture on Information Security Decision Making](#)  
Questions given: no
- Sommestad et al., 2014: [Variables Influencing Information Security Policy Compliance: A Systematic Review of Quantitative Studies](#)  
Questions given: example
- Ifinedo, 2012: [Understanding Information Systems Security Policy Compliance: An Integration of the Theory of Planned Behavior and the Protection Motivation Theory](#)  
Questions given: yes
- Johnston, 2010: [The Influence of Perceived Source Credibility on End User Attitudes and Intentions to Comply with Recommended It Actions](#)  
Questions given: yes

Papers using this construct cite the following papers for the source of this construct:

- Cheng et al., 2013: [Understanding the Violation of IS Security Policy in Organizations: An Integrated Model Based on Social Control and Deterrence Theory](#)
- Ifinedo, 2014: [Information Systems Security Policy Compliance: An Empirical Study of the Effects of Socialisation, Influence, and Cognition](#)
- Tamjidyamcholo et al., 2014: [Evaluation Model for Knowledge Sharing in Information Security Professional Virtual Community](#)
- Witherspoon et al., 2013: [Antecedents of Organizational Knowledge Sharing: A Meta-Analysis and Critique](#)
- Woon, 2007: [Investigation of IS Professionals' Intention to Practise Secure Development of Applications](#)
- Bulgurcu et al., 2010: [Information Security Policy Compliance: An Empirical Study of Rationality-Based Beliefs and Information Security Awareness](#)
- Herath, 2009: [Encouraging Information Security Behaviors in Organizations: Role of Penalties, Pressures and Perceived Effectiveness](#)
- Herath, 2009: [Protection Motivation and Deterrence: A Framework for Security Policy Compliance in Organisations](#)
- Parsons et al., 2014: [Determining Employee Awareness Using the Human Aspects of Information Security Questionnaire \(HAIS-Q\)](#)
- Li et al., 2010: [Understanding Compliance with Internet Use Policy from the Perspective of Rational Choice Theory](#)
- Pahlila et al., 2007: [Employees' Behavior towards IS Security Policy Compliance](#)
- Ifinedo, 2012: [Understanding Information Systems Security Policy Compliance: An Integration of the Theory of Planned Behavior and the Protection Motivation Theory](#)
- Zhang et al., 2009: [Impact of Perceived Technical Protection on Security Behaviors](#)
- Bulgurcu et al., 2009: [Effects of Individual and Organization Based Beliefs and the Moderating Role of Work Experience on Insiders' Good Security Behaviors](#)
- Siponen et al., 2010: [Compliance with Information Security Policies: An Empirical Investigation](#)
- Venkatesh et al., 2003: [User Acceptance of Information Technology: Toward a Unified View](#)

The following constructs were (potentially) derived from this construct:

- [Information security knowledge sharing](#) by Sohrabi Safa et al., 2016: [Information Security Policy Compliance Model in Organizations](#)
- [Information security collaboration](#) by Sohrabi Safa et al., 2016: [Information Security Policy Compliance Model in Organizations](#)
- [Information security intervention](#) by Sohrabi Safa et al., 2016: [Information Security Policy Compliance Model in Organizations](#)
- [Information security experience](#) by Sohrabi Safa et al., 2016: [Information Security Policy Compliance Model in Organizations](#)
- [Attachment](#) by Sohrabi Safa et al., 2016: [Information Security Policy Compliance Model in Organizations](#)
- [Commitment](#) by Sohrabi Safa et al., 2016: [Information Security Policy Compliance Model in Organizations](#)
- [Personal norms](#) by Sohrabi Safa et al., 2016: [Information Security Policy Compliance Model in Organizations](#)
- [ISOP compliance behavioural intentions](#) by Sohrabi Safa et al., 2016: [Information Security Policy Compliance Model in Organizations](#)
- [Attitude towards compliance](#) by Sommestad et al., 2014: [Variables Influencing Information Security Policy Compliance: A Systematic Review of Quantitative Studies](#)
- [Intention to comply](#) by Sommestad et al., 2014: [Variables Influencing Information Security Policy Compliance: A Systematic Review of Quantitative Studies](#)
- [Intention to comply](#) by Sommestad et al., 2014: [Variables Influencing Information Security Policy Compliance: A Systematic Review of Quantitative Studies](#)
- [Normative beliefs](#) by Sommestad et al., 2014: [Variables Influencing Information Security Policy Compliance: A Systematic Review of Quantitative Studies](#)
- [Perceived severity of incident](#) by Sommestad et al., 2014: [Variables Influencing Information Security Policy Compliance: A Systematic Review of Quantitative Studies](#)
- [Perceived vulnerability](#) by Sommestad et al., 2014: [Variables Influencing Information Security Policy Compliance: A Systematic Review of Quantitative Studies](#)
- [Response cost](#) by Sommestad et al., 2014: [Variables Influencing Information Security Policy Compliance: A Systematic Review of Quantitative Studies](#)
- [Response efficacy](#) by Sommestad et al., 2014: [Variables Influencing Information Security Policy Compliance: A Systematic Review of Quantitative Studies](#)
- [Self-efficacy](#) by Sommestad et al., 2014: [Variables Influencing Information Security Policy Compliance: A Systematic Review of Quantitative Studies](#)

**Figure 3.4:** ConstructDB Screenshot for construct [Attitude towards compliance](https://verdi.cs.ucl.ac.uk/constructDB/constructs/attitude-towards-compliance.html) at <https://verdi.cs.ucl.ac.uk/constructDB/constructs/attitude-towards-compliance.html>

Each construct has its own website, although some constructs have been grouped

together as described in [Section 3.2.2](#). To continue the existing example of [Sohrabi Safa, Von Solms and Furnell \(2016, \[ConstructDB\]\)](#), one construct used in the research is Attitude toward ISSP compliance. This construct has been grouped together with Attitude towards compliance and a screenshot of the website can be found in [Figure 3.4](#). The webpage contains the names of all other constructs that it has been grouped together with, as well as the higher level categories. Further, all instances where this construct is used are listed. In this case there are six publications: [Sohrabi Safa, Von Solms and Furnell \(2016, \[ConstructDB\]\)](#), [Ifinedo \(2014, \[ConstructDB\]\)](#), [Parsons et al. \(2015, \[ConstructDB\]\)](#), [Somme stad et al. \(2014, \[ConstructDB\]\)](#), [Ifinedo \(2012, \[ConstructDB\]\)](#) and [Johnston and Warkentin \(2010, \[ConstructDB\]\)](#).

Given this construct, the next two lists show the dependency tree in two directions. First, there are the prior publications that are cited as sources for this construct. I have identified 16 such parent publications. The final table of the construct page lists the constructs that are potentially derived from this publication. These are the constructs that cite any of the publications that contain this constructs (i.e. Attitude towards compliance) as a source. In the case of this construct there are 18 potential derived constructs. From the construct names it is obvious that not all of these are actually related to Attitude towards compliance, however given most authors lack of specifying the origins of their constructs exactly this is the best that can be done from the meta-data alone. I will discuss this shortcoming further in [Section 3.6](#).

### 3.3.3 Category pages

#### Category: **Social bond**

Constructs in this category:

- Attachment
- Attachment to co-workers
- Attachment to immediate supervisor
- Attachment to job
- Attachment to organization
- attitude
- Attitude toward ISSP compliance
- Attitude towards compliance
- Attitude towards compliance with ISOP
- attitude towards information security policy and procedures
- attitude towards issp compliance
- Belief
- Belief in legal system
- beliefs about the purpose and value of different digital-security behaviors
- commitment
- Commitment to Education
- Depend on parents
- Involvement
- Involvement in Organisations
- Organisational commitment
- Organizational commitment
- parent attachment
- peer attachment
- personal norms

**Figure 3.5:** ConstructDB Screenshot for category *Social bond* at <https://verdi.cs.ucl.ac.uk/constructDB/categories/social-bond.html>

As this research has identified 789 unique constructs, it seems sensible to categorise these for simplified browsing. In most cases the categories are taken from the publications describing the constructs. [Figure 3.5](#) shows the constructs that are categorised under *Social Bond*.

### 3.4 Analysis

In this section, I present statistics on the constructs and their validation. Please note that the numbers in this section do not take the grouping of constructs into account. For [Tables 3.1 to 3.3](#) I also split the analysis for the different types of research. Here *Construct construction* includes validation studies performed separately to the creation of the construct. I separately identify studies that use additional data (i.e. non-survey data such as logs or performance metrics) to supplement their analysis. In addition, the *Survey with Scenarios* category contains surveys that are conducted in the context of specific scenarios, where the participant responses are in relation to the scenario presented rather than in an abstract, classical setting. Where relevant, I also split the constructs into *New*, *Maybe new* and *Not new* categories, indicating if the construct is newly created in the publication.

Type	Question Given?				
	Yes	No	Partially	English only	Other
Construct construction	30	12	0	0	0
Meta analysis	0	28	65	0	0
Survey with Data	46	5	3	0	0
Survey with Scenarios	51	15	4	0	0
Survey	474	179	24	17	16
Total	616	239	96	17	16
New	152	47	11	4	3
Maybe new	42	24	5	1	0
Not new	422	168	80	12	13

**Table 3.1:** Distribution of construct questions given by type of the publication.

[Table 3.1](#) lists how many publications include the questions that are used in their constructs. The majority of publications do include them, which greatly benefits the ability to reuse them in future studies. Surprisingly, not all publications focussing on construct creation and validation include the construct questions. Most meta-analyses only contain partial elements of the constructs.

[Table 3.2](#) shows the frequency of the high level categories of content validation I have assigned to the constructs. Sorting exercises vary from Thurstone or Likert's scale generation techniques to more simple group assignment exercises (Thurstone 1931; Likert 1932). A minor number of publications ground their constructs validity in preliminary interviews. About half of publications rely on reviews by experts. These vary from 2 PhD students reading the questions to multiple rounds of large expert panels debating them. However in none of these cases formal methodologies

Type	Content validation				
	Sorting exercise	Grounded in Interviews	Review by Experts	None	Other
Construct construction	2	0	1	37	2
Meta analysis	0	0	0	0	93
Survey with Data	0	0	8	44	2
Survey with Scenarios	0	8	33	29	0
Survey	26	30	186	386	82
Total	43	38	228	496	179
New	27	22	48	101	19
Maybe new	0	0	11	51	10
Not new	16	16	169	344	150

**Table 3.2:** Distribution of types of content validation of a construct by type of the publication.

are employed. I state the more precise form of content validation on each of the publications pages. The majority of constructs are not validated in the publications. In these cases the authors often claim that no validation is necessary as the construct is not new, i.e. they cite prior work for validation. However this claim is dubious considering that nearly half (101/217) of newly created constructs are not validated in their original publication to begin with.

Type	Pilot / pretesting		
	Any	None	Other
Construct construction	10	28	4
Meta analysis	0	0	93
Survey with Data	32	22	0
Survey with Scenarios	52	18	0
Survey	327	375	8
Total	436	443	105
New	97	108	12
Maybe new	38	34	0
Not new	301	301	93

**Table 3.3:** Distribution of types of piloting or pretesting for a construct by type of the publication.

Table 3.3 lists the pilot and pretesting methods used for the constructs. I have grouped all publications that do any pretesting or piloting together into the *Any* column of this table for lack of a more appropriate, fine grained categorisation. The



construct and publication pages of the website contain the specifics of the pretests and pilots performed. There is an approximately 50/50 split between publications that do some and no piloting.

Answer type	Frequency
5-point Likert agreement	141
7-point Likert agreement	152
9-point Likert agreement	2
3-point Likert	16
5-point Likert	83
6-point Likert	4
7-point Likert	163
Dichotomous	22
Other scales	115
None specified	210
Other	76

**Table 3.4:** Distribution of answer option types for constructs

Table 3.4 highlights the diverse set of answer options that researchers choose to present to participants. Likert scales ranging between *Disagree* and *Agree* clearly dominate this table, with 5 and 7 points being the most common granularity. 210 constructs do not have their answer options specified. This obviously impedes reuse of the metric. It is interesting to note that only one construct relies on free-text responses.

### 3.5 A case study of construct origins

In this section we take the publication of Sohrabi Safa, Von Solms and Furnell as a case study and analyse the dependency of the 9 constructs used (Sohrabi Safa, Von Solms and Furnell 2016) [ConstructDB]. The authors cite four publications as sources to their constructs: Cheng et al. (2013, [ConstructDB]), Ifinedo (2014, [ConstructDB]), Tamjidyamcholo et al. (2014, [ConstructDB]) and Witherspoon et al. (2013, [ConstructDB]). I will now discuss each one of these papers in order to trace the origins of Sohrabi Safa, Von Solms and Furnell’s constructs.

Cheng et al. (2013, [ConstructDB]) measure the following constructs:

**Two constructs of deterrence** adapted from D’Arcy, Hovav and Galletta (2009, [ConstructDB]) (who in turn draws on 3 different research articles for their constructs and survey questions), Li, Zhang and Sarathy (2010, [ConstructDB]) (who in turn draw on eight different research articles for their constructs), and

Siponen and Vance (2010, [ConstructDB]) (who designed their own three scenarios, but adopt the majority of their questions from four different research articles);

**Four constructs of social bond, attachment, involvement, commitment, and beliefs** which they adopted from Lee, Lee and Yoo (2004, [ConstructDB]) (who source all their constructs from six research articles) and Chapple, McQuillan and Berdahl (2005, [ConstructDB]) (no direct sources for their constructs are given, but various references describe similar survey designs);

**Two constructs of social pressure** which were adapted from Herath and Rao (2009, [ConstructDB]) (who in turn rely on seven prior research articles for their survey questions);

**Information Systems Security Policy violation intention** appears to have been created by the authors.

Of the constructs Sohrabi Safa, Von Solms and Furnell use, the constructs on Attachment, Commitment are likely sourced from Cheng et al. (2013), who in turn sourced from Lee, Lee and Yoo (2004) and Chapple, McQuillan and Berdahl (2005). Also, Sohrabi Safa, Von Solms and Furnell use the construct Attitude towards compliance with Information Security Organisational Policy, which may be derived from Cheng et al.

Ifinedo (2014, [ConstructDB]) adopts their survey design from nine previous sources:

**Attachment and involvement** scales from Lee, Lee and Yoo (2004) (dependent on six research articles, as above);

**Commitment** is sourced from Lee, Lee and Yoo (2004, [ConstructDB]) (similar to Cheng et al. (2013) above) and Herath and Rao (2009, [ConstructDB]) (who cite seven articles for the sources of their constructs);

**Personal norms** adapted from Li, Zhang and Sarathy (2010, [ConstructDB]) (who cite eight articles as their source, again as Cheng et al. above);

**Two constructs of attitude towards compliance with ISSPs** from Woon and Kankanhalli (2007, [ConstructDB]) (who cite 15 papers for their 9 constructs, but do validated their constructs through expert sorting), Bulgurcu, Cavusoglu and Benbasat (2010, [ConstructDB]) (two articles are cited for a quarter of the survey questions, the others were developed for their study and pre-tested), Herath and Rao (2009, [ConstructDB]) (7 articles, as above) and Herath and Rao (2009, [ConstructDB]) (7 articles, as above),

**Locus of control** scale from from Workman, Bommer and Straub (2008, [ConstructDB]) (who adopt their survey questions from five different articles);

**Self-efficacy** was adapted from [Compeau and Higgins \(1995, \[ConstructDB\]\)](#) (who rely on six prior research articles for their question design as well as designing some themselves), [Woon and Kankanhalli \(2007, \[ConstructDB\]\)](#), and [Workman, Bommer and Straub \(2008, \[ConstructDB\]\)](#) (both as described above); **Detection probability and sanction severity** scales adopted from [Herath and Rao \(2009, \[ConstructDB\]\)](#) and [Herath and Rao \(2009, \[ConstructDB\]\)](#) (again, who cite 7 articles each as the sources for their constructs).

Of the constructs Sohrabi Safa, Von Solms and Furnell use, the two scales on **Attachment** and **Commitment** appear to have been sourced from Ifinedo, or rather from Lee, Lee and Yoo in the first place. The construct on **Personal Norms** and **ISOP compliance intentions** also appear to originate from Ifinedo's article.

And next, [Tamjidyamcholo et al. \(2014, \[ConstructDB\]\)](#) cite thirteen articles that have contributed to their survey design:

**Knowledge-sharing behaviour** is derived from [Hsu et al. \(2007, \[ConstructDB\]\)](#) (who's questionnaire design relies on eleven prior works), and [Lin, Hung and Chen \(2009, \[ConstructDB\]\)](#) (their survey design is based on thirteen prior research articles);

**Usefulness** by [Cheung, Chang and Lai \(2000, \[ConstructDB\]\)](#) (adopted from six prior works, but with major changes), and [Al-Khaldi and Olusegun Wallace \(1999, \[ConstructDB\]\)](#) (their questionnaire seems to be developed independently);

**Social interaction based perceived consequence** by [Huang \(2009, \[ConstructDB\]\)](#) (their survey is based on seven prior works, but has gone through one round of validation with subject experts) and [Chang and Chuang \(2011, \[ConstructDB\]\)](#) (five references for their survey questions are given, but the resulting survey has been subject to one round of pre-testing);

**Reputation-based perceived consequence** by [Chang and Chuang \(2011, \[ConstructDB\]\)](#) (five references, as above), and [Hsu and Lin \(2008, \[ConstructDB\]\)](#) (six prior works are relied upon);

**Affect** by [Jeon, Kim and Koh \(2011, \[ConstructDB\]\)](#) (their survey design is adopted from nine prior works);

**Social factor** was operationalised according to [Bergeron et al. \(1995, \[ConstructDB\]\)](#) (their survey questions are based on seven prior works), and [Hsu and Lin \(2008, \[ConstructDB\]\)](#) (as above);

**Facilitating condition** was measured with items adapted from [Jeon, Kim and Koh \(2011, \[ConstructDB\]\)](#) (as above);

**Information security risk reduction expectation** was based on [Feledi and](#)

Fenz (2012, [ConstructDB]) (no sources are declared), and Tamjidyamcholo et al. (2013, [ConstructDB]) (their questionnaires is based on prior works, but a small pre-test has been carried out).

It appears that Sohrabi Safa, Von Solms and Furnell draw they knowledge sharing constructs from Tamjidyamcholo et al.'s work, although they appear to have been renamed and re-grouped. Further, Hsu et al. (2007) and Lin, Hung and Chen (2009) cite themselves 24 prior work as the source to their constructs.

Witherspoon et al. perform a meta-analysis of organisational knowledge sharing antecedents in 46 studies. It is unclear how Sohrabi Safa, Von Solms and Furnell's questionnaire is influenced by this meta-analysis.

Of the literature referenced by Sohrabi Safa, Von Solms and Furnell and their references in turn, the vast majority source their constructs and survey questions from further literature. A number of articles construct their own questions. The most rigorous of those papers in the chain to do any pre-testing validation of their question design is by Bulgurcu, Cavusoglu and Benbasat (2010), who conduct two rounds of card sorting by 11 students followed by two rounds of pilot testing by 110 individuals. Huang (2009) and Chang and Chuang (2011) also conduct some limited pre-testing.

While it is good scientific practice to rely on constructs that have been rigorously tested in prior works, only one of the papers (Ifinedo 2014) discussed above cites the primary literature which validates the constructs in their original setting (Bulgurcu, Cavusoglu and Benbasat 2010). Further, as throughout the literature, the questions are taken out of context of their original research premise, where the validity of the original validation should be revisited in each new context. It is understandable that a full pre-study is infeasible for every new questionnaire, yet augmenting the survey with additional questions (as described by my research in Section 6.10) to support the measurement of validity post-hoc would not be time-consuming and desirable.

## 3.6 Limitations

There are a number of limitations to this project. Most of these can be addressed through spending considerably more time on this project.

- I do not include exact question texts on the website because this would require extensive extraction and re-formatting of text from PDFs. The authors or publishers may also have to be consulted in order to receive permission to reproduce their scales.

- The identification of direct descendants and re-uses of constructs is poor. The main fault here lies with authors, who don't explicitly state the origins (and not just which publication the construct is taken from, but also stating which exact construct in the cited publication) for each of their constructs. This could be solved by computing the similarity between the child construct and all the constructs in the cited publication, either by hand or using machine learning.
- The grouping methodology that I currently use ([Section 3.2.2](#)) is extended from previous literature. However, as the authors only group constructs from 29 publications their methodology is potentially not suitable for our large set of constructs.
- Even though I have collected 688 relevant publications, there are still a lot of publications not included.

### 3.7 Conclusion

In this review, I identified 688 relevant publications on studying human behaviour in information security and the origins of the constructs used in these publications. In total, these publications instantiate 984 constructs of which 217 are new constructs. I have extract these constructs from the publications, and present the constructs alongside some analysis and their sources in the form of a website.

Our analysis has shown that a quarter of publications do not include the exact questions that were posed to participants. Likert scales, and especially Likert scales asking the participants agreement to a statement dominate the answer options. The majority of constructs are not validated in their studies and less than half of authors do not report on pre-tests or pilots for constructs.

This review motivates my study on measuring and understanding human behaviour in a reproducible, valid manner. The lack of rigorous validity evaluations conducted in the reviewed literature stands in contrast with my work on assessing the validity of our own survey tool in the Productive Security research in [Section 6.10](#).

The models in these studies attempt to explain a lot of security behaviours, but with considerable methodological weaknesses. All reasoning about behaviours is subject to a ground truth, which is the implied correctness of the policy. The policy is treated as good, workable, correct, and is not challenged at all in these works. The following chapter does analyse policies and challenges these assumptions.

If behaviour is measured against incorrect expectations, than resentment in interventions is predictable.

#### 3.7.1 Recommendations for researchers and practitioners

Given the review described above, there are some obvious recommendations to be given:

- When re-using constructs from prior work, explicitly reference the origins of *each question* of the constructs by specifying the original author and construct.
- Be clear about the definition of your construct. As an example, the security awareness construct sometimes covers aspects of an individual's cognitive state of mind, while others consider only procedural aspects (Jaeger 2018).
- Pre-test your survey to ensure that there are no unanticipated difficulties when executing the survey.
- Validate your survey. Construct validity can be measured by statistical analysis, preferably on a pilot dataset. Content validity is more involved (Section 3.2.1). In this review, not a single paper carried out the technique by Lawshe (1975). Sorting exercises should be the minimum validation for new research.
- Consider adding free-text questions to allow participants to add context to their responses. While manually coding many responses is time consuming, it provides opportunities for post-hoc external validation, for example as described in Section 6.10.

Independent of the review of constructs described in this section, an important consideration is the observation by Karlsson, Karlsson and Åström (2017) described in Section 2.1. Most surveys are grounded in an established theory, for example many of the constructs on Security are based on the Theory of Planned Behaviour. This leads to value-monistic measures, where survey questions treat each subject in isolation of all others, and the analysis is reliant on correlation. Yet many aspects of human behaviour are context-sensitive, especially in security.

To rectify, Karlsson, Karlsson and Åström used a more value-pluralistic measure. Perhaps future research could investigate the viability of using more rich questionnaire setups than the omnipresent Likert scales, such as the Q-methodology (Stephenson 1935; O'Leary, Wobbrock and Riskin 2013; Stewart 2007; Burkell et al. 2013).

# Bank Terms and Conditions: perceptions and reality

*‘A policy is a temporary creed liable to be changed, but while it holds good it has got to be pursued with apostolic zeal.’*

(M. K. Gandhi 1922)

The literature review in the previous chapter has challenged the rationality of behavioural security research and practice. Much of previous research considers user behaviour in isolation of context in which it would be applied, and holds users accountable to an uncontested policy. In this chapter, we question this notion by contrasting many policies for the same product to see if the advice given is sound, actionable and reasonable (Section 4.3). For this, we draw on two user surveys to understand individuals actual behaviour characteristics (Section 4.4) and their knowledge of and ability to understand the policy (Section 4.5).

## 4.1 Introduction

The ability to revoke fraudulent bank payments, or at least reimburse the victims of fraud, is one of the main selling points of the consumer banking system and particularly payment cards. The additional security feature is also used to justify the higher transaction fees associated with card payments, compared with payment systems where transactions are final, such as cash and cryptocurrencies. However, whether a customer who becomes a victim of fraud actually is reimbursed depends on the contract between the bank and its customers (which may in turn depend on

national or international legislation), and how a bank chooses to apply the Terms & Conditions (T&Cs) the customer has signed up to. In order to be reimbursed, a fraud victim may need to demonstrate that they have followed security practices set out in the T&Cs, and thus it is very important that customers 1) are able to understand them, and 2) are able to comply with the behaviours stipulated in them.

This research builds on previous research into the fairness of bank terms and conditions, particularly how the rules adapt to changes in technology. The first study, by Bohm, Brown and Gladman, reviewed the terms and conditions of online banking services, which at the time were still in their infancy (Bohm, Brown and Gladman 2000). They found that some bank contracts stipulated that a customer accepting an online banking password also accepted liability for any transactions that the bank claimed were made with that password, regardless of whether the customer had actually made them. Bohm et al. pointed out that the liability had shifted; a forged handwritten signature is null and void in most countries, so a bank cannot make customers liable for forged cheques using its terms and conditions. The banks took advantage of the technology change to escape nineteenth-century consumer protection law. In some countries, such as the US, pressure by consumer-rights advocates led to regulations that require disputed transactions to be refunded.

There is also the issue of affordances; banks permit customers to change PINs ‘so you can pick one that is memorable to you’. For infrequently used accounts, however, customers will often change them to the PIN on the most frequently-used account, because they are afraid they will forget it, and consider re-using a PIN more secure than writing it down. Many T&Cs, however, stipulate that the PIN for the account must be unique. It is technically straightforward for a bank to set a random PIN on every card issued to a customer, and not let them change it. By letting customers change PINs, but forbidding changes most customers will make in order to cope in the small print, the banks are inducing their customers to break the rules, and thus create ground for not being reimbursed in case of a breach. At the regulatory level, there is a tension between direct consumer protection (which might limit PIN change facilities), and the promotion of competition (for which PIN changes are a good thing, otherwise people will be less likely to start using different cards). But to what extent are there inconsistencies between banks in a country (that may lead to confusion), what do customers understand of their obligations, and whether such issues are hidden from the public behind the obscure contract language? This research attempts to find out.

In [Section 4.3](#) of this chapter, I conduct an expert examination of bank terms and conditions around the world, identifying consistency, or lack thereof, both within and between countries. We draw on our diverse research team to sample the major



banks and translate relevant passages into English. We focus on security advice on PINs, bank statements and telephone and online banking.

In [Section 4.4](#), I study how people actually use cards and PINs by conducting a survey with 241 responses. I am interested in whether customers can practicably comply with typical bank contract terms, and whether they actually do in reality.

In the third part of this chapter in [Section 4.5](#), I conduct a cross-cultural study between 151 individuals in the US, UK and Germany. We ask participants if the rules are sufficiently clear, and if they understand the obligation imposed on customers by the banks' terms and conditions. I focus this study on two common cases of bank fraud and supply the participants with the relevant sections of bank terms and conditions from their country.

I conclude in [Section 4.5](#) that if banking rules cannot be understood, then it is unreasonable to expect customers to comply with them. Indeed, they make matters worse: Adams and Sasse demonstrated a long time ago that traditional password and PIN policies require humanly impossible memory tasks (Adams and Sasse 1999). A recent NIST report (Steves et al. 2014) found that in a work context, over 50% of staff write their credentials down in some way. As disputes are often centred around the PIN being written down and kept with the card, with the customer saying they did do not this, and the bank saying that they must have done, we argue that stipulating a behaviour that we know most people cannot follow means the rules are out of date at best, and unfair at worst. Finally, if the liability shifts to the customer, banks face a less than socially optimal incentive to detect and prevent fraudulent activity on their systems.

### 4.1.1 My contribution in this chapter

A collaboration of researchers enabled the work in this chapter, in particular with Alice Hutchings, Ruba Abu-Salma, Ross Anderson, Nicholas Bohm, Steven Murdoch, Angela Sasse and Gianluca Stringhini.

The content of this chapter has been adapted from three of my publications: Murdoch et al. (2016) and Becker et al. (2016, 2017). The following research was conducted and authored primarily by myself:

- The review of the related literature in [Section 4.2](#),
- The analysis of the Terms and Conditions of three banks (from the total of 30 banks),
- The design, execution and analysis of both surveys in [Sections 4.4](#) and [4.5](#).

## 4.2 Related literature

Financial fraud remains an area of concern. In the UK, payment card fraud has increased by 6% to £618 million in 2016 (Financial Fraud Action UK 2017). In the US about 2 million customers actively reported fraud in 2012 (Cheney et al. 2014), and while it is difficult to accurately specify in the US, Sullivan estimates the value of unauthorised third-party fraud transactions through debit and credit card transactions as \$3.8 billion for 2012 (Sullivan 2014). While the direct monetary loss to customers is negligible in the US thanks to strong consumer protection, major data breaches have further repercussions as more personal identifiable information is stolen because fraudsters use this very data to attack customer accounts. Sullivan also compares the loss due to fraud between (amongst others) the US, UK and the Single Euro Payment Area (SEPA) (Sullivan (2014), Chart 9): on a per-transaction basis over the period of 2005–2012, SEPA had the lowest fraud loss, followed by the UK and US.

Despite the significant volume of payment fraud, research on understanding its implications is surprisingly limited. In the US, the Federal Reserve Bank of Philadelphia regularly runs its own conference series on *Consumer Credit & Payments* and the Federal Reserve Bank of Kansas City one on *Payments*. Stanley discusses work by Hogarth and Hilgert (primary text unavailable) where 18% of households had complaints with their credit card provider, but fraud is not mentioned (Stanley 2003). The focus of the article instead lies on maintaining a safe environment: improving banks' risk mitigation techniques, increased cooperation between banks and payment reversal: attempting to shift liability of consumers is '*not good politics*'. Despite discussing the impact of consumer regulation in financial markets, most of the discussion focuses on the impact of disclosures. On the topic of credit cards, fraud is not considered to be an important regulatory aspect, compared to regulation limiting the creditors ability to change interest rates at short notice and levy fines, as banks stand to lose a large revenue stream. The reports notes that customer complaint data (from the CFPB) relatively closely reflects overall consumer satisfaction with financial institutions. The conference also debated consumers' ability to understand financial disclosures. It concludes that consumers cannot be expected to read disclosures, but if a consumer is interested they should be easily accessible and comprehensible. The participants agree that further academic research on the understanding and comprehension of disclosures should be conducted, as indeed we do in this research.

When Sullivan considers payment fraud in the US, his recommendations for reducing payment fraud are mostly technical (Sullivan 2014). In the short term the

industry should focus on hardening systems to attack, followed by improving the security of payment cards themselves in the medium term. Only in the long term does he call for standardization of payment systems (and their security). Further he argues that the existing basic tort-law principle *‘that the entity in the best position to deter check fraud will bear the losses for a check it processes’* should be expanded to the rest of the payment industry as well.

This is in stark contrast with legislation in the UK and the European Union. Before the introduction of payment cards, consumer protection in the UK was comparable to the US as consumers were not liable for forged signatures on cheques. This liability was reduced slightly with the introduction of payment cards, where the consumer was liable for the first £50 (or thereabouts, depending on the bank) of fraud (Bohm, Brown and Gladman 2000).

The introduction of ATMs caused a significant shift in consumer liability, as the bank would claim a customer was negligent or collusive if their card and PIN appeared to have been used. This was sharpened with the move to Chip and PIN as chip cards are harder to forge (Hache and Ryder 2011), and by the Payment Services Directive which supported harmonisation of regulations across the EU for members of the Single Euro Payment Area (SEPA). Suddenly if the bank deems that a consumer has been negligent in handling their PIN the consumer is fully liable. This change in customer liability has caused banks to become careless and therefore caused a huge increase in fraud in the UK (Anderson and Moore 2006), as consumers are unable to fix the banks’ numerous security issues (Anderson, Bond and Murdoch 2006; Hache and Ryder 2011).

Yet academic study of the reasons and consequences of this shift in liability is limited. In a rare publication on payment fraud from the field of criminology, Jansen and Leukfeld apply Routine Activity and Protection Motivation theory in a study with 30 phishing and malware victims (Jansen and Leukfeldt 2016). They find that to some degree everyone is susceptible to online banking fraud victimization, regardless of their knowledge and skill. They identify that victims had taken adequate steps to protect themselves yet recommend more safety training for customers as well as education about the fraud risks. However, if the risks apply to all customers who find that protecting themselves is difficult, it makes more economic sense to allocate the risks to the actor who can do most to reduce the fraud overall, and who can also implement a straightforward procedure for recovering the funds (Anderson and Moore 2006; Anderson 2007; Hache and Ryder 2011).

### 4.2.1 Legal and regulatory context

Banking contract terms are regulated everywhere. In the US, Federal Reserve regulations E and Z limit consumer liability for fraudulent debit and credit transactions at \$50 (unless for debit transactions the customer did not promptly notify the bank of a lost or stolen card, in which case the limit is \$500). Liability does not depend on whether the customer was negligent, and the only way to refuse a refund is to argue that the customer actually authorised the transaction or authorised someone else to perform it. In practice, fraud victims are generally refunded unless the bank suspects they are in cahoots with the merchant.

Practice in the EU was harmonised in 2007 by the Payment Services Directive (PSD), which states that customers are not liable for unauthorised transactions when their card is not stolen, and liability would be capped at €150 if it was. However, these limitations do not apply if the customers *‘failed with intent or gross negligence to fulfil one or more of his obligations under Article 56,’* which requires that customers comply with banking terms and conditions and, in particular, to *‘take all reasonable steps to keep its personalised security features safe.’* So, what counts as *‘reasonable’*?

European banks typically use the *‘gross negligence’* exception when they choose to deny a refund, but its definition is left to national rules and practices. As an example, banks commonly state that it is gross negligence to write down a PIN and keep it with the card. However, in practice, for a customer to be held liable, it is only necessary for the adjudicator to believe that gross negligence is, on the balance of probabilities, the most likely explanation.

The PSD requires that there be a means of adjudicating disputes without going to court. This was considered necessary because many European countries practice *‘costs shifting’*, whereby the loser of a civil case pays the winner’s costs, which could far exceed the sums in dispute.

The UK adjudicator is the Financial Ombudsman Service, from whose decisions we can see what they consider to be gross negligence. For example, in one case (Financial Ombudsman Service 2014c), a stolen debit card was used, and while the customer denied writing down the PIN, the bank records showed that the correct card and PIN had been used. The adjudicator observed that the customer had not used the card on the day (reducing the likelihood of shoulder-surfing), and concluded that the most likely explanation for the transactions was that the PIN was kept with the card. The customer was, therefore, held liable.

In one of the few UK cases to get to court (Kelman 2009), the judge also found that the most likely explanation for disputed ATM transactions was that either the

customer had made the transactions, or the customer allowed them through intent or negligence. This decision was based on expert witness testimony from the bank, stating that there had never been a breach of the Chip and PIN system at an ATM, the disputed transactions were near the customer's home, and the total disputed amount did not exceed the available balance of the account. The customer was refused a refund and ordered to pay £15,000 of the bank's costs.

Of course, there are other explanations for how the PIN could have been obtained in both of these cases; the same PIN could have been used on another card or on the customer's mobile phone, or the PIN check could have been bypassed technically (Murdoch et al. 2010). The outcome may turn on whether the adjudicator believes the bank or the complainant, which in turn may depend on their access to independent expertise.

The facts that people tend to choose PINs that are easy to guess, and that they tend to set the same PIN on multiple cards, mean that guessing a PIN is possible for about 1 in 11 stolen wallets (Bonneau, Preibusch and Anderson 2012), but a bank could argue that this is only a result of poor PIN choice and still amounts to gross negligence. Therefore, we examine guidance given to customers, to see if customers are set sufficiently clear and consistent rules with which they can reasonably be expected to comply.

### **4.3 Review of banking Terms and Conditions internationally**

#### **4.3.1 Methodology**

In the first stage of this project, we surveyed the terms and conditions of 30 banks operating in 25 countries. The study's scope included Europe (Cyprus, Denmark, Germany, Greece, Italy, Malta, and the United Kingdom), the United States, Africa (Algeria, Kenya, Nigeria, and South Africa), the Middle East (Bahrain, Egypt, Iraq, Jordan, Kuwait, Lebanon, Oman, Palestine, Qatar, Saudi Arabia, UAE and Yemen), and East Asia (Singapore). Also included in the analysis is the Code of Consumer Banking Practice for Singapore. No banks were found operating online in Libya or Syria. We selected these banks in order to get a good representation across different cultures and regions. There was some convenience sampling in the countries selected, according to the language skills available within the research team. The bank documents included in this survey, and the codes used to identify them, are outlined in [Table 4.1](#). For example, for APS Bank in Malta there were two

relevant documents: APS1 and APS2. For other banks three or more documents contained relevant information.

Major banks were selected for the study. These are not always the largest banks, as some make their terms and conditions available only to account holders. In some cases, multiple documents were reviewed, as some banks had separate terms and conditions for telephone and Internet banking, as well as credit cards, debit cards and current accounts. All the documents reviewed were downloaded from the banks' websites and were for personal (rather than business) accounts. We found that terms and conditions for accounts that adhere to Sharia law, which prohibits charging interest on loans or paying it on savings, had identical security clauses to other personal accounts at the same bank. There were no further differences found for other types of personal account customers, such as high-wealth individuals.

The terms and conditions were reviewed to identify instructions or advice on security. This included how users should handle the PINs associated with their cards, as well as credentials for telephone and Internet banking. The documents reviewed were in English, German, Italian, Arabic, and Greek. The authors include native speakers of these languages, who translated the relevant sections and coded them in accordance with the categories set out in [Table 4.2](#). To ensure consistency, we developed written instructions outlining how to select banks for review, which types of document to access, and the types of data to extract (and translate) from the terms and conditions, as set out in the coding categories.

Bank	Country	Document name	Ref	Access date
Ahli Bank	United Bahrain	Security Information	AUB1	01/09/2015
APS Bank	Malta	APS 365 Online Service - Terms and Conditions Agreement - Personal Customers	APS1	01/09/2015
APS Bank	Malta	Cards - Terms and Conditions	APS2	01/09/2015
Arab Bank	Jordan	Privacy Statement	ABJ1	02/09/2015
Arab Bank	Jordan	Ways to Bank - ATM - Security Tips	ABJ2	02/09/2015
Arab Bank	Jordan	Ways to Bank - Internet Banking Services (Arabi Online) Terms and Conditions	ABJ3	02/09/2015
Arab Bank	Yemen	Ways to Bank - ATM - Security Tips	ABY1	02/09/2015
Arab Bank	Yemen	Ways to Bank - Internet Banking Service - Terms and Conditions	ABY2	02/09/2015
Arab Banking Corp.	Algeria	Online Security	ABC1	01/09/2015
<i>ibid.</i>	Algeria	Terms and Conditions	ABC2	01/09/2015
Association of Banks	Singapore	Code of Consumer Banking Practice	ABS1	01/09/2015
<i>ibid.</i>	Singapore	Code of Practice for Banks - Credit Cards Privacy and Security - Information Security Tips	ABS2	01/09/2015
Bank Audi	Lebanon	Information Security Tips	BAL1	03/09/2015
Bank Muscat	Oman	Cards - Good Practices (Card Usage)	BMO1	03/09/2015

### 4.3: REVIEW OF BANKING TERMS AND CONDITIONS INTERNATIONALLY

Bank Muscat	Oman	Internet Banking - Security	BMO2	03/09/2015
Bank of Baghdad	Iraq	Electronic Services - Visa Card Service	BB1	01/09/2015
Bank of Cyprus	Cyprus	Cards Terms and Conditions	BCC1	07/09/2015
Bank of Palestine	Palestine	Terms and Conditions	BPP1	03/09/2015
Citibank	USA	Client Manual Consumer Accounts	CUS1	01/09/2015
Co-operative Central Bank	Cyprus	Bank Card Agreement	CCB1	07/09/2015
Commercial International Bank	Egypt	Online Security	CIB1	01/09/2015
Dachverband der Volksbanken und Raiffeisenbanken	Germany	Sonderbedingungen für das Online-Banking	DVR1	18/09/2015
<i>ibid.</i>	Germany	Sonderbedingungen für die VR-BankCard	DVR2	18/09/2015
Danske Bank	Denmark	Conditions Cheque and Cash Card Accounts	DBD1	01/09/2015
Danske Bank	Denmark	Terms and Conditions for Access Agreement - Danske eBanking Consumers	DBD2	01/09/2015
Deutsche Bank Privat- und Geschäftskunden AG	Germany	Bedingungen für den Zugang zur Deutsche Bank AG über elektronische Medien	DBG2	06/09/2015
<i>ibid.</i>	Germany	Bedingungen für Debitkarten	DBG1	06/09/2015
Deutscher Sparkassenverlag	Germany	Allgemeine Geschäftsbedingungen	DSG1	04/09/2015
<i>ibid.</i>	Germany	Bedingungen für das Online-Banking	DSG2	04/09/2015
HSBC	UK	Banking Made Easy	HUK1	01/03/2016
HSBC	UK	Current Accounts and Savings Accounts	HUK2	01/09/2015
HSBC	UK	Terms and Conditions	HUK3	01/09/2015
HSBC	UK	Personal Interest Banking Terms and Conditions	HUK3	01/09/2015
Monte dei Paschi di Siena	Italy	Terms and Conditions for "Mondo Carta" - Electronic Debit Card	MPS1	07/09/2015
<i>ibid.</i>	Italy	Terms and Conditions for "Multicanalita Integrata" - Internet and Phone Banking	MPS2	07/09/2015
National Bank of Abu Dhabi	UAE	General Terms and Conditions	NBA1	03/09/2015
National Bank of Greece	Greece	Unified Booklet of Terms for Deposits by Individuals	NBG1	07/09/2015
National Bank of Kenya	Kenya	Terms and Conditions Credit Cards	NBKe1	01/09/2015
National Bank of Kenya	Kenya	Terms and Conditions Personal Account Openings	NBKe2	01/09/2015

National Bank of Kuwait	Kuwait	ATM Safety Tips	NBKu1	02/09/2015
<i>ibid.</i>	Kuwait	Support - Security - Card Security Tips	NBKu2	02/09/2015
<i>ibid.</i>	Kuwait	Support - Security - Online Safety Tips - Prevention Checklist	NBKu3	02/09/2015
National Commercial Bank	Saudi Arabia	Consumer Protection Code	NCB1	03/09/2015
<i>ibid.</i>	Saudi Arabia	Personal Banking - AlAhli Online - Security Awareness Tips	NCB2	03/09/2015
Nedbank	South Africa	e-Banking Service Terms and Conditions	NSA1	01/09/2015
Nedbank	South Africa	Terms and Conditions of Transactional Current Accounts	NSA2	01/09/2015
OCBC	Singapore	Online Banking Security	OSi1	01/09/2015
OCBC	Singapore	Terms & Conditions - Electronic Banking Services	OSi2	01/09/2015
Qatar National Bank	Qatar	Personal Banking - Credit Cards - Credit Card Safety	QNB1	03/09/2015
Unicredit	Italy	Terms of Service - Carte di Debito Internazionali A Doppia Tecnologica - Debit Cards	UIt1	10/09/2015
Zenith Bank	Nigeria	e-Banking Service Terms and Conditions	ZBN1	01/09/2015

**Table 4.1:** Banking documents included in survey.

### 4.3.2 Results

The terms and conditions relating to PIN, telephone banking and Internet banking are considered in turn. A summary of the findings relating specifically to customer obligations to secure PINs is shown in [Table 4.3](#).

#### 4.3.2.1 PIN Writing Clauses

It is very common for banks' terms of service to provide guidelines to their customers on writing down their PIN: 26 banks out of 30 have them. The most common instruction is to keep the written PIN in a different place from the card, and not to write it on the card itself: 15 banks have such a requirement. Only six banks forbid their users from writing the PIN down anywhere. Vague statements are not uncommon: five banks instruct the customer to keep the PIN in a 'safe' place. These banks include Ahli United Bank (AUB1), Bank Audi (BAL1), Bank of Baghdad (BBI1), Nedbank (NSA1) and Zenith Bank (ZBN1). Furthermore, three banks (Arab Banking Corp. (ABC2), HSBC (HUK2), and National Bank of Kenya (NBKe1)) allow PINs to be written down in an 'obfuscated' fashion that others cannot easily recon-



Category	Description
PINWrite	References to writing down PINs
PINChange	References to changing PINs
PINReuse	References to reusing PINs, whether it be within the same or across different banks
PINAdvice	What to do with the written letter from the bank that contains the PIN
ReceiptsStatements	What to do with the receipts and statements
TelephoneWrite	References to writing down telephone banking access codes
TelephoneChange	References to reusing telephone banking access codes, whether it be within the same or across different banks
TelephoneAdvice	What to do with the written advice from the bank that contains the telephone banking access code
OnlineWrite	References to writing down online banking access codes
OnlineChange	References to changing online banking access codes
OnlineReuse	References to reusing online banking access codes, whether it be within the same or across different banks
OnlineAdvice	What to do with the written advice from the bank that contains the online banking access code
OnlineSecuritySoftware	Requirements to install and keep up to date security software
OnlineNetwork	Use of the network that the customer can access online banking from, including public access points
OnlinePassword	Requirements relating to the use of password managers or saving passwords in the browser
OnlineDevice	Requirements relating to the type or status of devices (e.g., not shared/public access, jail-broken/rooted)

**Table 4.2:** Description of coding categories used.

Bank (country)	W	C	R	A	Bank (country)	W	C	R	A
HSBC (United Kingdom)	●	●	●	●	The Association of Banks (Singapore)	●	●	○	●
OCBC (Singapore)	○	●	○	○	Nedbank (South Africa)	●	●	○	○
Zenith Bank (Nigeria)	●	●	○	○	National Bank of Kenya (Kenya)	●	○	○	○
APS Bank Limited (Malta)	●	○	○	●	Danske Bank (Denmark)	●	○	●	●
Monte dei Paschi (Italy)	○	○	○	○	Unicredit (Italy)	●	○	○	○
Sparkassen (Germany)	●	○	○	○	Deutsche Bank (Germany)	●	○	●	○
Volksbank (Germany)	●	○	○	○	Citibank (United States)	●	●	○	○
Ahli United Bank (Bahrain)	●	●	●	○	Commercial International Bank (Egypt)	●	○	○	○
Bank of Baghdad (Iraq)	●	○	○	○	Arab Bank (Jordan)	●	●	○	○
National Bank of Kuwait (Kuwait)	●	●	○	○	Arab Banking Corp. (Algeria)	●	○	●	○
Bank Audi (Lebanon)	●	●	○	○	Bank Muscat (Oman)	●	○	○	○
Bank of Palestine (Palestine)	○	●	○	○	Qatar National Bank (Qatar)	●	●	○	●
National Commercial Bank (Saudi Arabia)	●	○	○	○	National Bank of Abu Dhabi (UAE)	○	○	○	○
Arab Bank (Yemen)	●	●	○	○	National Bank of Greece (Greece)	●	●	○	○
Co-operative Central Bank (Cyprus)	●	○	○	●	Bank of Cyprus (Cyprus)	●	●	○	●

**Table 4.3: Summary of banks’ T&Cs related to PIN security** “W” indicates clauses related to writing down and storing a written PIN, “C” indicates clauses related to changing the PIN, “R” indicates clauses related to reusing it, and “A” indicates clauses related to the destruction of the letter from the bank advising of the PIN. A ● indicates that such a clause is present in the terms of service, while a ○ indicates its absence.

struct. In contrast, National Bank of Greece explicitly states that *‘the Bearer is required to: memorize the PIN, not write it down—even in an obscured fashion—on the Card or on any other document[...].’* (NBG1). There is considerable variation in how PINs may be written down, and where. For example, Arab Banking Corp. in Algeria, HSBC in the United Kingdom, and National Bank of Kenya stipulate the following:

*‘Never writing the Customer’s password or security details down in a way that someone else could easily understand, or allowing anyone to observe the Customer inputting the Customer’s password details on any electronic media.’* (ABC2)

*‘Never writing down or otherwise recording your PINs and other security details in a way that can be understood by someone else[...].’* (HUK1)

*‘If the Customer makes a written record of any PIN Code or security procedure, the Customer must make reasonable effort to disguise it and must not keep it with the card for which it is to be used.’* (NBKe1)

It is not specified whether it is the PIN that should not be understood by someone else (such as by using a code to disguise the numbers), or whether it is the connection between the PIN and the card that should not be understood. A number of other

banks are similarly vague about proximity of PIN and card. Here are more examples from Nedbank in South Africa, and Zenith Bank in Nigeria:

*‘The client must [...] ensure that any record of the PIN is kept separate from the card and in a safe place.’* (NSA2)

*‘The customer [...] undertakes [...] not to write down the Passcode, Accesscode/Password in an open place to avoid third party coming across [...]’* (ZBN1)

Bank Audi (BAL1), Bank of Baghdad in Iraq (BBI1), Bank of Cyprus (BCC1), Deutsche Bank (DBG1), Sparkassen and Volksbank in Germany (DSG1, DVR2), and UniCredit in Italy (UIt1) state that the PIN should not be stored with or on the payment card. Moreover, Bank of Cyprus (BCC1) states that the PIN should not be recorded or stored on an electronic device that allows it to be identified with the card.

Qatar National Bank provides vague advice (QNB1): it requests its customers to only memorize their PINs. On the other hand, a number of banks, including Ahli United Bank in Bahrain (AUB1), Citibank in the United States (CUS1) and National Commercial Bank in Saudi Arabia (NCB2), provide more specific advice. The following appeared under the heading ‘*Security Tips*’ of Citibank, so is perhaps not binding:

*‘Keep your Personal Identification Number (PIN), Telephone Personal Identification Code (TPIC) and other codes used to access your accounts secret. Do not tell them to anyone. Do not write them on your Citibank Banking Card or keep them in your wallet or purse [...]’* (CUS1)

The advice from the National Bank of Kenya differs by the type of account. For credit cards, there is only a requirement to keep the PIN secret:

*‘The Card member shall exercise due care to ensure the safety of the Card and the Secrecy of the PIN at all times [...]’* (NBKe1)

In contrast, the following requirements are set out for current accounts:

*‘If the Customer makes a written record of any PIN Code or security procedure, the Customer must make reasonable effort to disguise it and must not keep it with the card for which it is to be used [...]’* (NBKe3)

Arab Bank in Jordan and Yemen (ABJ2, ABY1), Bank Muscat in Oman (BMO1), APS Bank in Malta (APS2), Co-operative Central Bank of Cyprus (BCC1), National Bank of Greece (NBG1), and National Bank of Kuwait (NBKu2) forbid customers from writing down the PIN anywhere at all. For example, the following is from APS Bank:

*‘Not writing down the PIN on the Card or anywhere, or disclosing it to anyone else including the Police officers and/or the Bank’s personnel [...]’* (APS2)

Danske Bank in Denmark does not allow the PIN to be kept with the card. It does offer ‘PIN memorisers’ for recording obfuscated PINs:

*‘Do not keep your PIN with your card or write it on your card. For security reasons, you should memorise your PIN. If you are unable to do so, keep it in a safe place, preferably a PIN memoriser. PIN memorisers are available free of charge from any of our branches [...]’* (DBD1)

In Singapore, OCBC does not appear to specify how customers might record PINs (OSi2), even though its trade association has a Code of Practice which stated that customers should be told that *‘they should never write the PIN on the card [...]’* (ABS2), and the Code of Consumer Banking Practice which states that *‘you should [...] never write and/or keep record of your PIN together with your card’* (ABS1).

Finally, a few banks do not provide guidelines to customers on how PINs might be written down, such as the Bank of Palestine (BPP1), Monte dei Paschi di Siena in Italy (MPS1), and the National Bank of Abu Dhabi in the United Arab Emirates (NBA1).

#### 4.3.2.2 PIN Change Clauses

Half of the banks (15 out of 30) specifically indicate whether they allow users to change their PIN, or provide advice on how to choose a PIN. The rules varied across banks, with HSBC being concise, but general:

*‘These precautions include [...] not choosing security details that may be easy to guess [...]’* (HUK1)

One bank (Nedbank in South Africa) requires customers to change their PIN on receipt of a payment card, with no stated restrictions on PIN choice:

*‘The client shall [...] immediately change any temporary PIN and password allocated by the bank for the purpose of allowing the client to access the services for the first time [...]’* (NSA2)

One other bank (Bank of Cyprus (BCC1)) mandates customer PIN change, but also provides advice on how to select a PIN. The Ahli United Bank in Bahrain (AUB1), and OCBC in Singapore (OSi2), as well as the Association of Banks in Singapore (ABS2) set out requirements for selecting a strong PIN, telling users not to use telephone numbers, birth dates, personally identifiable information, or certain sequences of numbers as their PINs. For example:

*‘The Customer may change the Customer’s ATM-PIN from time to time. The Bank shall be entitled at the Bank’s absolute discretion to reject any number selected by the Customer as the Customer’s substitute ATM-PIN without giving any reason [...] When selecting a substitute ATM-PIN, the Customer shall refrain from selecting any series of consecutive or same or similar numbers or any series of numbers which may easily be ascertainable or identifiable with the Customer [...]’* (OSi2)

It is odd to see such a requirement in a contract, as ATM systems support a denied PIN list and the bank could simply add PINs such as 1234, 2345, ..., 9999 to this list to block them completely, along with commonly-blocked values such as 0000.

Seven other banks (Ahli United Bank (AUB1), Arab Bank (ABJ2), Bank Audi (BAL1), Bank of Palestine (BPP1), Citibank (CUS1), National Bank of Kuwait (NBKu2), and Zenith (ZBN1)) suggest their users change their PINs periodically. Finally, National Bank of Greece states that the *‘Bearer can replace [the PIN] with another number of his choice at any of the Bank’s ATMs, following the on-screen instructions’* (NBG1).

Citibank tells its customers not to choose PINs that begin with a zero:

*‘The PIN you select must consist of four numbers and cannot begin with a zero [...]’* (CUS1)

We have not been able to test whether this condition is enforced by Citibank ATMs on their own customers. We also do not know if the banks that do not set PIN-change conditions (including banks in Algeria, Cyprus (Co-operative Central Bank) Denmark, Egypt, Germany, Iraq, Italy, Kenya, Malta, Nigeria, Oman, Saudi Arabia and the UAE) offer a PIN change facility or not.

### 4.3.2.3 PIN Reuse Clauses

Even fewer banks provide advice on not reusing a PIN for multiple cards: only five out of 30. For example, HSBC states that customer precautions include *‘keeping your security details unique to your accounts with us [...]’* (HUK1). This is actually in conflict with the advice given earlier by the UK banks’ trade association which recommended customers to change all their PINs to the PIN issued for one of their cards. The UK banks allow cardholders from any bank to change their PIN at any bank-operated ATM.

Danske Bank allows customers to have a unique PIN sent to them, or to use a PIN for a personal card that has already been issued by the same bank (DBD1). The bank does not stipulate whether the PIN has to be unique to them, and in any case it does not appear to offer a PIN change facility. Arab Banking Corp. explicitly specifies that the PIN chosen has to be unique to the bank, while Ahli United Bank only states that the PIN used must be unique, under the heading *‘Security Information’* (AUB1).

### 4.3.2.4 PIN Advice Clauses

Seven banks stipulate that the original letter containing the PIN (the PIN advice letter) must be destroyed. HSBC demands this *‘immediately after receipt’*:

*‘Safely destroying any Card PIN advice we send you immediately after receipt, e.g., by shredding it [...]’* (HUK1)

In Cyprus (Co-operative Central Bank), Malta (APS Bank Limited), and Qatar (Qatar National Bank), the banks allow customers to memorise the PIN before destroying the advice:

*‘Memorise the PIN and immediately destroy the document [...]’* (CCB1)

*‘Destroying the PIN notification sent to him by the Bank immediately after memorising the PIN [...]’* (APS2)

*‘Upon receiving your credit/debit card, memorise the PIN and destroy the PIN mailer [...]’* (QNB1)

The customers of Danske Bank have no set time limit:

*‘You must also remember to destroy the letter containing your PIN.’* (DBD1)

#### 4.3.2.5 Clauses Relating to Bank Statements and Receipts

Fourteen of the 30 banks include clauses relating to bank statements and/or receipts. Overall, these banks have notably differing requirements regarding the retention of bank statements and receipts. Only HSBC in the UK and the National Bank of Kuwait insist that customers shred their bank statements if they dispose of them:

*‘Keeping card receipts and other information about your account containing personal details (such as statements) safe and disposing of them safely. People who commit fraud use many methods such as searching in dust bins to obtain this type of information. You should take simple precautions, such as shredding paper containing such information.’*

(HUK1)

*‘Save receipts: remember to take your receipts and shred them before discarding. It is best not to ask for receipts at all.’*

(NBKu1)

The advice from the National Bank of Kuwait that receipts should not be asked for differs from other banks, which require customers to retain receipts for reconciliation with bank statements (AUB1, BAL1, BMO1). The Qatar National Bank specifically states:

*‘Ensure that you received a copy of the receipt and keep it safe [...] Never throw away your transaction receipts.’*

(QNB1)

This requirement regarding the retention of records also differs across banks. The Arab Bank in Jordan requires customers to *‘ensure that your account records are properly disposed’* (ABJ1), while at the other extreme, the Arab Banking Corp. in Algeria recommends that *‘the customer prints off and keeps or electronically saves all electronic statements’* (ABC2). Three banks (Monte dei Paschi di Siena (MPS1), Unicredit (UIt1) and National Bank of Kenya (NBKe1)) provide vague statements, such as inviting their users to apply *‘common sense’* when dealing with card transactions, or using *‘due care’*.

#### 4.3.2.6 Clauses Relating to Telephone Banking Security

Clauses relating to telephone banking security are found for 13 of the 30 banks. Some are found in contracts specifically for telephone banking; others are in general terms and conditions; and yet others in a combination. Until July 2015, HSBC’s general UK contract set out requirements for safeguarding all credentials, including *‘PINs, security numbers, passwords or other details including those which allow you to use PIB [Personal Internet Banking] and TBS [Telephone Banking Service]’*

(HUK2). Further requirements for telephone banking were found in a document called ‘Banking Made Easy’. These documents were later revised; now, the terms and conditions simply require customers to follow the advice in the Banking Made Easy brochure. Some requirements for PINs also apply here: credentials cannot be written down in a way that can be understood by someone else, and they must be unique to the bank. The security code for telephone banking is a number of 6 to 10 digits created by the customer during registration, so there is no advice letter to destroy.

The OCBC (OSi2) and Monte dei Paschi di Siena (MPS2) do not specify whether telephone banking credentials may be written down, or whether they have to be unique. However, customers are permitted to change their telephone banking PIN. The OCBC specifies that:

*‘When selecting a substitute T-PIN, the Customer shall refrain from selecting any series of consecutive or same or similar numbers of any series of numbers that may easily be ascertainable or identifiable with the customer.’* (OSi2)

Citibank customers can set up a ‘Telephone Personal Identification Code (TPIC) by calling the bank’ (CUS1); the online instructions do not discuss limits on code selection, or demand that the TPIC be unique.

Many banks’ terms and conditions for PINs also apply to telephone banking, including that credentials should not be written in an ‘open place’ (ZBN1), ‘should not be kept with the card for which they are to be used’ (NBKe2), and ‘should be changed periodically and be kept confidential and private’ (ABJ3, BPP1, ABY2, NSA1).

#### **4.3.2.7 Clauses Relating to Internet Banking Security**

As with telephone banking, some banks have specific contracts for Internet banking, while others include this in general contracts. Some go still further to impose conditions on the security of the network, the security of the device including the use of security software, and the use of online password managers or browsers to store credentials.

The most onerous conditions are set out by HSBC in the UK. Under its Personal Internet Banking Terms and Conditions (HUK3), credentials for Internet banking must not be written down in a way that can be understood by someone else, they cannot be easy to guess, and they have to be unique to the bank. The customer must always access Internet banking by typing the address into the web browser and use antivirus, antispyware and a personal firewall. If accessing Internet banking from



a computer connected to a LAN or a public Internet access device or access point, they must first ensure that nobody else can observe, copy or access their account. They cannot use any third-party software, such as browsers or password managers, to record passwords or other security details. Finally, all security measures recommended by the manufacturer of the device being used to access Internet banking must be followed, such as using a PIN or biometric to lock a mobile device.

The OCBC, in Singapore, insists that the card and PIN must not be kept together, yet elsewhere that PINs must be memorised and not recorded anywhere. Customers were advised not to repeat any digits in the 6-digit PIN more than once, that it should not be based on the User ID, telephone number, birthday or other personal information, that it should not be used for different websites, applications or services, and that it should be changed ‘regularly’. Customers of the Singapore bank also have to install antivirus, antispymware and firewalls, and ensure they were updated and patched. File and printer sharing also have to be disabled, and customers cannot use public or Internet cafe computers. Browsers cannot be used to store credentials. What’s more:

- 10. Do not install software or run programs of unknown origin [...]*
- 14. Do not use a computer or device which cannot be trusted [...]*
- 16. You are advised not to access Online Banking using jailbroken or rooted mobile devices (i.e. the phone Operating System has been tampered with), as it poses potential risk of malicious software infection.’*

(OSi1)

The other banks reviewed do not impose such aggressive restrictions. Clauses specific to online banking include: using a firewall, antivirus and/or antispymware software (AUB1, APS1, ABJ1, ABY2, ABC1, BAL1, AMO2, CIB1, DBR1, DBG2, DSG2, NBA1, NBG1, NBKu3, NCB2); using a modern browser (ABC1, MPS2); patching the browser and/or operating system (AUB1, ABC1, BMO2, CIB1, DVR1, DBG2, DSG2, NBA1, NBG1, NCB2); not saving passwords in password managers or browsers (ABC2, BMO2, CIB1, NBA1, NBKu3, NCB2); not using public access computers (AUB1, ABJ3, ABY2, BMO2, CIB1, NBA1, NCB1, NCB2); encrypting wireless networks (CIB1); clearing the cache after each banking session (BMO2); using a password to access the computer (NCB2); and disabling file and printer sharing capabilities (NCB2). The National Bank of Kuwait and the Commercial International Bank in Egypt refers to particular firewalls and antivirus programs:

*‘Common commercial examples include Zone Labs, symantec.com and Computer Associates. The leading free firewall is “Zone Alarm” from Zone Labs and there are many others to choose from. Zone Alarm*

*is now used on over 20,000,000 PCs and has been awarded the PC World 2003 “World Class Award” for Best Firewall.’* (NBKu3)

*‘There are many effective programs to choose from, but the most common commercial products include McAfee, Symantec (Norton) and Sophos. It is also possible to obtain free anti-virus protection. A search for “free anti-virus” on Google will provide a list of the most popular.’* (CIB1)

Danske Bank stipulates that customers should not leave the mobile phone on which they receive codes and their payment card number with others, including members of their household (DBD2).

### 4.3.3 Discussion

This review of bank terms and conditions demonstrates that banks take a variety of approaches to the security advice they offer to, and the demands they make of, customers. The approach varies not just between jurisdictions but between banks in each jurisdiction. Advice on writing down PINs ranges from a strict ‘no’ to ‘yes, but’, with requirements on obscuring the PIN and safekeeping. About half the banks allow users to change PINs; but there is a range of advice on PIN choice. The PIN advice letter may have to be destroyed at once, or eventually, or not at all; and a similar range of advice is given for bank statements and receipts. The clauses regarding the safekeeping of authenticating information for telephone banking and internet banking are no different, except in that different banks take the different extreme positions. Internet banking security is just as diverse but much more complex; there are many more kinds of advice that banks can and do give about preventing malware infection of devices used for online banking, many of which are outdated and difficult to follow.

While some of the advice is drafted so as to be helpful, many of the instructions are demands. It is not clear from the prevent survey which of these demands are also designed to minimise the risk to the customer, and which are there to minimise the risk to the bank, by enabling it to ask a whole series of hard questions of customers who complain of fraud, and reject many claims on the basis of non-compliance. Such behaviour will be discussed further in [Section 4.5](#) below. Surveying how banks in different countries actually treat fraud victims and how this relates to local banking regulation would be a fascinating research project, but a very much larger one than the work reported here.

#### 4.3.4 Limitations

There are number of limitations to our methodology that should be mentioned. Primarily, our scope is limited to languages that our diverse range of collaborators speak. We did use search engines and site maps when searching for T&C documents, so while care was taken to retrieve the terms and conditions documents of the banks studied, there were a number of banks which did not seem to publicly list their T&Cs.

Further, banks regularly update their T&Cs. Our research here presents a snapshot. We have made our dataset publicly available (see [Section 4.6.1](#)), as outdated T&C can be hard to retrieve.

### 4.4 Survey of payment card PIN usage

In the previous section we have established that there is a diverse range of demands on customers security behaviours. Advice on PIN and password safety varies and is sometimes contradictory. To see if the expectations are realistic, I conducted an on-line questionnaire study of how people use payment cards, and, in particular, how many PINs they have, and how they are remembered. We also investigated their behaviour towards storing, resetting and sharing PINs. We exclusively investigated 4-, 5- and 6-digit PINs. For each individual PIN we asked the participants for its frequency of use (on a 7-point scale between every day and never), and the origin of the PIN (with options: *'self-chosen'*, *'bank chosen and allowed to change'*, and *'bank chosen and not allowed to change'*). For each PIN length, we enquired if, and if so, where participants had written down any of these PINs, if these PINs were used on other devices, and if so, on which ones, and if any of these PINs have been shared with other people, and if so, what the relationship between the sharers is. We also enquired for each PIN length if any of the PINs had been forgotten, and how they had been recovered if applicable, as well as if the participants have a pattern or strategy for creating PINs.

Questions that required categorical responses by the participants had a set of predefined choices as well as a free-text response field. The predefined choices were sourced from a small qualitative preliminary study. In nearly all cases the participants did not make use of the free text response field.

#### 4.4.1 Questionnaire setup

The questionnaire was set up using LimeSurvey<sup>1</sup>, and the participants were recruited using Prolific Academic<sup>2</sup>. We restricted submissions to British residents aged 18 or over. Participants were paid £1.50 for successfully completing the questionnaire. The questionnaire took five minutes on average to complete. We received 241 (out of 250) valid responses, and verified that the IP address used was from the UK in all but 5 cases<sup>3</sup>. 9 submissions were removed for administrative reasons.

#### 4.4.2 Results

Of the participants, 61% are female and 39% are male. The age of the participants spans 18 to 71 years, with a mean of 31.2. Our participants are better educated than average bank customers: 38% have at least an undergraduate degree (BSc, BA or similar), while a further 17% have postgraduate education. 30% did not attend higher education, and a third of these (10%) have a General Certificate of Secondary Education (GCSE) as their highest qualification. 49% of the participants are employed; a further 13% are self-employed; 24% of participants are students; only 13% are unemployed.

	Number of PINs										
	0	1	2	3	4	5	6	7	8	9	mean
4 digits	1	88	65	41	31	8	5	1	1	0	2.28
5 digits	233	5	3	0	0	0	0	0	0	0	0.05
6 digits	228	8	4	1	0	0	0	0	0	0	0.08

**Table 4.4:** Distribution of the number of 4-, 5- and 6-digit PINs the participants have.

The participants report having 1 to 9 payment cards (mean = 2.53). This contrasts with the number of 4-digit, 5-digit, and 6-digit PINs each participant has in [Table 4.4](#). The vast majority of customers have only 4-digit PINs, but the mean number of PINs is 2.28, statistically significantly lower than the mean number of cards per customer (dependent t-test,  $t = -4.38$ ,  $p < 0.0001$ ). This suggests that some participants have reconciled their PINs, but the majority use different PINs for different cards.

[Table 4.5](#) analyzes how often participants use their PINs. No participant had more than eight 4-digit PINs, or more than two 5-digit ones or three 6-digit ones.

<sup>1</sup>[www.limesurvey.org](http://www.limesurvey.org)

<sup>2</sup>[www.prolific.ac](http://www.prolific.ac)

<sup>3</sup>IP address geo-location has a non-trivial error rate, but this still confirms that our sample is predominantly from the UK, as intended.

#### 4.4: SURVEY OF PAYMENT CARD PIN USAGE

	4-digit PINs								Sum	5-digit PINs			6-digit PINs			
	#1	#2	#3	#4	#5	#6	#7	#8		#1	#2	Sum	#1	#2	#3	Sum
Every day	34	0	0	1	0	0	0	0	35	0	0	0	1	1	0	2
Several times a week	117	30	3	3	0	0	0	1	154	1	0	1	5	2	1	8
Once per week	59	35	12	3	0	0	0	0	109	2	1	3	0	0	0	0
Once per month	21	37	24	8	3	0	0	0	93	4	2	6	3	2	0	5
Several times a year	6	24	24	12	2	2	1	0	71	1	0	1	3	0	0	3
Once a year or less	1	14	10	10	4	1	0	0	40	0	0	0	1	0	0	1
Never	2	12	14	9	6	4	1	0	48	0	0	0	0	0	0	0

**Table 4.5:** Distribution of how frequently our participants use each of their PINs.

	#1	#2	#3	#4	#5	#6	#7	#8	Sum
I chose it myself	75	56	28	15	3	3	1	0	181
Assigned to me, I decided not to change it	161	94	56	31	11	3	1	0	357
Assigned it to me, I am not allowed to change it	4	2	3	0	1	1	0	1	12

**Table 4.6:** Source of participants' 4-digit PINs

We see at once that as the number of PINs increases, their usage drops. Only one participant uses more than one unique PIN on a daily basis. About half (48%) of the PINs are used at most once a month, and PIN #4 is used on average around twice a year. This supports the bank industry ‘folk wisdom’ that if you want customers to use cards other than their main card you must let them change their PINs.

Table 4.6 documents PIN change, and we see that two-thirds of PINs are left as their default. Interestingly, there is no correlation between frequency of PIN use (Table 4.5) and PIN origin (Table 4.6). Virtually all participants are allowed to change their PINs in the UK. The details for 5- and 6-digit PINs have been omitted here for brevity. We also investigated the reasons for PIN change. Of the participants that set their own PIN, 61% reported changing their PIN on first receipt, a further 23% stated they changed their PIN because they felt it was compromised, but only 6% claimed to change their PIN on a regular basis.

Our participants keep their PINs for a long time: only 13% of PINs were changed in the last year, with over 39% having been in use for over 5 years.

## 4.4.2.1 Remembering PINs

A quarter of the participants reported forgetting a 4-digit PIN at least once. Of these, 48% remembered or retrieved their PIN themselves, 24% were issued with a new PIN, and 15% used the bank’s services to retrieve their PIN. Finally, 10% did not bother retrieving the forgotten PIN; half of these said they transferred their money to a different account.

	4-digit	5-digit	6-digit	Sum
On the card	1%	0%	0%	1%
I kept the original PIN slip	0%	0%	0%	0%
On paper – kept in desk	16%	25%	0%	17%
On paper – kept in wallet	10%	0%	0%	10%
In a notebook/diary/planner, etc.	41%	25%	0%	40%
File on computer	10%	25%	0%	11%
File on phone	42%	25%	0%	41%

**Table 4.7: Location of written down PINs by participants** 79 (32.9%), 4 (50.0%), and 0 (0.0%) wrote down their 4-, 5-, and 6-digit PINs, respectively.

As many banks insist that PINs must not be written down, we decided to investigate this in reality. [Table 4.7](#) describes our participants’ strategies towards writing down PINs. Not a single participant kept the original PIN mailer. The prevailing method of PIN storage is on mobile phones, typically disguised as a phone number. 13% of the participants use a mnemonic for 4-digit PINs, the most common technique being the derivation of the PIN from a specific date. This was also reported by [Bonneau, Preibusch and Anderson \(2012\)](#).

	4-digit	5-digit	6-digit	Sum
Unlocking mobile phone	49%	0%	0%	48%
Burglar alarm	2%	0%	0%	2%
Voicemail	15%	0%	0%	14%
SIM card unlock	7%	0%	0%	7%
Unlocking computer	5%	0%	100%	7%
On-line Banking	25%	0%	0%	24%

**Table 4.8: A variety of locations where participants’ PINs are re-used** 55 (22.9%), 0 (0.0%), and 1 (7.7%) re-used their 4-, 5-, and 6-digit PINs, respectively.

#### 4.4.2.2 PIN re-use

16% of our participants stated they use the same PIN on many payment cards; when a PIN was re-used, it was used on 2.8 payment cards on average, with the maximum being 9 cards! PINs were also used in a variety of other locations (Table 4.8): 22.9% of participants are re-using their 4-digit payment card PINs, half of whom use a payment card PIN to unlock their mobile phone.

	4-digit	5-digit	6-digit	Sum
Stranger	0%	0%	0%	0%
Family member	37%	0%	100%	37%
Flatmate (if accommodation shared)	3%	0%	0%	3%
Spouse/partner	75%	100%	0%	74%
Casual acquaintance	1%	0%	0%	1%
Close friend	14%	0%	0%	13%

**Table 4.9: Sharing of PINs by participants** 102 (42.5%), 2 (25.0%), and 1 (7.7%) shared their 4-, 5-, and 6-digit PINs, respectively.

#### 4.4.2.3 PIN sharing

Finally, an impressive 42.5% of our participants share their PINs, in many cases with more than one person (Table 4.9). Sharing predominantly occurs with a spouse or partner (32% of participants) or other family members (16%), but also, in some cases, close friends (6%).

#### 4.4.3 Discussion

In general, it is difficult for customers to be certain whether they are complying with bank rules, as these rules lack detail and can even be contradicted. For example, HSBC prohibits PIN re-use, whereas the UK bank trade body recommends this (Bowerman 2007). Vague and contradictory guidance puts customers in a weak position should a bank claim that a disputed transaction must have been caused by a failure to comply with its rules.

Our survey of PIN use confirms the practical difficulty of remembering PINs. Customers are commonly asked to remember four or more PINs, some of which they only use every month at most. The combined effect of forgetting over time (Squire 1989), as well as interference between the different PINs remembered (Anderson and Neely 1996), makes unaided recall of these PINs a difficult task. In one study after 3 weeks, the majority of participants were unable to remember a PIN, and even

after 1 week 45% had forgotten it (Renaud and Ramsay 2014). In current usage scenarios, customers can only cope by re-using PINs or writing them down.

The 4-digit PIN system worked adequately in the environment for which it was originally designed: a single regularly-used ATM card. Today's usage scenarios are different, but the mechanism, and terms and conditions, remain unchanged. Not considering the usability implications pushes customers towards insecure PIN practices banned by the banks' contracts. Each PIN re-use allows a thief another 6 guesses, and mobile phone touchscreens can give away PIN digits directly. Not all customers can be expected to disguise PINs securely, but remembering an infrequently-used PIN is impractical without some kind of assistance. Customers who do not comply with their banking contract are then blamed for security failures that are actually caused by a system that is not fit for purpose.

### 4.5 Survey of understanding and interpretation of banking T&Cs

The third contribution of this research is a cross-cultural study of the understanding and interpretation of banking terms and conditions (T&Cs). As we have described in the previous section, there are significant differences in the legal setting of banking between countries, as well as between banks within the same country. However, these may appear rather theoretical. In order to distil out the practical effects of the banks' contracts, we conduct a survey with participants from Germany, the United Kingdom and the United States. Consumer protection for fraudulent transactions in Germany and the UK is governed by the same law, the EU Payment Services Directive (PSD)<sup>4</sup> (soon to be replaced by PSD2<sup>5</sup>), whereas US disputes are governed by the more consumer-friendly federal regulations E<sup>6</sup> and Z<sup>7</sup>. The PSD allows banks to refuse refunding a customer if the most likely explanation for the fraud is considered to be that the customer was grossly negligent in complying with bank security rules. Regulations E and Z require that customers be refunded in almost

---

<sup>4</sup>Directive 2007/64/EC of the European Parliament and of the Council of 13 November 2007 on payment services in the internal market amending Directives 97/7/EC, 2002/65/EC, 2005/60/EC and 2006/48/EC and repealing Directive 97/5/EC (Text with EEA relevance), available at <http://eur-lex.europa.eu/legal-content/EN/ALL/?uri=CELEX:32007L0064>

<sup>5</sup>Directive (EU) 2015/2366 of the European Parliament and of the Council of 25 November 2015 on payment services in the internal market, amending Directives 2002/65/EC, 2009/110/EC and 2013/36/EU and Regulation (EU) No 1093/2010, and repealing Directive 2007/64/EC (Text with EEA relevance)

<sup>6</sup><https://www.federalreserve.gov/bankinfo/reg/gecg.htm>

<sup>7</sup><https://www.federalreserve.gov/bankinfo/reg/zcg.htm>



all circumstances, and demonstrating gross negligence is not sufficient to refuse a refund. The aims of this survey are threefold:

1. identify the perceptions and prejudice of participants towards banking terms and conditions;
2. measure the ability of our participants to understand the banks' terms and conditions and act on them;
3. on a country specific basis and as a cross-cultural study.

### 4.5.1 Related literature

While there are to our knowledge no cross-cultural examinations of the understanding of Terms and Conditions, many methodologies of cross-cultural studies have been explored. Similarly, text comprehension is an established research branch.

#### 4.5.1.1 Cross cultural study design

The field of psychology has devoted much research into conducting valid cross-cultural studies. Jones and Kay lay out a number of challenges that cross-cultural research faces (Jones and Kay 1992). In this study we are interested in the comparative use of the scores deduced from the responses from each of the countries, as we intend to perform analysis on the differences between the groups. This requires us to aim for a construct-referenced meaning across the different languages: the study's *aim* needs to be translated to the cultures (Hui and Triandis 1985). This is not a purely translational issue: social norms and concepts may well be different. Hence in our study we have opted for a symmetric translation: we have adjusted cultural symbols (as far as these are present in contractual terms) to the best of the translator's knowledge of the target cultures.

The difficulty of conducting studies is not just limited to the textual content. The steps on Likert scales may have different meanings in different languages (Heine et al. 2002). We have taken great care that our study is accurately represented in the three cultures studied, but a full sociological validity analysis is outside the scope of this research.

#### 4.5.1.2 Reading comprehension

There is an existing body of research on the comprehension of legal texts. The initial research seems to have been carried out by Masson and Waldron, who simplified Terms and Conditions in three steps. Each simplification increases comprehension, but absolute comprehension values were still very low (Masson and Waldron 1994).

Further studies on the topic point out individuals don't read contracts (Wogalter et al. 1999) before signing them, or don't read Software Licensing Agreements before clicking OK (Wogalter and Hayes 2014). Shorter End User License Agreements (EULAs) may actually decrease number of software installations (Good et al. 2007), as can politely asking for user permission (Böhme and Köpsell 2010). Paraphrasing may well aid readability (Waddell, Auriemma and Sundar 2016), but just as in (Masson and Waldron 1994), absolute values of readability remain low.

These empirical findings are supported by one analytical study: Prichard and Hayden analyse the EULAs of freeware software using a number of readability metrics (Prichard and Hayden 2008). They (perhaps unsurprisingly) support the findings from the literature: the vast majority of EULAs were very difficult to read and very long.

#### 4.5.2 Survey design

The survey is divided into four stages. We begin with some demographics on the participants as well as some statistics on the payment methods our participants use, such as the number of different payment cards and bank accounts the participants have and the countries these are issued in. We also ask about previous experiences of fraudulent transactions in order to identify prior experiences with the T&Cs. This is followed by two scenarios on the conflicts between the T&Cs and customer. We ask the participants what they expect the resolution of the conflict to be. We elicit the responses to the scenarios twice, before being aware of the T&Cs and again afterwards. The scenarios are sourced from the UK Financial Ombudsman newsletter (Financial Ombudsman Service 2014a, 2014b, 2014c). The scenario texts presented below have not been changed in meaning from the Financial Ombudsman newsletter. The Financial Ombudsman Service (FOS) is an arbitration service that was set up by the UK banks as an alternative to using the court system to resolve disputes with customers, and that now adjudicates according to their interpretation of the requirements set out in the PSD. Arbitrations are binding on the bank, and while a customer may instead take a case to the courts, the prohibitive costs involved make this rare, so the FOS interpretation of the PSD is dominant in the UK. Its quarterly newsletter publishes examples of its common dispute resolutions.

After asking a number of questions regarding the scenarios, we introduce a set of terms and conditions on payment safety and fraud (as discussed in Section 4.3.3). We next ask the participant a number of questions to gauge their understanding of these T&Cs on a following page, without allowing them to page back to access the terms.

We then reintroduce the two scenarios, this time giving them access to the text of terms and conditions, and again enquire about their expectation of the outcome. The full survey text can be found online, see [Section 4.6.1](#).

Most of the responses in the survey are collected using free-text responses. There are several reasons for choosing this method. As we are interested in the perceptions of the participants and their understanding of the contract terms, we wanted to remove any form of prompt in order to get unbiased responses. These free-text responses are then manually grouped using Thematic Analysis (Fereday and Muir-Cochrane 2006). The raw counts are normalised, in most cases by the number of participants per country. As each participant may have mentioned multiple themes, each theme may range between 0% and 100%.

The participants for this study were recruited using Prolific Academic. The survey took on average 18 minutes to fill out, and we paid each participant £2.50. We trialled the study on German, British and American native speakers and ran an initial online pilot that helped us resolve some minor ambiguities.

Conducting the survey in two languages across three countries posed several challenges. Firstly, the financial legislation is very different between the EU and the US (the US being significantly more consumer-friendly). This had a direct impact on the responses, but we were nevertheless able to measure the impact of the treatment of the terms and conditions on the two scenarios. Second, there are significant cultural differences regarding privacy and data protection between the three countries.

The survey and the two scenarios were translated into German by a native speaker, and checked by a second native speaker in order to ensure that the intent of the questions was preserved as closely as possible. Minor changes were also made between the British English and American English version in order to aid comprehension.

### 4.5.3 The scenarios

Our two scenarios were presented in a random order to the participants. The order they were shown in did not lead to any statistically significant variations in answers. The scenarios shown below are those shown to the participants in the UK.

#### 4.5.3.1 Scenario 1: card loss

The first scenario is based on a typical story of theft (Financial Ombudsman Service 2014c). The scenario reads as follows:

*‘Miss K travels to work on the Tube. When leaving the Tube at the destination station, Miss K notices that her purse is missing. In the Tube station is a police office, where she reports her purse as stolen. When she gets to work, she phones her bank to cancel her debit card. But, by this time, the thief has made several large cash withdrawals using the card.’*

In the original article, the Financial Ombudsman decided that the most likely reason the thief could withdraw cash is that Miss K stored her PIN with the card. The Ombudsman concluded that Miss K had likely been grossly negligent and is denied a refund. We do not tell the participants this outcome.

#### 4.5.3.2 Scenario 2: phone scam

The second scenario is based on a combination of Ombudsman News stories (Financial Ombudsman Service 2014a, 2014b). The scenario reads as follows:

*‘Mr L received a phone call from his bank. The person he spoke to said there had been some “suspicious activity” on his account, and asked him if he had made certain purchases. When Mr L said he hadn’t, the person on the phone said that he should call a different department at the bank straight away to sort the problem. Mr L called the number on the back of his debit card. The person he spoke to asked him some security questions and then confirmed that suspicious activity had taken place. They said that Mr L should immediately transfer all the money from his account to a different account, and he gave him the details of that account over the phone. Mr L transferred the money straight away.*

*When Mr L told his partner what had happened, she was worried. She suggested he call his bank to check he’d done the right thing. It turned out that Mr L had been the victim of a scam. The fraudster had put a technical fix in place so that when Mr L ended the first call and rang the number for his bank, he’d actually just reconnected with the fraudster.’*

In this scenario, the Ombudsman ruled that Mr L should not be reimbursed, as he was deemed to have authorised the transaction. The Financial Ombudsman chose the wording of the first line of the scenario to intentionally highlight that the customer is unable to verify the identity of the bank. There is a large number of similar cases in the UK (although this type of fraud is less common in Germany), which all vary slightly on the exact manner the fraudulent transaction is processed. In some cases, the Ombudsman decides in favour of the customer; in many others, she does not.

#### 4.5.4 The Terms and Conditions

It is infeasible to have our participants work through an entire document of Terms and Conditions as part of a study, as these documents range from 20 to 40 pages. In order to get a realistic assessment of the ability of our participants, we specifically presented only the passages most relevant to the study, but left whole paragraphs intact. For the UK, we chose HSBC’s ‘General Terms and Conditions (HUK2)’, and in particular Section 9. ‘Important Security Information’, and Section 27.5 ‘Liability for Unauthorised Transactions’. As discussed previously, HSBC’s T&Cs are representative of T&Cs in the UK.

For the American participants, the survey focused on Citibank’s ‘Client Manual Consumer Accounts’ (CUS1). In particular, we chose the sections on ‘Lost or Stolen Banking Cards or Other Access Devices’ and ‘Unauthorised Electronic Transactions and Security Tips’. Again, our choice of the bank followed from our analysis in [Section 4.3](#).


Following the same argument, we chose the Terms and Conditions for Debit Cards of Deutsche Bank (DBG1) for Germany. Here, we focused on Section 6. ‘Geheimhaltung der persönlichen Geheimzahl (PIN)’ (Keeping your PIN secret), Section 12. ‘Erstattungs- und Schadensersatzansprüche des Kontoinhabers’ (Reimbursements and claims for damages of the account owner), and Section 13. ‘Haftung des Kontoinhabers für nicht autorisierte Kartenverfügungen’ (Liability of the account owner for unauthorised card charges).

#### 4.5.5 Demographics

We recruited 151 participants in total: 41, 56 and 54 participants from the DE, UK and US respectively. An overview of the age and gender of all recruited participants can be found in [Figure 4.1](#) and [Table 4.10](#). There are some surprising differences in these demographic distributions between the three countries, considering that all participants were sourced from the same platform. There is a strong gender bias of around 3:1 in Germany and the US. The participants in the UK are, however, gender-balanced.

Gender	DE	UK	US
Female	24%	52%	27%
Male	73%	48%	71%
Other	2%	0%	2%

**Table 4.10:** Gender of our participants



**Figure 4.1:** Histogram of our participants' age

Figure 4.1 highlights the age distributions between the participants from the three countries. We checked the participants' location by geo-locating their IP address used to access the survey. IP geo-location is far from accurate; however, all but 3 participants' IP addresses matched their declared country. We decided to include the answers from these outliers, as the answers were well-written and showed no other anomalies. The mean ages across the three countries are 27.0, 33.7 and 30.4 years for the DE, UK and US respectively, with standard deviations of 6.6, 12.7 and 9.8 years respectively. In general, Prolific Academic seems to have the most representative demographics for the UK. The age demographics for these countries are reasonably consistent with the participant pool available through Prolific Academic Prolific (2018), so there does not seem to be a selection bias for participants based on these demographics.

Employment Status	DE	UK	US
Employed	22%	48%	57%
Student	63%	30%	14%
Unemployed	2%	4%	12%
Self-employed	7%	13%	16%
Retired	2%	6%	0%
Prefer not to say	2%	0%	0%

**Table 4.11:** Employment demographics of our participants

There are distinct differences in employment across the US, DE and UK, as can be seen in Table 4.11. Almost all our participants claim to be native speakers in our study; 100%, 92% and 92% for the DE, UK and US respectively. There is an above-average distribution of educational statistics, which shows over 50% of our participants from each of the countries have finished at least a bachelor's degree or equivalent (see Table 4.12). The translation of education levels is not


Highest Qualification	DE	UK	US
GCSE Level education (e.g., GCSE, O-Levels or Standards) or lower	7%	15%	0%
A-Level education (e.g., A, AS, S-Levels, Highers)	24%	11%	12%
Some undergraduate education (e.g., No completed degree)	10%	19%	18%
Degree/Graduate education (e.g., BSc, BA)	32%	35%	43%
Postgraduate education (e.g., MSc, MA, MBA, PhD)	22%	19%	16%
Vocational education (e.g., NVQ, HNC, HND)	5%	0%	5%
Other	0%	2%	4%

**Table 4.12:** Educational demographics of our participants


straightforward, which may explain the 0% value for GCSE level education in the US. Two participants revealed that they had learning disabilities.

It is obvious that our participants' demographics could be better aligned for a crosscultural study. Unfortunately Prolific Academic does not offer the functionality to sample participants to a specific demographic distribution. However, while there are strong differences in the employment demographics (Table 4.11) and gender, participant's age and educational demographics (Figure 4.1 and Tables 4.10 and 4.12) are fortunately similar.

#### 4.5.6 Payment demographics

**Figure 4.2:** Participants' number of payment cards

In order to meaningfully compare the responses to our questions, we have to check that the participants have similar levels of financial development. One measure is the number of bank account and payment cards. Figure 4.2 shows 3 histograms for the number of payment cards our participants have in the DE, UK and US respectively. The means are here 2.0, 2.7 and 3.1 respectively. Similarly, Figure 4.3 displays the distribution of bank accounts of our participants with means 2.0, 2.5 and 1.8 respectively. While many credit cards are prevalent in the US, our participants there also have the smallest number of bank accounts.



**Figure 4.3:** Participants’ number of bank accounts

Frequency	DE	UK	US
Every day	0%	19%	20%
Several times a week	63%	65%	55%
Once per week	22%	13%	20%
Once per month	5%	2%	4%
Several times per year	7%	2%	0%
Once per year or less	0%	0%	2%
Never	2%	0%	0%

**Table 4.13:** Frequency of use of any of our participants’ payment cards

We also investigate the frequency of payment card use. Here, the UK participants use their cards the most, followed by the Americans. No participant in Germany uses a payment card on a daily basis (Table 4.13), yet payment card penetration rates are still high. Virtually, no participant manages a week on average without using a card.

#### 4.5.6.1 Fraud experience

Frequency	DE	UK	US
No	88%	72%	66%
Yes	12%	28%	34%

**Table 4.14: Frequency of Fraud** Have you ever experienced fraudulent transactions or incidents on any of your payment cards or bank accounts?

We hypothesise that people who have been a victim of fraud previously are more likely to pay attention to the details of payment contracts. We ask the participants if they have been fraud victims, and to explain the experience to us if they have. In



Table 4.14, we list our participants’ frequency of fraud experience. In order to get as complete description as we could, we solicited free-text responses.

Code	DE	UK	US
Fraud identified at a later stage	28.6%	55.0%	60.7%
Transaction before card blocked	0.0%	0.0%	3.6%
Transaction after card blocked	0.0%	5.0%	0.0%
Transaction blocked by bank	42.9%	30.0%	21.4%
Other/No idea where fraud occurred	42.9%	30.0%	42.9%
Offline transaction	14.3%	15.0%	17.9%
Online transaction	14.3%	40.0%	14.3%
Cash withdrawal	0.0%	0.0%	7.1%
Card stolen	0.0%	5.0%	7.1%
Online account hacked	0.0%	5.0%	0.0%
New card	28.6%	30.0%	32.1%
Full refund	14.3%	80.0%	82.1%

**Table 4.15: Thematic analysis of the description of fraud experienced by participants** The first four codes describe the identification of fraud, the next six codes describe the type of fraud, and the last two describe the follow-up actions that happened. Each percentage represents the proportion of statements from participants from that country that have been annotated with that code.

We analysed these responses using Thematic Analysis, and the results can be seen in Table 4.15. The table is divided into three sections: fraud identification, type of fraud, and resolution. As these were manually annotated free-text responses, the absolute percentages are approximate, but the relative differences are worth noting. There is a clear trend in the stage where fraud is identified: in Germany, more fraud is identified automatically than is noticed by the customers. This is reversed for the UK and the US, where almost two-thirds of fraud is identified by the customer. For American customers, this may be an annoyance, but a minor one: Federal Regulations E & Z ensure that the customer will get his money back. In the UK, this may be a greater worry as the refund is dependent on whether the bank considers you to have been ‘*grossly negligent*’.

#### 4.5.7 Scenario overview

In the following Sections 4.5.8 and 4.5.9, the participants consider the two scenarios in two different combinations: once before seeing the relevant terms and conditions, and once afterwards. Each time they are asked if they think the protagonist should be reimbursed by the bank and why. The results of the binary question can be found in Table 4.16. We find that in all but one case, the participants are more likely to

Question	DE	UK	US
Scenario 1: Card Loss	41.5%	81.5%	76.8%
Scenario 1: Card Loss after T&Cs	70.7%	66.7%	96.4%
Scenario 2: Phishing	31.7%	37.0%	35.7%
Scenario 2: Phishing after T&Cs	43.9%	46.3%	42.9%

**Table 4.16: Percentage of participants that say that the money should be returned in each of the scenarios** McNemar’s test is significant with  $p < 0.05$  for both Scenario 1: Card Loss and Scenario 2: Phishing.

have the protagonist reimbursed after reading the terms and conditions. This is statistically significant with  $p < 0.05$  for both scenarios using the McNemar’s test for binary variables. We will now consider each of these four conditions in isolation, and analyse the qualitative responses.

#### 4.5.8 Scenario 1: card loss

For each of the two settings, there are two sets of answers to consider: those that argue for the reimbursement of the protagonist, and those against it.

##### 4.5.8.1 Prior to revealing Terms and Conditions

Code	DE	UK	US
Banks have good security that should have prevented fraud	0.0%	4.5%	0.0%
Depending on the T&C of the bank	0.0%	0.0%	9.3%
Insurance will compensate her	0.0%	2.3%	14.0%
People are protected from fraud by the bank	35.3%	38.6%	48.8%
She did not authorise the transaction	17.6%	6.8%	2.3%
The theft was reported swiftly	52.9%	50.0%	41.9%
Yes, because the bank can prove it wasn’t her, due to CCTV at ATM	5.9%	11.4%	7.0%

**Table 4.17:** Thematic analysis of the answers in support of reimbursement in Scenario 1: Card Loss.

Tables 4.17 and 4.18 show the results of the *Card Loss* scenario before revealing the terms and conditions. The respondents who supported reimbursement gave a wide range of reasons (Table 4.17). The most recurring reasons across the German, UK and US surveys are: (1) the theft was reported immediately, cited by 52.9%, 50.0% and 41.9% of respondents respectively, and (2) banks are expected to protect their customers from fraud, with 35.3%, 38.6% and 48.8%. Additionally, some of the

UK respondents (4.5%) were more specific, and said that good security measures are deployed by banks to defend against fraud. 17.6% of the German respondents said that Miss K did not authorise the transaction and, hence, she should be reimbursed; only 6.8% and 2.3% mentioned the same reason in the UK and US surveys. Another interesting reason for reimbursing Miss K is that it can be easily proven that she did not make the transaction because CCTV cameras are widely deployed at ATMs; this reason was mentioned in all three surveys. Only 2.3% of the UK respondents believe that the insurance company is responsible for compensating Miss K, whereas 14.0% provided the same reason in the US survey. Interestingly, only US participants, with about 9%, mentioned that reimbursement depends on Miss K's bank terms and conditions.

Code	DE	UK	US
Common perception that the customer loses	4.2%	20.0%	23.1%
Debit, as opposed to credit, cards do not have fraud protection	0.0%	10.0%	23.1%
Don't know/unsure	4.2%	0.0%	0.0%
Her mistake	25.0%	20.0%	30.8%
Her purse is not insured, thief must be caught	0.0%	20.0%	0.0%
Money cannot be retrieved once it leaves someone's account	4.2%	20.0%	23.1%
She may have been grossly negligent	29.2%	10.0%	0.0%
She waited too long before notifying her bank	58.3%	10.0%	0.0%

**Table 4.18:** Thematic analysis of the answers not in support of reimbursement in Scenario 1: Card Loss.

Table 4.18 presents the reasons provided by the respondents who did not support the reimbursement across all three surveys. About 58% of the German respondents mentioned that Miss K waited too long before reporting the incident to her bank; only 10.0% of the UK respondents provided the same reason, whereas this reason was not mentioned by any of the Americans. Some of the German (29.2%) and UK (10.0%) respondents believed she was grossly negligent without explaining what '*gross negligence*' means. Another reason given is that it was her mistake because she forgot her purse in the train; this reason was shared by many respondents, namely 25.0% (DE), 20.0% (UK) and 30.8% (US). Interestingly, only 4.2% of the German respondents believed that a bank customer is destined to lose, but a much higher percentage provided the same reason in the UK (20.0%) and US (23.1%) surveys. Also, the same distribution was found for another reason: that once the money leaves someone's account, it cannot be retrieved. Another interesting perception is that debit, as opposed to credit, cards are not protected against fraud; this reason was given by 10.0% of the UK respondents and 23% of the US ones. Finally, 20.0% of the UK respondents mentioned that since Miss K's purse is not insured, the only

way to retrieve her money is to catch the thief (as if they simply assumed that the bank would not bear the loss). About 4.0% did not know (or were not sure about) whether Miss K should be reimbursed or not.

#### 4.5.8.2 After revealing relevant Terms and Conditions

Code	DE	UK	US
However, it's hard for debit, as opposed to credit cards	0.0%	0.0%	1.9%
Insurance will reimburse her	0.0%	0.0%	1.9%
She reported the card stolen within the time limits	31.0%	61.1%	98.1%
She used the landline to report the incident	0.0%	2.8%	0.0%
The card was stolen, the transaction was unauthorised, it's fraud	86.2%	63.9%	7.4%
Yes, if it can be proved that the card was stolen	0.0%	16.7%	0.0%

**Table 4.19:** Thematic analysis of the answers in support of reimbursement in Scenario 1: Card Loss, after the participants have seen the T&Cs.

After revealing the terms and conditions to our participants, we were interested in their comprehension. [Table 4.19](#) presents the reasons provided by the respondents who believed that Miss K should be reimbursed (after reading the terms and conditions). About 86% of the German respondents and 64% of the UK ones believed that the victim should get a refund because the card was stolen, and the transaction was unauthorised; only 7% of the Americans provided this reason. On the other hand, 98.1% of the Americans mentioned that Miss K reported the incident within the time limits specified by the Terms and Conditions; this reason was given by 31.0% and 61.1% of Germans and British. No other reasons were mentioned in the German survey. In contrast, 16.7% of UK respondents believed that it can be proved that the card was stolen. One of the Brits reported that Miss K used the land-line to report the incident (2.8%). Another was unsure whether Miss K would be reimbursed or not. One American said that insurance can actually reimburse Miss K, and another believed it would be possible to retrieve the money if the stolen card was a credit card, and not a debit card.

Code	DE	UK	US
It is difficult to recover the money	0.0%	11.1%	100.0%
PIN might have been written down in her purse	16.7%	66.7%	0.0%
She was grossly negligent as she lost her card and failed to immediately cancel it	83.3%	38.9%	0.0%

**Table 4.20:** Thematic analysis of the answers not in support of reimbursement in Scenario 1: Card Loss, after the participants have seen the T&Cs.

In contrast, [Table 4.20](#) displays the reasons mentioned by the participants who said that Miss K should not be reimbursed, after seeing the terms and conditions. Most Germans (83%) said Miss K was grossly negligent because she lost her card and failed to cancel it swiftly. The same reason was provided by 39% of Brits. In contrast, most Brits (67%) believed that Miss K must have written her PIN down on a piece of paper, and left that in her purse; only 17% of Germans reasoned this way. All the Americans who opposed reimbursement said that it is difficult to recover the money; only 11.1% of the Brits gave this as their reason for refusing a refund.

### 4.5.8.3 Analysis

The arguments from both sides are interesting, considering that the protagonist's claim in the UK was denied due to the Ombudsman deciding that the most likely explanation for the fraud was that she had stored her PIN with her card and hence was grossly negligent. Only 10% of the UK participants who argued against the protagonist being reimbursed gave this reason. This changes drastically after the participants have read the terms and conditions: now two-thirds of those who oppose reimbursement give the same reason as the Ombudsman.

We do not know how this case would have been decided in Germany and the US, but we can still analyse the change in perceptions. In the case of Germany and the UK, the percentages in favour of reimbursement did not change with the revelation of the terms and conditions. In the US, however, there was a significant shift to '*She reported the card stolen within the time limits*' from 41.9% to 98.1%. This strongly suggests that our participants read the terms and conditions carefully. In contrast to the American T&Cs ('*sixty days after the statement was mailed to you*'), the German terms do not give a definite time frame as to when a transaction has to be reported as fraudulent. This may have motivated the high response rate in [Table 4.20](#).

## 4.5.9 Scenario 2: phone scam

### 4.5.9.1 Prior to revealing Terms and Conditions

[Table 4.21](#) presents the reasons provided by the participants who initially supported reimbursing Mr L in the *Phone Scam* scenario. A common theme across all three surveys is that banks should secure their systems properly; this was the view of 53.8%, 50.0% and 55.0% of DE, UK and US respondents. Second, banks should be insured, should be ethical, and should be able to reverse any unauthorised transaction; support was 30.8%, 35.0% and 25.0%. Third, Mr L was tricked, but did

Code	DE	UK	US
Banks have good security that should have prevented fraud	53.8%	50.0%	55.0%
Don't know/unsure	0.0%	0.0%	5.0%
He was tricked into phoning the number on the back of his card	30.8%	35.0%	15.0%
If the fraud can be proven	15.4%	10.0%	10.0%
The bank should be insured/reverse the transaction/be ethical	30.8%	35.0%	25.0%
The scammer can be someone working in the bank	0.0%	0.0%	10.0%

**Table 4.21:** Thematic analysis of the answers in support of reimbursement in Scenario 2: Phone Scam.

the right thing by phoning the number on the back of his debit card (30.8%, 35.0% and 15.0%). Additionally, 15.4% of the Germans said that as long as fraud can be proven, Mr L should get his money back; this reason was mentioned by 10% of Brits and Americans each. Only Americans (with 10.0%) said that Mr L should be reimbursed because the scammer might have been a bank employee.

Code	DE	UK	US
Banking accounts have no protection	7.1%	11.8%	16.7%
Banks tend not to care about customers	3.6%	8.8%	5.6%
Difficult to recover the money	7.1%	5.9%	16.7%
Don't know/unsure	0.0%	0.0%	2.8%
His own fault, he was scammed	75.0%	8.8%	19.4%
May have acted fraudulently	17.9%	64.7%	33.3%
No one can tell the difference between the fraudster and the real customer	0.0%	17.6%	30.6%

**Table 4.22:** Thematic analysis of the answers not in support of reimbursement in Scenario 2: Phone Scam.

Table 4.22 shows the reasons given by the respondents who initially opposed reimbursing Mr L. Three-quarters of Germans believed that it was his fault because he fell for a scam; in contrast, most Brits (64.7%) said that Mr L had most probably acted fraudulently; this reason was given by one-third of the Americans but only one tenth of the Germans. Another one-third of the Americans said that Mr L cannot be reimbursed because no one can differentiate between him and the scammer. Some other reasons were mentioned as well across all surveys, such as bank accounts are generally not protected, banks do not tend to care about their customers, and it is hard to recover the money.

## 4.5.9.2 After revealing relevant Terms and Conditions

Code	DE	UK	US
Don't know/unsure	16.7%	0.0%	4.2%
He could not have been aware that there was a technical fix in place	27.8%	48.0%	8.3%
He followed the security procedures as documented for telephone calls	22.2%	28.0%	4.2%
He was not grossly negligent	22.2%	20.0%	0.0%
If the fraud can be proven	0.0%	4.0%	0.0%
It is not an authorised transaction	16.7%	28.0%	75.0%
Phishing not covered by the T&C	16.7%	8.0%	8.3%
The bank can retrieve the money	0.0%	4.0%	4.2%

**Table 4.23:** Thematic analysis of the answers in support of reimbursement in Scenario 2: Phone Scam, after the participants have seen the T&Cs.

After revealing the relevant terms and conditions, the respondents who supported reimbursement provided the reasons shown in [Table 4.23](#). 27.8% in the DE survey said that Mr L would not have thought that a technical fix was in place; this reason was given by almost one-half of the UK participants but only 8.3% of US ones. Another 22% of DE respondents said the Mr L followed the security procedures documented for a phone call, a view shared by 28.0% and 4.2% in the UK and US surveys. Most of the US respondents believed Mr L should be reimbursed because he was not the one who authorised the transaction, a view shared by only 16.7% of Germans but 28.0% of Brits.

Code	DE	UK	US
Difficult to recover the money	8.7%	3.4%	12.5%
Don't know/unsure	4.3%	0.0%	3.1%
He gave his details out on the phone to the fraudsters	13.0%	10.3%	3.1%
It is gross negligence	60.9%	48.3%	28.1%
Mr. L transferred the money himself	34.8%	37.9%	43.8%
Phishing not covered by the T&C	8.7%	0.0%	9.4%

**Table 4.24:** Thematic analysis of the answers not in support of reimbursement in Scenario 2: Phone Scam, after the participants have seen the T&Cs.

Finally, [Table 4.24](#) documents the reasons for why Mr L should not be reimbursed. About 60% and 50% in the DE and UK surveys believed that Mr L was grossly negligent; 28% of the US participants who opposed reimbursement provided the same reason. Another common reason is that Mr L transferred the money himself, given by 35% (DE), 38% (UK) and 44% (US). Other reasons included that Mr

L was the one who gave his details out to the fraudsters, that it is hard to recover the money, and that social engineering attacks, such as phishing are not covered by the bank terms and conditions.

### 4.5.9.3 Analysis

In the UK, the Ombudsman decided that the protagonist was not to be reimbursed as he was deemed to have authorised the transaction and has, hence, been grossly negligent. While the majority of participants from the UK shared the view that he should not be reimbursed (Table 4.16), the majority of participants were unable to give the same reason after reading the terms and conditions, they only decided that he had been grossly negligent. There was a significant shift in the opinion of the German participants after reading the T&Cs: previously the majority had reasoned that *'it was his own fault'*, but this changed to the vaguer but more consistent with the T&Cs view of *'gross negligence'* (Tables 4.22 and 4.24). Interestingly, even though *'gross negligence'* is not mentioned in the terms and conditions shown to the American customers, still 28.1% gave this reason (or one to that effect). But, only in the US did the majority of participants gave the same reason as the Ombudsman, although it is uncertain if the decision would have been the same in the US.

The participants that decided that the protagonist should be reimbursed changed their reasoning significantly after reading the T&Cs. In Germany and the UK, the previously most frequent response—that the bank should have prevented the fraud (Table 4.21)—does not appear as a reason in favour at all in Table 4.23. Instead, the reason has shifted towards the fact that the customer acted with best intentions, and could not have known that he had been reconnected to the fraudsters after following the prescribed security procedure by calling the number on the back of his card. While the participants in the US initially gave the same reasons as those from the UK and Germany, after reading the terms and conditions the vast majority (75.0%) agree that the protagonist did not authorise the transactions. This must have been a clear feature of the terms and conditions presented to the participants from the US.

### 4.5.10 Understanding of Terms and Conditions

The terms and condition documents are not easily accessible, and to be sure that our participants actually spend some time reading them rather than glossing them over, they were shown the terms on a separate page and were instructed to read carefully because they would be asked questions on these terms on the following



page. Participants were unable to return to the terms page once they had left. On average, the participants spend 204 seconds on reading the T&Cs.

A set of comprehension questions followed on the next page. It seems that many of the participants had never read their bank’s terms and conditions before. One comments: ‘*Why am I responsible for closing the door to an ATM lobby as I leave? Why am I being told as a customer to not let people into banks after hours?*’ Each of the comprehension questions solicited a free-text answer, and we subjected the responses to Thematic Analysis.

Code	DE	UK	US
Don’t know	2.4%	7.4%	7.1%
Notified not quickly enough	19.5%	13.0%	80.4%
Shared details	7.3%	27.8%	3.6%
Violate T&Cs	7.3%	18.5%	1.8%
Fraudulently	0.0%	16.7%	0.0%
Always	7.3%	5.6%	3.6%
If you notice something suspicious	0.0%	0.0%	1.8%
Not kept details safe	19.5%	9.3%	3.6%
Been phished	2.4%	1.9%	0.0%
Gross negligence	53.7%	27.8%	0.0%

**Table 4.25:** Thematic analysis of the answers to the comprehension question: ‘*When are you liable for an unauthorised transaction?*’

In [Table 4.25](#), the participants analyse liability. The responses clearly represent the peculiarities of the contracts: for American customers, the only reason to get a non-fraudulent claim turned down is to miss the deadlines. In contrast, in Germany and the UK, the focus is on gross negligence, with 54% of participants from Germany correctly stating that gross negligence is the reason for becoming liable.

Code	DE	UK	US
Don’t know	4.9%	3.7%	12.5%
Carelessness	4.9%	31.5%	46.4%
Not being careful with details	53.7%	48.1%	8.9%
Your fault	2.4%	11.1%	5.4%
Ignoring warnings	2.4%	1.9%	0.0%
Not informing your bank of loss	7.3%	14.8%	14.3%
Negligence beyond reasonable practice	17.1%	9.3%	10.7%
Harmful misconduct	7.3%	0.0%	3.6%
Not following the T&Cs	7.3%	5.6%	0.0%

**Table 4.26:** Thematic analysis of the answers to the question: ‘*What is gross negligence?*’

[Table 4.26](#) follows through by diving into the participants’ understanding of

‘*gross negligence*’. British and German participants agree that ‘*gross negligence*’ is mostly about being careful with details, where details may be any form of credentials or cards. Conversely, the participants resident in the US equate it with the more traditional meaning of carelessness, clearly because their T&Cs do not mention ‘*gross negligence*’ at all. The more legally correct version of ‘*harmful misconduct*’ is mentioned only infrequently.

Code	DE	UK	US
Write down	17.1%	11.1%	26.8%
Change periodically	0.0%	0.0%	21.4%
Memory technique	36.6%	14.8%	16.1%
Use existing/memorable numbers	9.8%	31.5%	14.3%
Choose unique	4.9%	1.9%	0.0%
Just remember it	26.8%	27.8%	25.0%
Write down encrypted	4.9%	3.7%	1.8%
Don't know	7.3%	11.1%	5.4%

**Table 4.27:** Thematic analysis of the answers to the question: ‘*What can you do to remember your PIN?*’

Next, we asked how one was supposed to remember PINs (Table 4.27). Writing down PINs is more accepted in the US than in Germany or the UK with over a quarter of participants stating that the terms and conditions allowed them to do so. Unfortunately, it was difficult here to find sample terms and conditions whose intentions were actually made clear in the extract. Still further insights can be gained: there is a tendency in the US to change PINs frequently, something that was only mentioned in the extract for the American participants. Interestingly, PIN reuse is seen favourably in the UK with 32% of participants noting it as acceptable, an even higher proportion than we found in previous research [41]. We also note that Germans tend to use memory techniques (36.6%) while Brits are more likely to change their PIN to an existing or memorable number (31.5%). We already noted that the UK banks’ association encourages PIN changes, and all banks provide the facility. However, some banks in Germany do not allow customers to change their PINs at all.

In contrast to these tables, we asked the participants to self-judge their own understanding of the terms and conditions. Table 4.28 shows that the vast majority of participants claimed to understand the majority of the terms although less than a quarter of participants from Germany claimed to understand them fully. Given that our participant pool has above average education, it is likely that most bank customers do not fully understand the contract terms of their bank accounts. However, it should be noted that the subject pool from the US thought they understood their

Level	US	DE	UK
Understood nothing	0%	0%	0%
Understood the minority	2%	7%	6%
Understood half of it	4%	12%	2%
Understood the majority	50%	59%	54%
Understood everything	45%	22%	39%

**Table 4.28:** Responses to the question ‘How confident are you that you have understood the T&Cs?’

terms to a much greater extent (although they are about equally well-educated). Perhaps the better consumer protection makes them less cautious. It is also noteworthy that after reading the T&Cs, participants actually realised that they had even stronger rights than they thought.

Code	DE	UK	US
Tips useful	0.0%	0.0%	1.8%
All ok	36.6%	51.9%	73.2%
Complicated	29.3%	13.0%	17.9%
Unclear	51.2%	13.0%	19.6%
Abbreviations, special terms	24.4%	25.9%	1.8%
Gross negligence	0.0%	13.0%	0.0%
Negligence limits unclear	0.0%	0.0%	5.4%

**Table 4.29:** Thematic analysis of understanding issues of the T&Cs of the participants.

Diving into more detail, [Table 4.29](#) lists the broad themes that the participants were struggling with. What most stands out is that in Germany the T&Cs were branded as unclear, needlessly complicated and full of special terms and abbreviations. One participant noted: ‘*Everything is overcomplicated. The terms actively avoid using clear, simple language.*’ We concur; German T&Cs do actually appear much more difficult to understand than the UK’s.

## 4.6 Discussion

Fifteen years ago, when online banking was in its infancy, many banks sought to shift liability explicitly by making customers liable for any transaction where they said the customer’s credential was used. This led to complaints about liability shifting. The situation now is for banks to give instead a variety of different advice, much of it so vague that it is unclear how customers are to set about complying with it, or indeed whether their behaviour is likely to be changed by it at all. In some cases, advice given by banking trade associations is contradicted by member banks’

small print. In the case of the most aggressive banks (in the UK and Singapore), it is probably infeasible for customers to comply with the stated contract terms, and later work will test this on a panel of representative users.

Customers disputing transactions frequently contact the authors to discuss their case. We find that a common approach of the bank is to request that the customer answers a checklist of whether they complied with security recommendations taken from the bank T&C. This advice includes recommendations that never appear in bank publicity, and even contradict advice from banking trade bodies, and so the checklist will likely be the first time the customer has ever seen the recommendations. This creates a climate of expectation in which a court or Ombudsman will be tempted to run through the checklist, in effect asking the customer to prove they were not careless. This also explains a possible reason for banks to list some security recommendations only in terms and conditions that they know the vast majority of their customers will not read. But, rather than a blanket assertion to this effect, the actual argument for refusing a refund will usually be one based on the facts of the case where the bank says *‘Your password was used so you must have been negligent.’* In the US, where consumer laws are held to discourage such an argument, the bank can argue instead *‘As your password was used, you authorized this transaction’* (Patrick 1996). The exceedingly onerous UK bank terms and conditions are particularly worrisome in this context.

The initial draft of the PSD would have restricted UK and other EU banks from using their terms and conditions in this way, as the bank would be required to find additional evidence of negligence, in addition to their own records showing that the transaction was performed with a payment instrument they issued (such as a card). During the development of the PSD, we know from banking industry submissions (Barclays PLC 2002) and our discussions with the individuals involved in the drafting process, that industry lobbied to continue to be permitted to treat their own records as authoritative statements showing that customers are liable for disputed transaction. The banking industry proved successful in doing so by amending Article 59(2) to insert the word *‘necessarily’* and so, in the view of regulators enforcing the PSD who the authors met with, nullifying the original version’s effect. The final version of the PSD Article 59(2), substantially replicated in PSD2 Article 72, reads

*‘Where a payment service user denies having authorised an executed payment transaction, the use of a payment instrument recorded by the payment service provider shall in itself not necessarily be sufficient to prove either that the payment transaction was authorised by the payer or that the payer acted fraudulently or failed with intent or gross negligence to fulfil one or more of his obligations under Article 56.’*

The PSD2 does add the requirement that *‘The payment service provider, including, where appropriate, the payment initiation service provider, shall provide supporting evidence to prove fraud or gross negligence on part of the payment service user.’* Banks may however may continue to decide that their records showing that a transaction was authorised are such sufficient supporting evidence.

Most developed countries have unfair-contract laws, so the question to ask may be: *‘are bank contracts fair?’* Our initial investigation shows that in many cases they are too vague for a firm view to be taken one way or another, and so an assessment will come down to a study of actual dispute resolution practice. However, where contract terms require user behaviour that is far from normal, a usability assessment may provide an answer; and where a banking association advises customers to change all their cards to the same PIN, while some of its member banks have small print forbidding the practice, that is clearly unfair. The unfairness that results from obfuscation does vary, however. Americans tend to be reassured when they actually read their bank contract terms and conditions, while most Germans find them too hard to understand. Overall, the data we have collected gives a number of insights into the effects that the differing approaches to bank regulation have had on consumer expectations between countries. There is much more work to be done here, by researchers and regulators alike.

#### 4.6.1 Data availability

The terms and conditions and their translations that we studied in Section 4.3 can be found at <http://dx.doi.org/10.14324/000.ds.1554770>. The survey, survey data and analysis script for Section 4.4 can be downloaded at <http://dx.doi.org/10.14324/000.ds.1473489>, and for Section 4.5 at <http://dx.doi.org/10.14324/000.ds.1489747>.



# Studies on perceptions of security

*‘There is nothing either good or bad, but thinking makes it so.’*

(W. Shakespeare 1602)

The previous chapter has demonstrated that there are differences between the content of policies, people’s perceptions of the policies and individual behaviours. Policies by different banks for similar products had different and sometimes contradictory security advice. Additionally, participants struggled to comprehend the content of the policies.

This chapter continues to explore the differences between expected and actual behaviours. In [Section 5.1](#) we study how security technology can be explained through the use of metaphors to ordinary people and how their perceptions of security processes differ from reality. I study the intentions of security professionals in [Section 5.2](#).

This chapter helps to establish the discontinuity between users perceptions and knowledge, their actual behaviours, and the security intended by professionals.

## 5.1 Security metaphors

This section focusses on my work in collaboration with Albesë Demjaha, Jonathan Spring, Simon Parkin and Angela Sasse on the effectiveness of metaphors for explaining end-to-end encryption (2018). My contribution in this work is the empirical evaluation of the metaphors. Metaphors are powerful tools for explaining concepts without relying on specialised vocabulary, and are thought to be especially useful to security practitioners for conveying important aspects of security to individu-

als without requiring the exact technical knowledge. This research which fits well into the scope of this thesis as a motivation for the subsequent analysis of security cultures.

The purpose of cryptography has expanded from benefiting primarily the military to securing systems for the general public (Katz and Lindell 2014). For example, widely used messaging applications such as WhatsApp and Telegram have embraced the use of end-to-end (E2E) encryption. This recent flourishing of technology gives non-expert users free access to encrypted person-to-person communication. That general users can take advantage of encryption is however in question: in 1999, Whitten and Tygar claimed that cryptographic details confuse users. Concerns persist about the usability of E2E-encryption, and continue to generate a considerable body of research (Fahl et al. 2012; Bai et al. 2016; Herzberg and Leibowitz 2016; Abu-Salma, Krol et al. 2017; Abu-Salma, Sasse et al. 2017).

Saltzer and Schroeder have long encouraged designers to bridge the gap between a user's mental image of a protection system and the system's specification language (1975). However, few attempts have been made to address this gap. Application properties are communicated to users in technical jargon, potentially obstructing their comprehension of security features (including *encryption*, which itself is a technical cryptographic term). The focus of communication efforts appears to be to teach users how the system works, with the tacit assumption that improved understanding will allow users to complete their tasks. In *'Analogy Considered Harmful'*, Halasz and Moran argue that analogies by themselves are insufficient for teaching people about computing systems (1982). We expand their ideas to metaphors and provide empirical evidence.

Beyond cryptographic tools, systems can be explained to users through accustomed metaphors (Clark and Sasse 1997), which can be extended into design models. diSessa distinguishes between a 'structural' and a 'functional' mental model (1986). A structural model provides a detailed understanding of the system, whereas a functional model provides certain properties of the system which are necessary to complete a real-life task. A functional model is similar to a task-action mapping model, defined by Young as an internalised representation of the system to the real-world task which users have to perform (1983).

After Whitten and Tygar concluded in their *'Johnny'* paper that the 'key' metaphor was misleading (1999), Whitten responded by visually transforming the 'key' metaphor for a secure e-mailing tool (2004). This was effectively a revised structural model intended to cue users' understanding of cryptographic functions. Whitten does not question the appropriateness of the 'key' metaphor itself, nor whether the strategy of cuing structural models will lead to better user outcomes.



There is a dearth of studies that specifically address the issue of inadequate terminology and metaphors to describe cryptographic systems such as E2E-encryption. Furthermore, no attempts have been made to generate and test metaphors in a comparable, repeatable manner. Here we break this impasse and adapt HCI methodology to rigorously generate metaphors for E2E-encryption, using Alty et al.'s model (2000). Alty et al. propose a set of six design steps to provide designers with a practical approach to the application of metaphor in the design of interactive systems. At each step, the authors give an explanation of the activities required. We aim to generate metaphors that cue functional mental models in users (diSessa 1986). Rather than requiring users to explicitly learn all the relevant security tasks, we aim to cue users' already-existing models. Metaphors leverage these existing models to approximate a working mental model for performing a specific task.

In [Section 5.1.1](#) I will describe the metaphors that were constructed using Alty et al.'s framework (2000) and grounded in a number of qualitative studies. These metaphors are tested using the methodology described in [Section 5.1.2](#) and the findings are described in [Section 5.1.4](#).

### 5.1.1 Metaphors considered

Alty et al. propose several approaches for generating metaphors (2000). We adopted three of these approaches: *Design Metaphors*, *Brainstorming* and *Extension*. Design Metaphors emphasises the role of users as a useful source of metaphors. Users often apply familiar metaphors from their everyday life to their language to aid their understanding of a system when undertaking a task.

Brainstorming suggests mapping real-life functionalities to the functionalities of the technical system in order to identify potential metaphors. Finally, Extension promotes the idea of recycling metaphors that are currently used by similar computer systems and extending them in a manner that appropriates the metaphor to the new system.

We introduce five new metaphors in total. The metaphors *Special Language* and *Treasure Hunt* are a product of the Design Metaphor approach, the metaphors *Colours* and *Banknote* are a product of the Brainstorming approach, whereas *Owl* is a product of the Extension approach.

**Special Language** *'Messages and calls with this person will be translated to a special language for which only the two of you know the dictionary.'*

**Treasure Hunt** *'Messages and calls exchanged with this person are like a treasure hidden in a place to which only the two of you know the map.'*

**Colours** *‘Messages and calls you exchange with this person are like colours. Before sending them, you mix them with another colour, known only by you two. Nobody else can retrieve them unless they know the secret colour.’*

**Banknote** *‘Messages and phone calls shared with this person are matched like a ripped banknote, each piece being owned by one of the two people, therefore in order to access the message both pieces are needed.’*

**Owl** *‘Messages and calls with this person will be delivered by your owl which will not share them with anyone else but the two of you.’*

**Telegram/Viber** *‘Secret chats have end-to-end encryption.’*

**Whatsapp** *‘Messages to this chat and calls are now secured with end-to-end encryption, which means the application creators and third parties can’t read or listen to them.’*

### 5.1.2 Survey testing of metaphors

The metaphors are tested through interaction with users (Alty et al. 2000; Becker, Parkin and Sasse 2017b). In order to do this, we design a survey to test whether our new metaphors cue better understanding of E2E-encryption than the existing explanatory metaphors. Within the survey we test the following metaphors: Special Language, Colours, Treasure Hunt, Telegram/Viber, and Whatsapp.<sup>1</sup> We use LimeSurvey for designing the survey and the crowd-sourcing platform Prolific for recruiting participants and distributing the survey.

As our study is purely observational and does not involve any sensitive or personal identifiable data, it is considered a service evaluation by our institution’s Research Ethics Committee (REC) guidelines and is exempt from REC review.

We target a population of users that reside in the United Kingdom, and have heard of at least one messaging tool that adopts E2E-encryption. The demographics that we have collected are age, level of education, adoption of messaging tools, and frequency of use.

#### 5.1.2.1 Functionalities and non-functionalities

To primarily evaluate participants’ current understanding of E2E-encryption, they respond to four statements with either true or false; two of the statements are functionalities of E2E-encryption, the remaining two are not.

##### **Functionalities:**

**Statement 1:** *Only you and the recipient can read your messages* (True)

---

<sup>1</sup>To avoid biases in the survey we remove the words ‘Telegram’, ‘Viber’ and substitute the word ‘WhatsApp’ with ‘application makers’ in the respective industry descriptions.

**Statement 2:** *Other people can send a message pretending to be you* (False)

**Non-functionalities:**

**Statement 3:** *Only you and the recipient can know the messages were sent* (False)

**Statement 4:** *If somebody hacks your phone, they will be able to read your messages* (True)

The purpose of the chosen survey statements was to test understanding of E2E-encryption and simultaneously explore conceptual baggage. Thus, Statement 1 (True) and Statement 2 (False) were chosen to reflect functionalities of E2E-encryption. On the other hand, Statement 3 (False) and Statement 4 (True) were included as conceptual baggage because they are non-functionalities of E2E-encryption. The latter allows us to test whether the metaphors over-promise.

### 5.1.2.2 Survey process

Participants followed the following process:

1. Select which messaging tools they use and how frequently they use them;
2. Respond to the four statements above (true/false, in random order) in the context of using E2E-encrypted messaging applications in general;
3. Read one randomly selected metaphor and answer whether they have encountered the metaphor previously, and if yes, where;
4. Repeat the step of responding to the same four statements outlined above with true or false. The metaphor is still being shown. This evaluates whether the metaphor cues a change in the participants' understanding of E2E-encryption;
5. Ask for any comments or feedback.

### 5.1.3 Reliability and validity

It is crucial to design a survey that is both reliable and valid. Reliability refers to *'the extent to which repeatedly measuring the same property produces the same result'* whereas validity refers to *'the extent to which a survey question measures the property it is supposed to measure'* (Thayer-Hart et al. 2010). To reduce ambiguity in the deployed survey, we conducted several rounds of revision of the questions and descriptions included in the survey, showing them to several pilot participants in each stage. Furthermore, we explore different ways of testing changes in the participants' understanding of E2E-encryption. We give one metaphor per participant to prevent the result of one metaphor impacting the result of another. In addition, by measuring a change in understanding from their initial understanding, we can test whether the results are dependent on prior understanding rather than on the metaphors.

Unintentionally introducing biases is relatively easy when designing a survey. We take a number of steps to ensure that biases are minimal:

1. Only one metaphor is allocated per participant;
2. The allocation of a metaphor to a participant is random;
3. The order in which each statement appears is random (both between surveys and within the same survey).

#### 5.1.4 Findings

We collect a total of 211 valid responses from the survey. One participant is disqualified because of not using any messaging apps and 19 responses are not considered because they come from non-unique IP addresses. All five metaphors appear in the survey: Special Language (39 participants), Colours (47 participants), Treasure Hunt (48 participants), Telegram/Viber (41 participants), and WhatsApp (36 participants). From the survey participants, 57 are male and 153 female. Their ages range from 18 to 64 (average age 35). Two of our participants had no formal qualification, 36 have gone to secondary school, 73 have gone to college, 72 hold an undergraduate degree, 24 have a graduate degree, and two of our participants hold a doctorate degree. This imbalance is typical of studies conducted on Prolific Academic. It should however not bias our results, as previous work has shown that technical people don't understand the encryption either (Abu-Salma, Sasse et al. 2017).

Only 16 participants say they have seen *Telegram/Viber* before and when asked where, their answers include: Facebook (2), WhatsApp (8), BBC News (2), terms and conditions of WhatsApp (1), Facebook Messenger (1), WhatsApp group screen (1), and on a news site like BBC (1). Similarly, only 18 participants say that they have seen *WhatsApp* before and when asked where, almost all answers are associated with WhatsApp (16), in addition to: in app description or updates (1), Facebook (1), and Facebook Messenger (1). One participant says they have seen the *Treasure Hunt* metaphor online. Two participants say they have seen the *Special Language* metaphor on WhatsApp. Two participants say they have seen the *Colours* metaphor on a news story talking about encryption and on tech websites discussing secure messaging.

##### 5.1.4.1 Statistical results

Table 5.1 compares the changes in participants' responses to the four statements due to seeing the descriptions. The first four rows indicate the difference in the number of correct responses to the statements. For example, the first 12% indicate

Statement	Telegram / Viber	WhatsApp	Treasure Hunt	Special Language	Colours
1	12%	25%	6%**	8%**	15%**
2	0%**	11%**	-4%**	-3%**	0%**
3	-5%**	-11%**	4%**	13%**	-4%**
4	-12%	-25%*	-4%**	-10%**	-11%
mean	-5%	0%	2%	8%	0%
negative mean	-52%**	-53%**	-25%	-18%	-30%*
changes	0.90**	1.06**	0.56**	0.59**	0.60**

**Table 5.1: Distribution of the changes in the participants’ responses** Statistical significance is indicated by \*\* for  $p < 0.01$  and \* for  $p < 0.05$ , calculated by a Fisher’s exact test.

that 12% more participants answered statement 1 correctly after seeing the Telegram/Viber description than before seeing the statement. There are a number of interesting trends to note: in general (and especially for the *WhatsApp* description), participants’ understanding of the applications functionality (Statements 1 & 2) were improved, while the non-functionality statements suffered. The *Treasure Hunt* and *Special Language* metaphors appear to be most balanced in the effect on participants understanding.

This is followed by the mean of the previous scores for each description. Here, a Wilcoxon signed-rank test was performed, but no description displayed statistically significant variations. While for individual statements the descriptions show positively and negatively statistically significant variations, these variations balance out upon aggregation. Hence we find no evidence that any of the descriptions give rise to an improved understanding of E2E-encryption.

However, when we penalise descriptions for causing participants to change their previously correct responses, descriptions 1 (Telegram/Viber), 2 (WhatsApp) and (less so) metaphor 5 (Colours) appear to cause harm (row *negative mean*, Wilcoxon signed-rank test). These descriptions appear to be undoing participants’ existing mental models.

The bottom row indicates the mean number of responses that participants have changed after seeing the metaphor (with a maximum of 4). The mean participant

has changed on average 0.73 of their responses (out of a maximum of 4 changes, statistically significantly according to a Wilcoxon signed-rank test at  $p < 0.01$  for all descriptions). This confirms that participants have changed their perceptions in response to the statements, and that the results are not based on a small subset of participants.

We also perform pairwise tests for each of the statements' change scores. Here we find not a single statistically significant result, indicating that in pairwise testing no description outperforms any other (the sample size exceeds the requirements for observing a medium effect size ( $f = 0.25$ ) in pairwise testing). This does not contradict the variations described in [Table 5.1](#), as those tests were performed in-sample. So while there is evidence that some descriptions cause harm, we cannot conclude that some descriptions actually perform better than others in improving understanding of E2E-encryption.

We find a correlation between over-promising the capabilities of E2E-encryption and this decrease in performance. Essentially, the industry descriptions are more likely to make participants believe E2E-encryption does more than it actually can provide. The WhatsApp and Telegram/Viber descriptions carry more conceptual baggage than the other tested metaphors. Since baggage refers to properties of the metaphor not in the system, this is one possible explanation. A hypothesis is that the industry metaphors are more likely to influence incorrect answers post-exposure for the two questions which had a correct answer of 'no' (statements 2 and 3). The change scores for these two metaphors and statements are: 0, -2 and 4, -4, giving an aggregate of -2. A Fisher's exact test supports the hypothesis with  $p < 0.01$  and a large effect (Cramer's  $V = 0.55$ ).

Lastly, our analysis indicates that there is no statistically significant correlation between the participants' age, their frequency of use, and if they claim to have seen the statement previously and the descriptions' change scores, i.e. none of the additionally captured demographics and app usage statistics have any impact on the participants' responses.

### 5.1.5 Discussion

It is unsurprising that if users do not have the same meaning assigned to these terms as the developers do, there will be problems. For E2E-encryption generally, (1) participants do not know the meaning of the system, (2) they give incorrect non-technical descriptions of the system, and (3) they give incomplete non-technical descriptions of the system. This state of participant understanding helps explain why the metaphors based on technical jargon and structural mental models do not

perform well.

The most natural question is to ask why users do not have the correct understanding. From the data we have available here, there are a number of avenues which can be explored. For one, other work has suggested that security concepts are too complex to explain in a simple metaphor, certainly for the purposes of risk communication (Camp 2009). Future work may explore ways to address specific elements of how E2E-encryption works which require functional contribution from users; that is, focus on explaining correct user behaviours to enact, as opposed to requiring users understand the system structure correctly. This can be put into action by developing task-action statements to test metaphors against when applying our rigorous methodology.

Complementary to this, our results strongly suggest that metaphors for security communication must be rigorously tested. We found that existing explanations created by domain experts fall short due to their attempt to engender structural mental models, and this is not doable in the available space and time. In addition, Sasse, Brostoff and Weirich argue that the goal of effective security can be achieved by considering the primary task, context of use, as well as strengths and weaknesses of users (2001). These factors are still widely ignored even when designers attempt to incorporate usability; they ultimately develop systems based on their perception of what is usable, where it may be more productive to focus on the mental models and capabilities of users (Cranor and Garfinkel 2005).

Furthermore, both new and old metaphors cause measurable harm to participant understanding. This harm appears evenly distributed whether participant beliefs start off accurate or not, and appears independent of participant experience with the technology. Related works examining metaphors for E2E-encryption (e.g., Whitten (2004)) have not included in their methodology design the capacity to check for negative impact upon participants' prior beliefs. This omission creates a clear opportunity for bias, in which researchers selectively measure positive changes without checking what the collateral damage to understanding has been. Testing for harm as well as improvement is a clear contribution of our own approach which we recommend be adopted in future studies.

Transferring knowledge or understanding to novices is an important aspect of usable security. We can view the task of devising explanatory metaphors for E2E-encryption via work on resituating knowledge more generally (Morgan 2014). The knowledge is local to experts, and the target audience is novices. The jargon-based explanations essentially try to bridge the knowledge directly to the novices. The more friendly metaphors serve as an intermediate generalisation which the novice can then attempt to localise to their personal situation.

## 5.2 SANS analysis

The SANS Securing The Human group conducts an annual questionnaire on the state of security awareness in the industry. The questions in general focus on broad trends throughout the community, however in 2016 and 2017 I was part of the survey team and helped design and analyse the survey.

The final output of the survey is glossy marketing (SANS Securing The Human 2016, 2017), with the main take-away being:

1. Support is essential: the more money a program has, the more mature it is;
2. Soft skills are lacking: programs are run by people from a technical background;
3. Security awareness is still in its infancy.

However there are a few questions that are interesting in the context of this thesis. I will discuss two questions in more detail in this section; with other (less relevant) analysis included in [Appendix B](#) for completeness.

Response	2016	2017
Non-existent	7.5%	7.4%
Promoting awareness and behaviour change	52.3%	54.9%
Compliance focused	25.3%	27.0%
Robust metrics framework	2.3%	0.9%
Long-term sustainment and culture change	12.7%	9.9%

**Table 5.2:** Comparison between 2016 and 2017 for question: *‘How would you classify the maturity of your organisation’s security awareness program?’*

The survey defines a scale for the maturity of security awareness programs (see [Table 5.2](#)). The majority of responses point towards security awareness for the sake of it. A quarter of responses rationalise security awareness as a means to improve compliance. The responses support the motivation of our research: security awareness is used for encouraging compliance to the security policy.

Second, the survey asks the professionals regarding their organisations approach when individuals fail to exhibit the correct security behaviour (see [Table 5.3](#)). There is a near equal split of respondents that state that their organisation punishes repeat ‘offenders’ for non-compliance and respondents who’s organisation don’t penalise employees for incorrect security behaviours.

However only a tiny fraction (2.3% in 2016 and 0.9% in 2017) are supported by actual measurements of employees’ behaviours. While security awareness and behaviour change are promoted in more than half of all responses, their effectiveness is rarely measured.



Response	2016	2017
Punitive: security policies and secure behaviours are the ‘flaw of the land’ and people who violate them should be punished.	4.4%	4.3%
Progressive: security policies and secure behaviours are enforced on a sliding scale. First time offenders are not punished, but people are punished the more times they fail.	49.5%	47.1%
Flexible: security policies are more like guidelines and should not get in the way of effectiveness.	46.1%	48.6%

**Table 5.3: Frequency of answers to question 16 in 2016 and 2017** ‘What is your organisation’s approach when people fail to exhibit the correct secure behaviors?’

	Non-existent	Compliance focused	Promoting awareness and behaviour change	Long-term sustainment and culture change	Robust metrics framework
Punitive	0.0%	0.9%	2.2%	1.3%	0.0%
Progressive	1.6%	10.1%	27.4%	7.6%	1.3%
Flexible	5.7%	13.6%	21.1%	3.5%	0.9%

**Table 5.4: Frequency analysis of 2016’s Q14 vs Q16** Q14: ‘How would you classify the maturity of your organisation’s security awareness program?’ vs. Q16: ‘What is your organisation’s approach when people fail to exhibit the correct secure behaviors?’

When analysing [Tables 5.2](#) and [5.3](#) in combination ([Table 5.4](#) for 2016, and [Table B.12](#) for 2017), it emerges that as programs advance on the awareness maturity metric, their response to incorrect behaviour moves from *flexible* to *progressive*. However only a tiny fraction (1.3% in 2016, and 0.7% in 2017) of *progressive* approaches are supported by a metrics framework.

This is counter-intuitive: we would expect that compliance-based programs rely on punitive behaviour as this is their only option to influence behaviour. Similarly a program with a *robust metrics framework* should be able to accurately assess weaknesses and adjust training accurately in order to educate better rather than punish. Clearly better tools for measuring the impact of security awareness campaign are wanted.

This divergence highlights a potential miss-conception of security awareness professionals: as the programs mature and evaluation become available, accurate punishment becomes possible and is hence used in order attempt to promote change.

### 5.3 Conclusion

In this chapter we have explored the differences between expected and actual behaviours. Security concepts are often complex, and attempts to overly simplify their explanations may even have adverse affects. While real-world analogies are often thought to be useful, we demonstrate that they can in fact adversely affect our participants understanding (Section 5.1).

Directly related is the study of security awareness professionals (Section 5.2), who hunt for short, catchy phrases that improve employees' security posture. Their awareness material may well be an analogy similar (in concept) to the ones in Section 5.1. Yet this community does not evaluate the effectiveness of their messaging, and when its impact is measured, it is used for disciplinary actions.

In the next chapter I will address this gap through a methodology that measures organisational security and provides tools to dissect the organisational structures by security behaviours.

## Productive Security

*‘Most discussions of decision making assume that only senior executives make decisions or that only senior executives’ decisions matter. This is a dangerous mistake.’* (Drucker 2004)

In the previous two chapters I have established the gap between policies and practice (Sections 4.3 and 4.4), the importance of taking user perceptions into account and designing security with these in mind (Sections 4.5 and 5.1), and identified the serious demand to support interventions by measuring their impact (Section 5.2).

In this chapter I focus on measuring and understanding organisational security. I present a methodology for gathering large scale data sets on employee behaviour and attitudes via scenario-based surveys. The survey questions are grounded in rich data drawn from interviews, and probe perceptions of security measures and their impact. I study employees of two large multinational companies, demonstrating that our approach is capable of determining important differences between various population groups.

First, in Section 6.3, I describe the overarching methodology. The survey responses are analysed in Section 6.4. I develop a validation technique for the survey mapping by coding optional free-text responses in Section 6.10. Based on the survey responses I reframe Security Champions as a tool for targeting policies and interventions in Section 6.8.

## 6.1 Introduction

Technically-focused sources of data—such as system logs—may be used to support analysis of policies (and indeed I do, in [Chapter 7](#)), they do not provide an insight into employees’ thought processes. Security systems are not just the sum of their technical components. User co-operation plays a critical role in providing organisational security, which highlights the need to consider the relationships between people, process, and technology (Dhillon and Backhouse [2001](#)). In addition, an over-reliance on technical solutions can hinder an organisation’s capacity to support employees in their productive tasks (Tariq, Brynielsson and Artman [2014](#)). Behavioural data is therefore an important factor for effective security management, and a goal of this work has been to create a set of repeatable metrics capable of assessing employee attitudes and behaviour around security.

In particular, we develop a methodology capable of identifying areas in which the security policy itself creates incentives for negative behaviour. Rigid systems can force compliance with policy but risk causing disgruntlement (Beautement, Sasse and Wonham [2008](#)). Where conflict exists between security systems and productive tasks, friction results. Workarounds and ‘*circumvention strategies*’ (Adams and Sasse [1999](#)) are then likely to develop as users take advantage of system flexibility to modify how technology and procedures work. This reduces security effort but often introduces security vulnerabilities as a side-effect (e.g., using the same password for a number of accounts across both work and personal life). Managers may even be complicit in supporting workarounds if secondary tasks (such as security) stand in the way of business continuity (Röder et al. [2014](#)). Different threat models exist within different areas of life, so the vulnerabilities in one space can weaken security in others (e.g., carrying unencrypted USB data devices in transit between work and home) (Beautement et al. [2009](#)).

The methodology is a multistage process designed to elicit realistic responses from employee populations at scale. As it is necessary for our data to support the decision-making process of different organisations of all sizes, both of these points are of great importance. Data that does not closely represent the operational reality of the organisation cannot be used to drive decision-making, as it is not a reliable predictor of future states and outcomes. Likewise a data collection method that is overly time-consuming or does not scale (potentially up to tens of thousands of participants) quickly becomes impractical for larger organisations. Both of these concerns are addressed by the Productive Security (ProdSec) methodology.

## 6.2 Related literature

A number of existing works use surveys and/or interviews to explore the relationship between an organisation's information security policy and employee behaviour. These works examine the impact of attitudes and perceptions on behaviour and consider both intrinsic and extrinsic influencing factors. For example, Pahlila, Siponen and Mahmood (2007) found that attitudes towards security<sup>1</sup> and the habits of individuals can have a significant effect upon the intention to comply with security policies. They also assert that the social environment around an individual will have an effect upon their propensity to comply with policy.

Expanding on intrinsic motivators, Rhee, Kim and Ryu (2009) use social cognitive theory to model the influence of experience with security incidents upon self-efficacy, and the role of self-determination upon the outcome of security-related scenarios. A large-scale survey completed by ~400 students found that individuals with high self-efficacy used more security tools and were more vigilant to security, and experience of security compromises negatively impacts self-efficacy.

The notion of competence was also investigated by Workman, Bommer and Straub (2008), who explored the *'knowing-doing'* gap in individuals who have appropriate security skills and knowledge, but who do not apply these skills consistently. Based on the results of a survey in which 588 members of a technology services company participated, the paper concludes that security technology should be user-centred to avoid a tension between assessing threats and use of coping responses.

Siponen, Pahlila and Mahmood (2007) utilise Protection Motivation Theory (PMT) to reason about employee compliance with information security policies. The work considers component parts of PMT, namely threat appraisal and coping appraisal (where this includes response costs). A survey was conducted with 917 employees of Finnish companies. Amongst the findings, threat appraisal was found to have a significant impact on intention to comply with information security policies. Employee beliefs about their ability to adhere to policy influence their intention to comply. This finding stresses the importance of perception; the authors assert that if policies are not perceived as relevant by an employee, compliance to policy will be diminished.

Perception was also the focus of work by Bulgurcu, Cavusoglu and Benbasat (2010), who infer that the perceived costs and benefits of compliance (or non-compliance) are formed by the perceived consequences. The authors find that intention to comply is heavily influenced by attitude, beliefs and ability to comply. The relationships between these factors are explored using a survey of 464 employees

<sup>1</sup>As previously, constructs/abstract ideas that are measured are underlined throughout.

across a number of organisations. The study identified three belief classes relating to consequences of compliance decisions: benefit of compliance, cost of compliance, and cost of non-compliance.

The prevalence of attitude and perception as themes throughout these works strongly influenced our survey design. However, these surveys all rely on some sort of rating (e.g., Likert) scale, or a sliding scale (e.g., keeping information safe is beyond, or within, a person's control).

We build on these themes but opt to take a more immersive scenario-based approach in this chapter.

A related approach is taken by Albrechtsen and Hovden (2009), utilising the differences in skills, perceptions, and interpersonal relationships to characterise the 'digital divide' between information security managers and end-users. The researchers analysed interviews with 11 managers and 18 employees alongside complementary web-based surveys exploring how 87 managers and 151 users assess security threats and vulnerabilities. The study acknowledges that users prioritise other work tasks, that policy is potentially impenetrable and hard to find for the non-expert, and that security provisioning is often one-way. The methodology chosen here extends this approach by grounding survey questions in interview outcomes, towards greater resonance with real-world user experiences.

Gaw, Felten and Fernandez-Kelly characterise the adoption of encrypted emails at a large US cooperation through 4 vignettes (2006). Employees were found to be either cautious, borderline paranoid habitual user of email encryption, technology-loving self-senders, ephemeral users or uninitiated user who believe they have nothing to hide. Given that the organisation's business was centred around secret, non-violent direct political action, universal, routine use of encryption was seen as paranoid. These vignettes correspond well to the behaviour types we formulate for Company A (Section 6.4.4.2), however with the focus on privacy preserving behaviours rather than security more generally.

Other methods of constructing scenario content have been attempted. Both D'Arcy, Herath and Shoss (2014) and Parsons et al. (2014) generate survey questions by drawing on existing literature and interviews with experts. While this makes good use of general information, it does not allow for surveys to be tailored to the specific context of deployment.

D'Arcy, Herath and Shoss use their survey to explore links between stressful information security demands and intentional violation of security policies, to identify workplace factors which contribute to policy violation, including overload, complexity, and uncertainty (2014). Stressful conditions contribute to security coping strategies, as behaviours are adapted in response to stress factors, which then have a knock-on

effect on productivity. Where security requirements are perceived as overloading, complex and uncertain, users then become disengaged, implying that high-effort policies can themselves promote insecure behaviour. The inclusion of productivity as a consideration is of particular interest here, mirroring our goal of ‘Productive Security’.

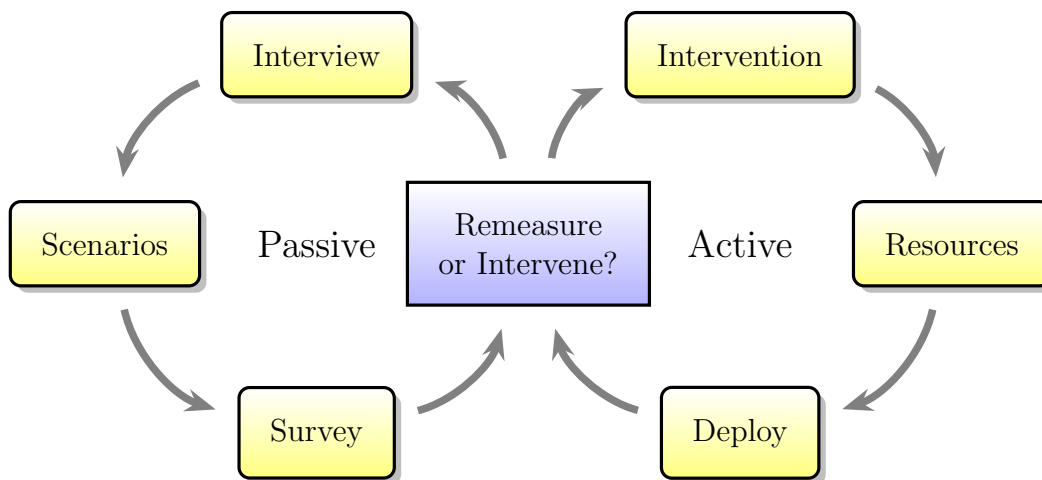
Counter to D’Arcy, Herath and Shoss (2014), Guo et al. (2011) propose a model of what is referred to as ‘Non-Malicious Security Violation (NMSV)’, validated by a survey, delivered in both paper and web formats, of employees and their working conditions. The authors look beyond visible behaviours and instead examine the role of attitudes toward security policy violations, such as the productivity advantage of non-compliance, perceived risks, and workplace norms. As in D’Arcy, Herath and Shoss’s work, scenarios are developed based on related literature and interviews with security practitioners and experts, where end-users of policy are not directly engaged. Results imply that job performance advantages, perceived security risk and workgroup norms are key predictors of intention to engage in NMSVs, with users favouring business tasks. The study also found that attitudes toward security policy itself were not significant in driving non-compliant behaviour, in contrast to Bulgurcu, Cavusoglu and Benbasat (2010). The authors recommend a user-centred security management strategy, where employees can satisfy productivity goals while also maintaining security.

Parsons et al. (2014) present the Human Aspects of Information Security Questionnaire (HAIS-Q), which determines the threat posed by employees by examining the relationships between knowledge, attitude, and security behaviour. HAIS-Q is tested with 500 employees working in a commercial environment. The goal was to focus on areas of policy that are relevant to both employers and employees, and also susceptible to non-compliance. They were determined by a policy review in consultation with senior management representatives. Areas include Internet and email use, password management, and information handling. Statements were derived for each area, focusing separately on knowledge, attitude, and behaviour. Respondents rated their agreement with the statements on a Likert scale. Results imply that security behaviour is driven by knowledge of, and attitudes toward, policy and procedure. The authors suggest that targeting attitudes will improve policy compliance, and that training should relate the importance of compliance to behaviours.

A scenario-based approach is also taken by Blythe, Coventry and Little (2015), who study how individual and organisational factors in the workplace impact secure behaviours, using interviews based on 16 ‘vignettes’. These cover security behaviours identified from information security policies. Vignettes were effectively used as a device for building a rapport with participants and eliciting attitudes and beliefs

relating to a specific subject, an approach reflected here in our interview technique. Results suggest that research should focus on individual security behaviours rather than beginning and ending with policy compliance, and that participants accepted responsibility for some elements of security, while leaving others to their organisation.

### 6.3 Methodology



**Figure 6.1:** Overview of processes in our methodology

The goal of the ProdSec methodology is to provide researchers studying organisations with a repeatable, scalable data gathering process that allows them to better understand the security-related issues facing their employees, and the behaviours and attitudes they adopt in response.

The full ProdSec methodology consists of two independent, iterative processes. [Figure 6.1](#) illustrates the steps involved. The two cycles represent a *passive* data collection and monitoring phase on the left, and an *active* intervention phase on the right. Unfortunately the ProdSec project never reached the intervention phase, although an overview of the active cycle is presented here for contextual clarity.

The passive cycle identifies the predominant security behaviours and attitudes within an organisation, along with specific points of friction between the business and security processes. In order to support real-world decision-making, data collection must both accurately represent the real-world environment where it is applied, and be sufficiently scalable so as to be of use to potentially very large, multi-national organisations. These two goals are to some degree in conflict. Rich, in-depth data capable of accurately representing a real-world context can require a greater invest-



ment of time and effort to collect, making it problematic at scale. ProdSec tackles this by utilising a two-stage method.

Firstly, semi-structured *Interviews* were conducted with a vertical cross-section of the organisations to capture attitudes and behaviours across as many roles, physical locations and demographic groups as possible. Based on interview findings, we carefully craft a scenario-based survey that reflects dominant security-related issues. By tailoring our survey to each operational context, we ensure that survey questions are relevant and recognisable to participants, with the aim of eliciting more realistic and genuine responses.

Once this cycle is complete, security practitioners then have the choice of monitoring the situation over time by repeating the measurement cycle at some future interval (e.g.,  $\approx$  6 months later), or actively engaging with any uncovered problems. The right half of [Figure 6.1](#) describes the active *Intervention* phase. Based on the conclusions drawn from the passive cycle we work with the organisation to prioritise the issues identified, and design and deploy optimal intervention(s), taking into account business as well as socio-technical factors [Section 6.4.7](#). Direct collaboration is important as interventions also need to be suitably centred around the human.

The Productive Security project never reached the intervention phase. The cooperating companies were unwilling to make the interventions recommended, citing other priorities. It was also difficult to convince the leadership of the organisations to work in a scientific manner by establishing a baseline and repeat measurements over time.

Visualising the methodology as cyclical is essential for understanding its intent, reflecting the notion that security is a process and not a fixed state. It is our experience that organisations often see the implementation of an intervention as the final step, whereas we consciously position this step as part of an ongoing process. The passive cycle can therefore be used to track changes in attitude and behaviour over time, as a consequence of the organisation's evolution or in response to specific interventions. Ongoing monitoring can inform decision-makers as to whether interventions are having the desired effect, and indeed whether interventions have themselves influenced security processes. Likewise, no one set of interventions will provide a 'silver bullet' solution, requiring repetition of the active cycle.

The analysis in this paper focuses on the data collection and analysis stages of the passive cycle. As this research spans three years, with two main phases of data collection, lessons have been learnt along the way. As such, we refined our methodology between the two rounds of data collections, both at multinational companies. I will discuss the methodological differences and results for each organisation separately in [Sections 6.4.4](#) and [6.4.5](#). As I did not conduct the semi-structured interviews

that the surveys are grounded in, they have been omitted from this thesis. The ten scenarios described in [Appendix A.3](#) for Company A directly correspond to issues that were identified in the semi-structured interviews, with topics ranging from file sharing, permission management to tailgating and external threats. There was some overlap in topics in Company B, which allowed us to re-use some of the scenarios. New scenarios allowed us to explore issues on password managers, VPN usage, as well as sector specific topics such as credit checks ([Appendix A.4](#)).

A number of other works analysed the interview transcripts for different aspects. Kirlappos, Parkin and Sasse have identified the concept of *shadow security* through the analysis of these datasets, where security conscious employees who think they cannot comply with the prescribed security policy create a more fitting alternative to the official policies and mechanisms (Kirlappos, Parkin and Sasse 2014; Kirlappos, Parkin and Sasse 2015; Kirlappos 2016). This insight of intentional policy violations in order to achieve more workable security inspired us to quantify this behaviour further.

[Section 6.4.1](#) describes the methodology from semi-structured interviews to scenarios. As a last step in the passive cycle we study the responses of employees to the scenarios, by way of the survey. We discuss the analytical approach in [Section 6.4.6](#) and [Boxout 1](#). The results from the survey and related analysis are described in [Section 6.5](#).

## 6.4 Surveys

### 6.4.1 Online scenario-based survey

The approach to scalable data collection was driven by an online survey. In order to efficiently reach large numbers of people, it was necessary to allow them to take part in the data collection exercise from any location, in particular from their usual work environment. Not only does this increase the response rate by minimising demands on participation, but it also furthers our aim of making data collection as naturalistic as possible, as the collection environment matches the operational environment being assessed. As the primary method of recruitment and communication with participants was through company newsletter emails, we embedded a link to the survey in an issue of the newsletter.

As described in [Chapter 3](#), many surveys present short questions with either a multiple choice or Likert scale-style answers. It is our view that this approach is unlikely to engage participants, in particular due to a widespread fatigue with questions of this style. Participants are likely to skim through the survey and apply

little thought to their answers. Also, short questions would not allow us to utilise the full value of the interview information. As such, we elected to build scenario-based survey questions, in which participants were presented with one of the common situations identified via our interview analysis.

Once a topic was selected the scenario was written, using organisation-specific details and terminology from the interviews (see [Section 6.4.2](#)). For each scenario, we also created four possible answers or outcomes, again drawing on the interview data to craft these so they appeared familiar and plausible to the participants. How these options were then used will be covered in more detail in [Section 6.4.5.3](#). Our goals here were to:

- Present scenarios to the participants that seemed both realistic and familiar;
- Offer answer options that were likewise realistic and familiar;
- Gather as much implicit data as possible to maximise the benefit of the survey while minimising the time taken per participant.

## 6.4.2 Scenario commonality

Although the companies that we worked with operate in different sectors there was some overlap in the issues that arose as the result of our interview analysis. Clear desk policies, tailgating through physical security and file sharing were present as issues in both environments. This suggests that as more companies are assessed, a database of scenarios could be developed that over time reduces or minimises the need for the interview stage. However, despite the similarities in topic area it was still necessary to alter key details in the text accompanying each scenario, in order to present to each participant set a scenario that approximated the reality of their environment. For example, one company observed tailgating through security doors, the other through turnstiles in their foyer. Accurately representing these details increases the realism of the scenario, with an aim to encourage honest responses.

### 6.4.2.1 Survey biases

A potential problem with surveys, especially one covering a sensitive topic such as security, is that participants attempt to give the answers they feel are expected of them, or are correct, rather than an honest reflection of their own thoughts and views. These deviations fall into two main categories, response bias and demand characteristics. We addressed these in different ways.

Response bias refers to biases introduced by the participant being influenced by sensory inputs and cognitive processes when answering the question, and thus unintentionally altering their response. For example, how a question is phrased can

alter the response (in particular if it is a leading question). Aside from eliminating instances of leading or priming language we also took care to phrase our attitude scenarios from the point of view of a fictitious colleague, as participants are often more comfortable reporting on the actions of others. So rather than asking, ‘what would you do in this situation?’, we asked, ‘what would Jessica do?’, intending this to counter some aspects of the response bias. This helps us obtain an accurate representation of the maturity levels of the employees.

Demand characteristics refer to the fact that research participants might speculate about the purpose of the research and give responses that they think align with what the researchers are trying to find out. This is something we were particularly concerned about, as security is a sensitive topic with potentially significant outcomes. As we were essentially asking people to report on their own rule breaking, the potential for participants trying to give ‘right’ answers was high. To tackle this we made sure that in each scenario there were no obviously correct options: each possible answer involved some difficulty or transgression.

### 6.4.3 Survey design decisions

We split the scenarios into two types, based on our interview analysis. When participants reported incidents it was either in the form of something they did themselves, or something they observed their colleagues doing. Our scenarios followed the same approach, and were divided into Behaviour and Attitude scenarios. The full scenarios, for both companies can be found in [Appendices A.3](#) and [A.4](#) respectively.

The scenarios present the actor with a situation that puts the requirements of the primary business process in conflict with some aspect of the security policy. Typically this involves the actor in the scenario needing to complete a specific task, the completion of which is being slowed down or prevented by a security process or mechanism. Four options were then given that presented courses of action that would resolve this conflict. Each of these options contained an element of non-compliance so as to avoid participants seeking to give the ‘right’ answer.

In pursuit of our goal to capture rich behavioural data, we linked each answer option as closely as possible to one of four behavioural or attitude risk types. This meant answer choices also allowed us to monitor the prevalent behaviour types. The linking for Company A was carried out by Adam Beutement, Simon Parkin, Iacovos Kirlappos and Angela Sasse (I was not a part), and for Company B by Adam Beutement, Ingolf Becker, Simon Parkin & Angela Sasse. For now we take the accuracy of this mapping for granted, however we devise and execute a methodology for validating this mapping in [Section 6.10](#).

Crossler et al. (2013) posit that cultural theory can be used as a predictor of the impact the norms of an organisation can have upon perceptions of security-related risks. This motivates the choice of two scales that attempt to contrast attitudes with behaviours. The following section describes the scales used in Company A (Section 6.4.4), and their evolution (Section 6.4.4.3) to the surveys used in Company B (Section 6.4.5).

#### 6.4.4 Company A

There were 10 scenarios for employees at Company A, 5 for attitude types and 5 for behaviour types. The full text and answer options for these scenarios can be found in Appendix A.1.

##### 6.4.4.1 Attitude types

The attitude scenarios attempt to elicit the employee's willingness to own security through the mechanisms that are available in the organisation or otherwise. The first three attitude types lend themselves to be ranked (where  $3 > 2 > 1$ ) in the order of the organisations preference of employee's actions when faced with a threat. Attitude type 4 includes some risk to the employee themselves, so it cannot be endorsed as a preferred outcome, and indeed organisations that have a large number of participants of Attitude type 4 should reconsider the effectiveness of their existing security measures.

The Attitude types are:

**Attitude type 1** Discount suspicions, cause no bother, passive;

**Attitude type 2** Report suspicions but take no direct action;

**Attitude type 3** Take direct action through official channels;

**Attitude type 4** Take direct personal action against the threat.

As for all types of scenario questions in this chapter, participants are asked to rank all for options. Additionally, they are asked to assign each type an *appropriateness* score. If an employee assigns a low appropriateness score to a type it may be unlikely that they would actually follow through with this option.

##### 6.4.4.2 Behaviour types

The second set of scenarios elicit the participant's preference between security and productivity.

The Behaviour types are:

**Behaviour type 1** Prepared to perform insecure acts to maximise productivity;

**Behaviour type 2** Show a minor priority for work over security when the two conflict;

**Behaviour type 3** Passive, expects others to take the initiative to ensure security;

**Behaviour type 4** Tries to remain secure wherever possible.

Behaviour types are again ranked by the participant, and each answer option is assigned a Severity score. Ideally the distribution of severity scores for the different answer option has a low variance, so that employees can choose the option that represent their actual behaviour rather than being influenced by options that have extreme consequences.

Additionally, the employees perceptions of the *Acceptability of failing the work task* are measured on these behaviour type scenarios on a 5 point likert scale (ranging from ‘*Not Acceptable At All*’ to ‘*Very Acceptable*’).

#### 6.4.4.3 Evolution of theoretical constructs from A to B

The survey scenarios presented in Company A focussed on employee’s responses to the perceived trade-off between security and productivity. This lead to scenarios that narrowly focussed on these two issues without considering the wider issues of employee non-compliance. Participants subsequently chose the answer they were knew was expected of them, essentially creating a social desirability bias. Our analysis of the acceptability scores in [Section 6.5.1](#) will highlight this further: there is a strong correlation between the acceptability of the answer option and their ranking (Kendall’s  $\tau = 0.72$  and  $\tau = -0.28$ , both  $p < 0.01$  for Attitude types and Behaviour types respectively, [Table 6.4](#)). This is addressed in the scenario design for Company B, where we generally aim to understand how security is perceived by employees, relative to their primary tasks. We avoid the potential desirability bias by drawing on more abstract, not orderable constructs. Kendall’s  $\tau$  for acceptability scores will have decreased in magnitude to 0.62 and  $-0.20$  for Company B for Maturity levels and Behaviour types respectively (both  $p < 0.01$ , [Table 6.4](#)). At the same time, my validation study in [Section 6.10](#) identifies that the variance of acceptability scores between the different answer option is lower in Company B than in A, indicating a smaller likelihood of social desirability bias.

### 6.4.5 Company B

#### 6.4.5.1 Behaviour types

The behavioural risk categories used in our surveys stem from Adams (2003), and their characteristics are given below:

**Individualists** rely on themselves for solutions to problems;  
**Egalitarians** rely on social or group solutions to problems;  
**Hierarchists** rely on existing systems or technologies for solutions to problems;  
**Fatalists** take a ‘naive’ approach to solving problems, feeling that their actions are not significant in creating outcomes.

Individualists may for instance feel less loyalty to others in the organisation and to policy, but may be more likely to report what they see as inappropriate behaviour to others (Crossler et al. 2013).

#### 6.4.5.2 Maturity levels

To further explore the security culture in an organisation, we complement the behaviour-type scenarios with attitude-type scenarios. Rather than presenting participants with a task, the actors described in the attitude-type scenarios observe an instance of non-compliance in their environment—such as finding a screen unlocked—and respondents are asked to indicate how they would react. The four options in this case represent distinct responses, such as to confront the transgressor, or dismiss the incident as commonplace. As with the behaviour-type scenarios, each response contained an element of non-compliance or an implicit cost. While it may seem like confronting a transgressor is an obvious right answer, there is in fact a high social cost associated with doing so (we find for instance that typically security-conscious individuals are regarded as paranoid by their peers).

The answers here were linked not to behavioural risk types but to a model of cultural maturity that has been developed in support of this work. The model considers security competence relative to an individual’s business tasks. Other works describe the need for competence in repeatable tasks which can form good security habits (Thomson and von Solms 2006). Here, we also consider the capacity to embody policy (where it is clear) and adapt it to new or complex situations that require a conscious response, a distinction that has been explored by Reason in the realm of safety (Reason 2008).

Our model contains a series of levels (see Appendix A.5) which attempt to articulate the maturing relationship between the individual and security policy. Those at the lower levels engage with security only as absolutely necessary, while those at the higher levels champion security in their local environment. The levels linked to the answers in the survey are as follows:

**Level 1: Uninfluenced** Is not engaged with security in any capacity.

**Level 2: Technically Controlled** Follows security policy only when forced to do so by external controls.

**Level 3: Ad-hoc Knowledge and Application** Understands that a policy exists and follows it by rote.

**Level 4: Policy Compliant** Has internalised the intent of the policy and adopts good security practises even when not specifically required to.

**Level 5: Active Approach to Security** Champions security to others and challenges breaches in their environment.

Although the model includes a level 1, practically speaking individuals at this level will not be found in an organisational environment, as there is typically infrastructure in place that at the least requires employees to have a registered username and password to facilitate access to IT resources. As such our survey utilises level 2 and upwards. Level 2 assumes that compliant behaviour must be imposed upon individuals to ensure that routine tasks remain secure, and so in turn the IT infrastructure constrains behaviour. This is analogous to ‘basic hygiene’ as described by (Stanton et al. 2005), acting to manage the ‘dangerous tinkering’ and ‘naive mistakes’ which might otherwise happen. However, organisational security is complex and technology-based solutions alone cannot anticipate and manage all situations. Level 3 assumes that employees have enough security knowledge to make some in-situ decisions, whereas toward level 5 employees know enough to apply their knowledge and skills to unforeseen situations, as well as to articulate workable solutions to those around them.

The maturity levels in Company B are in relation with the policy, whereas the attitude types in Company A relate to the security apparatus.

### 6.4.5.3 Scenario selection and distribution

For each company, 8–10 scenarios were created. However, for several reasons we did not wish for each participant to complete the entire set of scenario questions. First, it would be too time-consuming; our partner companies were generally concerned about the productivity impact of large-scale data collection involving any number of employees over a short period, and so we wished for our scenario survey to be completed in 10–15 minutes. Second, as the scenarios were tailored to specific topics not all of them would be relevant to all parts of the business. For example, giving a question about a retail environment to an engineering division will yield data based on guesswork rather than experience. During deployment, we would request a range of demographic information, including business role (the options for which were drawn from the company’s structure) at the beginning of the survey, then deploying a subset of 3–4 scenarios to each participant based on their responses. This is an example of where it is important, and necessary, to engage with a partner



**Boxout 1: Hotspots**

We refer to instances of ranking scores being positively correlated with severity rating scores, or negatively correlated with acceptability rating scores, as ‘hotspots’. Where these correlations are detected, it indicates that participants are favouring the use of options that they know carry high risk and which represent unacceptable forms of behaviour. Hotspots represent significant areas of concern for the company, as they show areas in which employees report that they have to choose (knowingly or unknowingly) insecure practices.

organisation at the right level, to ensure that survey tools can be managed in a way that fits naturally with activities within the business.

**6.4.6 Survey tasks scoring**

As discussed in [Sections 6.4.4](#) and [6.4.5](#), each of the options given with each scenario was linked to an attitude type, a behaviour type or a maturity level. In order to determine the prevalence of each categorisation, a scoring method linked to the ranking task was used. The position of the option in the ranking task determined the score of the associated type. This score was cumulative over the scenario questions. For example, if the first option was linked to Behaviour type then the ranking of this option determined the score given to type 1. The scoring was as follows:

**Rank 1:** 4 points

**Rank 2:** 3 points

**Rank 3:** 2 points

**Rank 4:** 1 point

As each participant answered a maximum of two behaviour and two attitude questions, these scores were normalised for each participant, enabling statistical analysis. This scoring system was also used to determine the popularity of the scenario options themselves.

**6.4.7 Selecting interventions**

Organisations may enact ‘interventions’ in order to influence a change in the regular security behaviours of employees. The rich picture of employee behaviours and attitudes provided by the interview and survey process means that the methodology described in this paper can support a more systematic and informed approach to the identification of interventions. While this paper does not cover the outcomes of this step in detail (the right-hand side of [Figure 6.1](#)), the intended use of the data and survey results are included here in the interests of completeness.

The interview and survey results provide security managers with information on the most pressing problems—the ‘hotspots’ (see [Boxout 1](#))—encountered by employees in their own company’s IT environment, as well as an idea of the factors that underpin an issue. Interventions can then be targeted at the motivating factors of an issue, rather than the symptoms or the elements of it which most relate to specific regulatory expectations. It is intended that researchers would engage with the organisation to determine which category of intervention is optimal, drawing on system (re)design, awareness and training, or technical controls where appropriate. Means of addressing tensions between security and productivity are further discussed in [Beautement, Sasse and Wonham \(2008\)](#). Having some sense of the scale of a policy hotspot is also useful (how many employees regularly enact an insecure behaviour), as organisations can then invest resources proportionally, and crucially consider the intended scale of an intervention to ensure that it is properly implemented and does not introduce problems of its own (for instance by updating only a subset of the awareness materials which employees are directed to use, which can in turn introduce inconsistencies).

### 6.4.8 Research ethics and data handling

The study successfully went through an ethics approval process at UCL (approval number: 3615/002) and was registered with the Data Protection Act (registration number: Z6364106/2012/11/08). We had a written agreement with management, which was distributed with the recruitment email, that employees would not face negative consequences for policy violations they reported. The audio-recordings were transcribed by an external company under NDA. Transcripts were redacted to remove any identifying information such as names of people and locations. The original audio recordings were deleted.

## 6.5 Survey results

In this section I will present the survey results from both companies. The analysis was carried out as a number of factor analysis, this being a statistical analysis of the various Appropriateness, Acceptability and Severity scores and Behaviour and Attitude rankings grouped by Business divisions, Age groups, employee types and Business locations. In order as not too flood this section with tables, only the analysis by Business divisions will be shown in this section ([Tables 6.2](#) and [6.4](#)). The remaining tables can be found in the Appendix ([Tables A.1](#) and [A.2](#)).

## 6.5.1 Company A

### 6.5.1.1 Participant demographics

Business Division	<25	25–29	30–34	35–39	40–44	45–49	50–54	≥55	Sum
Asset Management	19	37	25	21	29	38	20	18	207
Commercial	2	4	5	7	6	3	3	1	31
Construction	3	7	7	11	11	17	10	8	74
Customer	4	13	8	7	13	21	10	8	84
Customer Operations	11	10	15	11	9	11	3	7	77
Finance	3	14	22	19	22	26	6	6	118
Global IS	0	7	9	10	18	30	12	8	94
Head Office	1	9	10	5	14	29	12	14	94
Legal	2	3	12	2	7	10	1	2	39
Network Operations	9	10	15	14	25	21	4	5	103
Network Strategy	8	13	14	10	7	18	14	7	91
Operations	5	16	17	12	24	64	18	29	185
Other	6	7	7	10	9	13	7	4	63
Shared Services	1	2	2	3	7	13	6	5	39
Location									
Other	8	33	49	43	55	82	24	23	317
UK	49	70	74	72	82	103	42	39	531
US	17	49	45	27	64	129	60	60	451
Type									
External, Offsite	2	2	6	2	5	12	2	1	32
External, Onsite	5	11	18	10	27	23	10	11	115
Internal, Offsite	2	6	14	12	22	27	14	8	105
Internal, Onsite	54	109	106	96	104	179	66	74	788
Other	11	24	24	22	43	73	34	28	259
Sum	74	152	168	142	201	314	126	122	1299

**Table 6.1: Demographics for Company A** Number of participants by business division, location, employee type and age groups

Table 6.1 displays the demographic distribution of our participants in Company A. We break up the business divisions, location and types by age groups.

Similar tables for analysis by age groups, location and employee category can be found in Table A.1 in Appendix A.1.

#### 6.5.1.2 By business division

Table 6.2 displays the results of the factor analysis by business division. The attitude type questions are analysed in Table 6.2(a), the behaviour type questions are analysed in Table 6.2(b).

For the attitude type rankings the responses across the business divisions are uniform: attitude type 2 is ranked highest, followed by types 3, 4, and 1. The differ-

Business Division	Attitude 1		Attitude 2		Attitude 3		Attitude 4		$\tau$
	Rank*	Approp	Rank**	Approp	Rank	Approp	Rank**	Approp**	
Asset Management	-0.04	-0.01	-0.05**	-0.01**	-0.00**	0.03**	0.09**	0.13**	0.71**
Commercial	0.23**	-0.06	-0.12**	-0.19*	0.08**	-0.03**	-0.19**	-0.04**	0.67**
Construction	-0.01	-0.03	-0.12**	-0.05**	0.09**	0.08**	0.04**	0.10**	0.74**
Customer	-0.06	-0.02	0.03**	0.03**	-0.03**	-0.17**	0.06**	-0.11**	0.68**
Customer Operations	-0.04	0.01	0.08**	-0.03**	0.10**	0.11**	-0.14**	-0.29**	0.73**
Finance	0.03	-0.03	0.07**	-0.04**	-0.05**	-0.15**	-0.05**	0.08**	0.73**
Global IS	0.04	-0.05	-0.00**	0.10**	0.09**	0.23**	-0.13**	-0.13**	0.74**
Head Office	-0.03	-0.01	0.03**	-0.07**	-0.03**	-0.10**	0.03**	0.06**	0.71**
Legal	0.00	-0.12	0.05**	0.02*	0.13**	0.10**	-0.18**	-0.19**	0.78**
Network Operations	-0.01	-0.02	0.04**	0.10**	-0.07**	-0.03**	0.04**	0.11**	0.76**
Network Strategy	0.04	0.07	0.06**	-0.03**	-0.06**	0.02**	-0.04**	-0.09**	0.71**
Operations	0.00	0.04	-0.07**	0.02**	-0.01**	-0.02**	0.08**	0.05**	0.72**
Other	0.06	0.05	-0.03**	-0.02**	0.00**	-0.02**	-0.04**	-0.09**	0.72**
Shared Services	-0.01	0.15	0.13**	0.07**	-0.10**	0.01**	-0.02**	-0.05**	0.70**
mean	1.27	1.28	3.58**	4.58**	2.92**	4.13**	2.23**	2.56**	0.72**

(a) Attitude type rankings and Appropriateness score split by business division.

**Table 6.2: Factor Analysis in Company A** The values in each cell of the tables above describe the variation from the mean in their column, with the mean being shown at the bottom (the mean is the value for the organisation as a whole). Based on the scoring in Section 6.4.6, higher ranks imply more popular choices. Similarly, the higher the Approp/Sev score, the more appropriate/severe the participants take the option to be. In the second row, the \*\*/\* after Rank/Approp/Sev show statistical significant variations from the median rank or acceptability or severity score respectively based on the Kruskal-Wallis H-test for independent samples at  $p < 0.01/p < 0.05$  confidence respectively. If this Kruskal-Wallis test shows statistical significance, for each subgroup a two-sided Mann-Whitney rank test between this subgroup and the union of all other subgroups is carried out; the results of these tests are shown by further \*\*/\* at each number, showing statistical significance at  $p < 0.01/p < 0.05$  confidence respectively. **Caption continued below Table 6.2(b) on Page 127.**

ences between these ranks are statistically significant. In terms of Appropriateness scores, the results are similar with the same ranking shown as for the attitude types themselves.

The individuals in this organisation do care about security, but do not take personal responsibility for threats. Not a single employee group consider direct action through official channels as the preferred option.

The behaviour types are analysed by business division in Table 6.2(b). While the organisation in its entirety ranks behaviour type 4 highest, followed by types 2, 3 and 1, a number of divisions also rank other behaviour types equally highly. In terms of severity scores, Behaviour 1 is ranked most severe. The other three behaviour types have severity scores that indistinguishable at an organisational level. Nevertheless, the severity scores are a weak but statistically significantly inversely correlated to

Business Division	Scenario	Behaviour 1		Behaviour 2		Behaviour 3		Behaviour 4		$\tau$
	Accept**	Rank	Sev*	Rank	Sev	Rank	Sev**	Rank**	Sev**	
Asset Management	0.18*	-0.00	0.02**	-0.11**	0.07**	-0.09**	0.18**	0.20**	-0.09	-0.31**
Commercial	0.07	-0.11	0.23**	-0.01	-0.31	0.13**	0.03	-0.01**	0.08*	-0.27**
Construction	0.07	-0.08	-0.05**	-0.03**	-0.04**	-0.06**	0.07	0.17**	-0.32**	-0.39**
Customer	-0.09	0.19	-0.05**	0.01**	-0.06	-0.09	-0.03	-0.12**	-0.14	-0.24**
Customer Operations	-0.32*	0.13	-0.15**	-0.06	-0.02*	0.11**	-0.27**	-0.18**	0.23*	-0.13**
Finance	-0.27*	0.02	0.00**	0.03**	-0.08	-0.06**	-0.14	0.01**	0.13	-0.27**
Global IS	-0.00	-0.03	0.23**	0.11**	0.03	0.05**	0.15	-0.13*	0.35**	-0.26**
Head Office	-0.16	-0.11	-0.15**	0.05*	-0.01	0.14**	-0.22*	-0.07**	-0.01	-0.29**
Legal	-0.43**	-0.04	0.09**	-0.05	0.11	0.02**	-0.25	0.07**	0.02	-0.21**
Network Operations	0.13	-0.09	0.09**	0.18**	-0.08	-0.04**	-0.06	-0.05*	0.02	-0.30**
Network Strategy	-0.08	0.05	0.01**	-0.02*	-0.03	0.06**	-0.14	-0.08**	0.09	-0.27**
Operations	0.19*	-0.01	-0.09**	-0.03**	0.10**	0.03**	0.16*	0.02**	-0.20*	-0.31**
Other	0.06	-0.12	0.08**	0.07*	0.07	0.03**	0.03	0.02**	0.19	-0.29**
Shared Services	0.18	0.30	-0.04**	-0.08	-0.09	-0.05	-0.05	-0.17*	0.06	-0.22**
mean	2.43	1.93	4.35**	2.62**	3.51	2.34**	3.52	3.11**	3.46	-0.28**

(b) Behaviour types rankings and behaviour severity score split by business division.

**Table 6.2: Continued:** Further, the colours show the order of mean Rank/Approp/Sev for each of the groups (i.e., ranking them horizontally). The largest mean is given the darkest colour, and the colour changes to a lighter shade if there is a statistically significant difference between the distribution of ranks/scores of the current mean and the next largest mean, based on a one-sided paired Wilcoxon rank test. This statistical test is further shown by \*\*/\* at the value of the higher cell, showing  $p < 0.01/p < 0.05$  confidence respectively. If more than one cell contains the same colour, there is no statistical significant variation between the rankings/scores for these options.

Lastly, the rightmost column  $\tau$  lists Kendall's  $\tau$  correlation coefficients between the rank and the acceptability/severity score respectively for each of the groups. Kendall's  $\tau$  ranges from  $-1$  (perfect anti-correlation) to  $1$  (perfect correlation). \*\*/\* signifies rejecting the null hypothesis of independence (i.e.  $\tau = 0$ ) with statistical significance at  $p < 0.01/p < 0.05$  confidence respectively.

the ranking.

Overall, the employees of this organisation claim to remain secure whenever possible. Yet, behaviours where employees show a minor priority for work (type 2), or expect others to ensure security (type 3) are not described as any more severe. The mean Behaviour type scenario is ranked as slightly unacceptable of failing to complete the task (score of 2.43, where a score of 1 is *not acceptable at all*, and a score of 5 is *very acceptable*). There are minor, statistically significant, variations in this score between divisions: in particular, the *Legal* division considers the scenario less acceptable to be failed (2.0), while the *Operations* division considers it more acceptable (2.62).

### 6.5.1.3 By age groups

There is no discernible difference between any of the age groups and their ranking of either Attitude or Behaviour type, apart from the *Under 25* group who rank behaviours types 2 and 3 as equally likely.

### 6.5.1.4 By business location

There are minor differences between the two (US/UK) business locations in this company. While there are no differences in the overall ranking of the attitude or behaviour types, the US ranks Attitude type 3 (*'Take direct action through official channels'*) statistically significantly higher than the UK. In terms of Behaviour types, The UK considers type 4 (*'Take direct personal action against the threat'*) statistically significantly more important than the US employees. This is in line with the scenario's score for the acceptability of failing to complete the task: the UK considers it more acceptable than the US.

### 6.5.1.5 By employee type

The differences between the employee types are more pronounced. The company splits here between *Onsite* and *Offsite* employees; regardless if they are *External* or *Internal*. *Offsite* employees rank Attitude types 2 and 3 equally, whereas the rest of the organisation ranks  $2 > 3$ . This division is also noticeable in the Appropriateness score, where *External*, *Onsite* employees rank Attitude types 2 and 3 as equally appropriate, and *External*, *Offsite* employees even switch signs and rank type 3 as more appropriate than type 2.

## 6.5.2 Company B

Akin to our analysis for Company A, this section conducts a similar set of factor analysis for Company B.

### 6.5.2.1 Participant demographics

In total, 641 complete survey responses were recorded. The briefing document informed participants that any surveys completed in less than 5 minutes (minimum reading time in our pilot study) would not be included, which left us with 608 responses for analysis.

For the purposes of our study, the company is split up across 7 business divisions as well as a number of locations. The majority of responses originated from the *Sales & Services* (292), followed by *Operations* (152). The remaining divisions

Business Division	< 25	25–29	30–34	35–39	40–44	45–49	50–54	≥ 55	Sum
Business	2	3	6	5	10	4	1	2	33
Finance & Prof. Services	1	5	8	13	14	4	2	0	47
Human Resources	1	2	3	5	2	1	0	0	14
Marketing & Consumer	0	3	4	5	7	5	0	0	24
New Business	0	2	1	3	2	2	0	1	11
Operations	9	6	20	25	28	32	21	11	152
Other	1	6	10	10	3	3	1	1	35
Sales & Service	39	76	72	39	28	20	11	7	292
<b>Location</b>									
1: HQ	8	9	18	29	24	17	9	4	118
2	3	9	10	10	3	2	0	1	38
3	0	11	14	7	10	6	3	1	52
4	1	10	11	5	6	5	3	2	43
5	7	18	26	29	23	17	6	4	130
Homeworker	0	1	8	8	11	8	8	4	48
Minor offices	0	1	2	1	4	5	4	4	21
Other	34	44	35	16	13	11	3	2	158
Sum	53	103	124	105	94	71	36	22	608

**Table 6.3: Participant demographics in Company B** Number of participants by business division, location and age groups

were all significantly smaller, ranging from 11 to 47 responses. Participation was more equally divided between the business locations surveyed, with locations *1: HQ* and *5* (a large regional office) being the largest ones with 118 and 130 responses respectively. Further we analyse trends across 8 age groups. Survey respondents were approximately normally distributed across the age groups, with the age group 30–34 representing the largest share with 124 participants. The edge cases of < 25 and ≥ 55 were nonetheless sufficiently large with 53 and 22 responses respectively, to allow for potentially statistically significant results across all age groups. An full view of the demographic statistics can be seen in [Table 6.3](#).

The number of responses were sufficient to allow a factor analysis by business division ([Section 6.5.2.2](#)), age group ([Section 6.5.2.3](#)) and location ([Section 6.5.2.4](#)), with 8 factors each. A full factor analysis with 512 factors is outside the scope of this methodology at present as it would a sample size several orders of magnitude larger.

Each participant responded to at least one maturity type scenario ([Section 6.4.5.1](#) and [Appendix A.4](#)), by ranking the four options presented in order of preference as well as assigning an acceptability score on a Likert scale to each option (see [Section 6.4](#)). A comprehensive statistical analysis was then carried out on these responses, with the results in [Table 6.4](#). In total, three such tables have been produced, but only the first one is shown here for brevity. The remaining diagrams

Business Division	Level 2		Level 3		Level 4		Level 5		$\tau$
	Rank**	Accept**	Rank**	Accept**	Rank**	Accept**	Rank**	Accept**	
Business	0.33**	0.29*	-0.23*	-0.32*	0.10	0.33**	-0.21**	-0.14*	0.61**
Finance & Prof. Services	0.37**	0.24**	-0.12*	-0.23*	0.13*	0.37**	-0.38**	-0.34**	0.59**
Human Resources	0.23**	0.18	-0.21	-0.47	0.14	0.20**	-0.16**	-0.15	0.66**
Marketing & Consumer	0.43**	0.37**	0.12*	-0.15	-0.10	0.29	-0.45**	-0.42**	0.50**
New Business	0.33*	0.32	-0.32*	-0.23	0.19	0.47**	-0.21*	0.04	0.55**
Operations	0.23*	0.14	0.25**	-0.34**	-0.13**	-0.23**	-0.34	0.06**	0.45**
Other	0.30**	0.30**	-0.33**	-0.42**	0.25**	0.27*	-0.22**	-0.38**	0.65**
Sales & Service	-0.31**	-0.23**	-0.03**	0.34**	-0.00**	-0.06**	0.34**	0.12**	0.76**
mean	1.48	1.50	2.13**	2.19**	3.08**	3.98**	3.30**	4.51**	0.62**

(a) Maturity level rankings and acceptability score split by business division.

**Table 6.4: Factor Analysis in Company B** The values in each cell of the tables above describe the variation from the mean in their column, with the mean being shown at the bottom (the mean is the value for the organisation as a whole). Based on the scoring in Section 6.4.6, higher ranks imply more popular choices. Similarly, the higher the Accept/Sev score, the more acceptable/severe the participants take the option to be. In the second row, the \*\*/\* after Rank/Accept/Sev show statistical significant variations from the median rank or acceptability or severity score respectively based on the Kruskal-Wallis H-test for independent samples at  $p < 0.01/p < 0.05$  confidence respectively. If this Kruskal-Wallis test shows statistical significance, for each subgroup a two-sided Mann-Whitney rank test between this subgroup and the union of all other subgroups is carried out; the results of these tests are shown by further \*\*/\* at each number, showing statistical significance at  $p < 0.01/p < 0.05$  confidence respectively. **Caption continued below Table 6.4(b) on Page 131.**

are included in the the appendix (see Appendix A.2, Table A.2).

The last line of Table 6.4(a) shows the full company’s maturity level properties (please refer to the caption of Table 6.4 for the details of the statistical analysis carried out). The ranking and acceptability score of each of the maturity levels are all statistically significantly separated and increasing with the maturity score. Level 5 has an average rank of 3.30 and acceptability score of 4.51. These ranks are high: a perfect score would represent an average maturity rank of 1, 2, 3 and 4 for levels 2 to 5 respectively. Further, there is a strong positive correlation between rank and acceptability score: the more acceptable the option, the more likely the participant is to choose it.

The answer options of the four behaviour scenarios map to the four behaviour types. The participants were asked to rank the options in the order they would consider enacting them themselves as well as assign a severity score on a Likert scale to each answer option and to the scenario in general. The statistical analysis that follows is similar to the analysis of maturity levels described above. Again, we show only one analysis table here for brevity (Table 6.4(b)), the remaining tables



Business Division	Scenario Sev <sup>**</sup>	Individualist		Egalitarian		Hierarchist		Fatalist		$\tau$
		Rank <sup>**</sup>	Sev <sup>**</sup>	Rank <sup>**</sup>	Sev <sup>**</sup>	Rank <sup>**</sup>	Sev <sup>**</sup>	Rank <sup>**</sup>	Sev <sup>**</sup>	
Business	0.52 <sup>**</sup>	0.17	-0.10 <sup>**</sup>	0.59 <sup>**</sup>	-1.18 <sup>**</sup>	-0.53 <sup>*</sup>	0.29	-0.23	0.10	-0.22 <sup>**</sup>
Finance & Prof. Services	0.38 <sup>**</sup>	0.34 <sup>**</sup>	-0.24	0.50 <sup>**</sup>	-0.76 <sup>**</sup>	-0.67 <sup>**</sup>	0.29	-0.16	0.15	-0.19 <sup>**</sup>
Human Resources	0.62 <sup>*</sup>	0.53 <sup>*</sup>	-0.13	-0.09	-0.90 <sup>**</sup>	-0.16	-0.41	-0.29	-0.73	0.08
Marketing & Consumer	0.84 <sup>**</sup>	0.32	-0.74 <sup>**</sup>	0.86 <sup>**</sup>	-0.84 <sup>**</sup>	-0.47 <sup>*</sup>	0.34 <sup>*</sup>	-0.71 <sup>**</sup>	0.18	-0.22 <sup>**</sup>
New Business	0.58	0.59 <sup>*</sup>	-0.58	0.62	-0.39	-0.80 <sup>*</sup>	-0.11	-0.41	0.01	-0.39 <sup>**</sup>
Operations	0.03 <sup>**</sup>	-0.02 <sup>**</sup>	-0.33 <sup>**</sup>	-0.34 <sup>**</sup>	0.38 <sup>*</sup>	-0.40 <sup>**</sup>	-0.39 <sup>**</sup>	0.70 <sup>**</sup>	-0.96 <sup>**</sup>	-0.48 <sup>**</sup>
Other	0.24	0.35 <sup>**</sup>	-0.46 <sup>**</sup>	0.04	-0.53 <sup>**</sup>	-0.26	0.11	-0.13	-0.16	-0.28 <sup>**</sup>
Sales & Service	-0.28 <sup>**</sup>	-0.18 <sup>**</sup>	0.37 <sup>**</sup>	-0.06	0.25 <sup>**</sup>	0.48 <sup>**</sup>	0.10 <sup>*</sup>	-0.24 <sup>**</sup>	0.50 <sup>**</sup>	-0.17 <sup>**</sup>
mean	2.24	2.68 <sup>**</sup>	3.49	2.02	3.76 <sup>**</sup>	2.80 <sup>*</sup>	3.20	2.50 <sup>**</sup>	3.44 <sup>**</sup>	-0.20 <sup>**</sup>

(b) Behaviour types rankings and behaviour severity score split by business division.

**Table 6.4: Continued:** Further, the colours show the order of mean Rank/Accept/Sev for each of the groups (i.e., ranking them horizontally). The largest mean is given the darkest colour, and the colour changes to a lighter shade if there is a statistically significant difference between the distribution of ranks/scores of the current mean and the next largest mean, based on a one-sided paired Wilcoxon rank test. This statistical test is further shown by <sup>\*\*</sup>/<sup>\*</sup> at the value of the higher cell, showing  $p < 0.01/p < 0.05$  confidence respectively. If more than one cell contains the same colour, there is no statistical significant variation between the rankings/scores for these options.

Lastly, the rightmost column  $\tau$  lists Kendall's  $\tau$  correlation coefficients between the rank and the acceptability/severity score respectively for each of the groups. Kendall's  $\tau$  ranges from  $-1$  (perfect anti-correlation) to  $1$  (perfect correlation). <sup>\*\*</sup>/<sup>\*</sup> signifies rejecting the null hypothesis of independence (i.e.  $\tau = 0$ ) with statistical significance at  $p < 0.01/p < 0.05$  confidence respectively.

(Tables A.2(b) and A.2(d)) are included in Appendix A.2.

The last line of Table 6.4(b) shows the analysis of behaviour types for the company as a whole. The ranking of the behaviour types is Hierarchist (2.80 mean ranking), Individualist (2.68), Fatalist (2.50) and Egalitarian (2.02). All pairwise differences are statistically significant (see Table 6.4(b)). The ranking of the severity of the options for each to the behaviour types is less clear as they can only be divided into 3 statistically distinguishable categories (as indicated by the use of three shades of colour only), although the Egalitarian option is seen as statistically significant most severe at 3.76. Further, there is a statistically significant negative correlation between severity score and behaviour type, implying that the employees rank less severe options higher, as may be expected.

It should be noted that there is no inherent ordering between the behaviour types (as it was the case between maturity levels), hence when analysing the data and Table 6.4(b), care has to be taken not to infer a ranking of the types themselves relative to each other, but rather work with the ranking of the types by the participants.

While at the level of the whole company there are is a statistically significant ordering of the preferences of the behaviour types, this changes considerably when analysing across different subgroups as discussed in [Sections 6.5.2.2 to 6.5.2.4](#), where there are in many cases only 2 statistically different groups.

### 6.5.2.2 By business division

[Table 6.4\(a\)](#) illustrates the relationship between maturity levels and business divisions. The data is shown in terms of variations from the company's mean in order to facilitate comparisons across the business divisions. There are a number of interesting deviations from the organisational mean. Only the *Sales & Services* division ranks maturity level 5 statistically significantly above level 4, where the *Finance & Prof. Services* division ranks maturity level 4 highest and statistically significantly higher than level 5. The participants from the other divisions did not discriminate between level 4 and 5 options. Participants from *Sales & Services* opted for responses corresponding to level 5 statistically significantly more often than any other division in the organisation with a mean level 5 rank of 3.64.

The acceptability scores demonstrate a similar trend. Only *Operations* and *Sales & Service* discriminated between level 4 and 5. Yet none of the divisions inverted the ranking. *Operations* are noteworthy since they clearly distinguished between level 2 and 3 maturity as well as acceptability scores.

Switching over to the behaviour types in [Table 6.4\(b\)](#), in the *Business* division the Egalitarian and Hierarchist are ranked statistically significantly higher and lower, respectively. This is also the case in *Finance & Prof. Services*, but foremost the Individualist type is ranked highest here. *Marketing & Consumer* also agrees on the Egalitarian and Hierarchist differences, but here the Fatalist option is statistically significantly lower ranked than in the organisation as a whole. *Operations* are by far the most Fatalist: they rank this option statistically significantly highest and Egalitarian lowest, and are also much less Egalitarian and Hierarchist than the organisation generally. The *Sales & Services* team agree with the organisational ranking of the types, but they gave a significantly higher score to the Hierarchist option than any other division by at least 0.64.

The *Human resources* division represents the first Hotspot (see [Boxout 1](#)), as the division shows a non-negative correlation between the option's severity score and rank. This implies that employees choose which option to prefer independent of the severity they assign to that option.

Analysing the severity scores, the *Operations* division is alone in perceiving a full ordering of the options, ranking the fatalist score third most severe and statistically

significantly much less severe than the rest of the organisation. This is in stark disagreement with *Sales & Services*, who perceive the Fatalist option much more severely with a ranking difference of 1.44.

### 6.5.2.3 By age groups

There are three age groups that did not discriminate between level 4 and 5 maturity levels: 35–39, 50–54 and 55+. The 35–39 group also shows statistically significant lower average level 5 rank than the other age groups, but ranks level 4 statistically significantly higher than the other age groups. All age groups ranked the acceptability of the options according to the maturity levels.

There are no statistically significant variations between the different age groups for the Individualist and Egalitarian behaviour types. All the differences occur when considering the Hierarchist and Fatalist types: both the age groups 25–29 and 30–34 are statistically significantly more Hierarchist than all other age groups. The age group 50–54 shows the opposite, they are significantly less Hierarchist. When examining the Fatalist type, the picture changes: the 30–34 group is significantly less Fatalist, the 50–54 and the 55+ are significantly more so. In fact the 50–54 group rank Fatalist highest, followed by a statistically significant difference by the Hierarchist, an opposite order to the organisation at whole and unique to this group. It is interesting to note that the middle three age groups from 35 to 49 (as well as the under 25 group) have little or no preference between the behaviour types and also rank them nearly equally on the severity scales.

Between the different age groups there are no statistically significant variations of the severity scores for any of the behaviour types.

### 6.5.2.4 By location

Responses from location 1: *HQ* rank maturity level 4 higher than level 5 as well rank level 2 significantly higher than all other locations. This is also evident in the acceptability score: level 5 is perceived as statistically significantly less acceptable and level 2 as more acceptable than at all other locations. Employees at locations 4, 5 and at minor offices were unable to distinguish between levels 4 and 5. Location 3 achieved the highest average level 5 rank of 3.6, statistically significantly higher than the average.

Acceptability scores only varied significantly for staff at the minor offices, which collectively scored level 5 with an extremely high score of 4.91. Further, the level 3 score was significantly lower, with a mean of 1.48, making it indistinguishable from level 2's score.

The predominant behaviour types vary widely by business location. Both locations 1: *HQ* and *Homeworker* rank the Individualist options highest, in the case of 1: *HQ* because it ranks the Individualist and Hierarchist types statistically significantly higher and lower than the other locations, respectively.

Locations 5 and *Minor Offices* rank Fatalist first; this is followed by a statistically significant lower score for the Hierarchist type at these locations. Locations 2, 3, 4 and *Other* show an opposing trend, ranking the Hierarchist type higher than other locations and the Fatalist type lower. It is worth noting that the *Other* category represents mostly retail workers spread across the company's various sites.

Interestingly, there are also a large number of statistically significant variations in the severity scores, with all four types rejecting the null-hypothesis of equal distribution of the Kruskal-Wallis test. This is also reflected in strong variations in the severity score of the behaviour scenarios across the locations. Employees at location 3 saw all four options as significantly more severe, increasing the severity scores of each option by over 20%, but the scenario's severity score remains unchanged. The opposite effect is portrayed by *Homeworkers*, who rate the scenarios 0.51 more severe than the average, but show no variations for any of the behaviour type severity scores.

There is also a second hotspot present in this comparison: location 2 shows no statistically significantly negative correlation between severity scores and rank, implying that employees at this location choose which option to take independent of the severity they assign to the option.

## 6.6 Discussion of survey results

This research applied a scenario-based survey to assess both security maturity levels and self-reported security behaviours, and employee understanding of how risky certain behaviours are. A statistical analysis of the results of the survey conducted at the two companies allows us to draw several key conclusions regarding the security culture within the organisations. In line with the existing literature, we found that assessing attitudes provides a solid approach to understanding how employees interact with security policy. However, the scenario-based survey approach allows us to go further and detect intra-population differences within the organisation, showing that there are significant differences between different employee groups in how they respond to security-related challenges in the workplace. The salient outcomes are discussed below.

### 6.6.1 Company A

Company A is extremely homogeneous. The responses point to behaviours that are split between external and internal employees, as well as minor differences between the US and UK. The majority of the organisation displays a *party line* behaviour: the prevalent behaviour intention is type 4 (Tries to remain secure wherever possible), yet the actual attitude appears to be type 2 (Report suspicions but take no direct action). This points at employees that aim for compliance, but don't actually actively participate in good organisational security.

### 6.6.2 Company B

Our analysis of the survey has shown that Company B in general has a very positive security posture: the majority of employees are at maturity level 5 and there is a downwards gradient of the ranking of the lower maturity levels. This combines well with a founded understanding of the acceptability and severity of the options presented to the employees of the organisation; employees in general choose what are in their opinion the more acceptable and less severe options. This strength is based on the willingness of the majority of the organisation's employees to engage actively with security. The predominant attitude within the company is to adopt good security practices, even when not specifically required to by technology or policy prescriptions. Many members of the organisation reported that they would challenge any breaches of policy they observe in their environment, with older employees being less likely to do so. Where friction exists between the business and security processes, employees take a predominantly Individualist approach to conflict resolution, meaning they rely on their own skills and knowledge. This echoes the results of Rhee, Kim and Ryu (2009) and Siponen, Pahlila and Mahmood (2007) who both recognise the role of self-efficacy in decision making. Individually-derived approaches to security, driven by personal perception of what constitutes secure practice, can also manifest when policy and support is not known or visible to the individual (Kirlappos, Parkin and Sasse 2015).

The ranking of the behaviour types is also positive, but the differences in their respective rankings are weaker. Hierarchists are unlikely to challenge the existing structures, and while they may follow security policies to the letter, the Individualist that innovates may identify and solve new challenges before they become problematic (Kirlappos, Parkin and Sasse 2015). Some CISOs might think that it is desirable if all employees were Hierarchists, but it could be argued that it would be counter-productive for an organisation to be exclusively one behaviour type, as there are many benefits in diversity. From a productivity point of view, the or-

organisation requires diversity and even from a security point of view, variation has benefits. A diverse mix of behaviour types may even be essential, as security issues are embedded in, and deeply influenced by, social context such as corporate and national culture (Pfleeger, Sasse and Furnham 2014). In this sense, these issues have to be understood and addressed before any successful intervention program can be introduced.

These points will be revisited in [Section 6.8](#), where we reformulate the organisations champions of security and culture change based on the data gathered in this section.

### 6.6.3 Comparisons between Companies A and B

The two companies cannot be compared directly; as the different surveys did not measure the same aspects. The methodological differences however are interesting. Where we elicited in A only the potential actions and responses of employees in one specific set of circumstances, the more abstract behaviour types allow for extrapolation of behaviours in general, and indeed I will utilise this generalisation for my definition of security champions in [Section 6.8](#).

### 6.6.4 Limitations

Each engagement with an organisation is time-consuming, involving interviews which are used to generate scenarios specific to the organisation. We envisage that the cost will decrease with further iterations of the methodology, but may present a high barrier of entry. Especially in organisations business processes, products and structures may change very quickly. This then requires an iteration of the methodology to be executed in a short timeframe, something that we have so far struggled to do.

We would like to emphasise that from an organisational point of view however, employing our methodology is worthwhile because it creates a benchmarking tool that the organisation can use to re-evaluate and monitor over time to compare to previous iterations. As researchers working with many organisations, we envisage that the organisations where we conduct interviews yield a library of questions that we may be able to reuse for other organisations that are broadly similar. This benchmark of a company's security posture can then be used to compare responses to particular security mechanisms across organisations and sectors as a whole.

Our survey did not capture many contributing factors to the participants responses that may have helped to explain their answers. The respondents background (e.g., computer literacy, previous jobs, other relevant experience) would have

provided hints at a number of other relationships worthwhile studying, and potentially allow us to tailor interventions even more specifically. We collected free-text responses at the end of the survey that we will analyse as part of future work, they might help us shed more light on employees' reasoning and justification for their choices. Data collection does not stop once the intervention phase is reached, the methodology presented here supports decision-makers to identify broad employee categories and hotspots to target for improvements. A follow-up intervention may in itself involve data collection to identify contributing factors to particular behaviours.

## 6.7 Survey conclusion

The methodology presented here allows organisations to take steps towards empirically assessing the security culture, as well as gaining an understanding into the predominant behaviours and attitudes found within the organisation. We address the issue of scalability by deploying scenario-based surveys that employees can complete in 10–15 minutes but can therefore be deployed to a large enough fraction of the organisation to be representative. All the scenario details, and answer options, are grounded in information gathered from a series of semi-structured interviews with employees of the organisation. We demonstrate that this approach allows us to detect statistically significant differences between employee groups that can inform targeted interventions. Business area, age, and geographical location all provide axis of differentiation. Giving an organisation an understanding of these details can potentially allow them to plan their future training, communication, awareness and policy making strategies more effectively. Enabling targeted interventions that focus on particular employee groups can save employees from both being involved in non-targeted interventions and needing to determine if they apply to them. Targeted interventions are then a good step towards reducing the draw on employees' compliance budget (Beautement, Sasse and Wonham 2008).

## 6.8 Security Champions

So far in this chapter I have described a rather dry, statistical analysis of the characteristics of two organisations. Based on the data gathered, I now take a holistic view of the security culture in Company B. Organisations need employees to participate in the construction of workable security, by identifying where policies causes friction, are ambiguous, or just do not apply. However, current efforts to involve employees in security act to identify employees who can be local representatives of

policy—as with the currently popular idea of ‘security champions’—rather than as a representative of employee security needs.

Security managers in large organisations will define policies to encourage a shared approach to IT security for all members of the organisation. Policies can refer to a mix of procedures and technical controls, which employees in the organisation will interact with, and are expected to use according to the rules of the policy. The fit of these policies to employee work and business processes is often assumed to be good, but the need to involve employees in identifying oversights, bad security procedures, and grey areas in policy is under-appreciated. My work on Security Champions establishes the framework for answering Research Question 2 (Section 1.5) by giving a cross-sectional view of security behaviours in the organisation that can inform targeted interventions.

The promotion of *security champions* is seen as a way to find local representatives who can promote and monitor security policy at a local level, acting as an extension of company’s security management team (Gabriel and Furnell 2011). However, security champions may only be effective in this way if the policy itself is workable (Beris, Beutement and Sasse 2015).

This brings us to examine the role of policy effectiveness from the perspective of security usability. Employees may feel that policy is too cumbersome, that it actually asks the impossible of them, or that the relevance of security mandates to their work is unclear (Kirlappos, Beutement and Sasse 2013). This can then lead to deliberate or unwitting non-compliance, and workarounds to prescribed security processes. I look to explore the way that the organisation can engage employees, to give them a role in identifying and solving shortcomings of security policy.

In Company B, attitude to policy and behaviour types—the prevailing security cultures—vary greatly (see Section 6.5.2). I focus on four business divisions (Sales & Services, Operations, Business, and Finance & Professional Services) and examine these in further detail. As part of the survey 267 participants chose to elaborate their choices by giving additional free-text responses, not previously analysed. We combine both the two-dimensional security culture dataset and the free-text responses, conducting a novel analysis of the dependencies between maturity levels and behaviour types in the organisation.

There is a role in contributing to the effectiveness of security policies not only for those who follow policy, but also for those who question policy, socialise solutions, or expect security to justify itself as a critical part of their productive work. This demonstrates that security champions cannot be uniform across the organisation, but rather that organisations should re-think the role of security champions as diverse ‘bottom-up’ agents to change policy for the better, rather than communicators



of existing ‘top-down’ policies. My approach can identify local pockets of security expertise and indicate how to engage with those employees to create workable security solutions (Becker, Parkin and Sasse 2017a).

### 6.8.1 Free-text survey responses


A participant could provide additional comments on each question, scenario and the available options, via an associated free-text field in the survey. The security-work dilemmas and answer options were based on a large-scale interview study with employees at the same organisation, and included an element of non-compliance or an implicit cost that had been described and justified by multiple interviewees. Participants may then feel that there are other solutions available, informed by their local work environment (and which the security function may not be aware of). There was no direct incentive associated with providing additional comments; where employees provided further comment they were in effect being proactive toward security.

### 6.8.2 Source data

All 608 complete responses from Company B were analysed. The survey captured business division, but also each participant’s main place of work and age (as discussed in Section 6.5.2). We focus our analysis on business division, since the security mechanisms and rules that an employee interacts with vary with roles. Surveys were distributed to seven business divisions (where one was a group of smaller divisions) across a larger number of physical locations. The majority of responses originated from Sales & Services (292), followed by Operations (152) (see Table 6.3). The number of responses was approximately proportional to the size of each of the divisions, but we were unable to further control the sampling within each division.

We complement analysis of the combined attitude-level and behaviour-type with examination of the free-text responses. These responses also illustrate the kind of information that security managers could use in policy formulation should they involve employees, who may have differing relationships with organisation security, more directly in the process. This relies on security managers believing that employees may adopt a range of security behaviours in the workplace which are within and outside of policy, but that this may be done in response to practical challenges in fitting policy and other security mechanisms to the business more effectively (as with ‘shadow security’ behaviours (Kirlappos, Parkin and Sasse 2015)).

## 6.9 Culture analysis



**Figure 6.2: Kiviat diagram for the entire organisation** Maturity level 2 (red), 3 (blue), 4 (yellow), 5 (green) compared to the distribution of behaviour types.

We cross-analyse behaviour types with attitude levels for groups of employees. To support analysis and to visualise the data, we employ Kiviat diagrams. [Figure 6.2](#), for example, describes the distribution of behaviour types and attitude levels for the entire organisation. These diagrammatic representations expose the interplay between different behaviour types at varying levels of attitude toward policy, where together these varieties can help to strengthen the security posture of the organisation as a mapping of the reach and influence of policy. As we examine different subsets of the organisation, we can compare the security culture in each group to understand better how employees can be engaged to improve security effectiveness.

[Figures 6.2](#) to [6.6](#) all have four tetragons plotted, one for each attitude level. The diagonal axes represent the fraction of participants that exhibit each personality type, with the centre of the diagram being 0%. To aid readability, Kiviat plots are scaled to fill the entire chart. Hence, each corner of the tetragon represents the fraction of participants of that attitude level that exhibit the behaviour type of that diagonal. As an example, for the case of [Figure 6.2](#), the blue tetragon represents participants at attitude Level 3 (‘Ad-hoc Knowledge and Application’). Of these, 4.5% (4 employees) exhibit as *Egalitarian*, 36% (29 employees) as *Individualist*, 34% (27 employees) as *Fatalist* and 26% (20 employees) as *Hierarchist* behaviour types. This distribution is very different to Level 5 attitude types, where 45% (212 employees) of participants belong to the *Hierarchist* group.

The quotients of participants labelled with each behaviour type and attitude level can be found in [Table 6.5](#). While only a minority of participants exhibit attitude Levels 2 and 3, the variations between different divisions remain strong enough for detailed analysis. The Kiviat diagrams illustrate the relationship that employees may be having with the IT-security infrastructure around them, knowingly or unknowingly, as part of their working lives. These interactions may be the result, or the root cause, of their behaviour type, where security-related skills are also a mediating factor.

It is interesting to note the difference in distributions of behaviour types for different attitude levels. *Individualists* have a larger proportion of lower attitude levels, and *Hierarchists* emerge at the highest level, Level 5 ('Active Approach to Security'). Referring to the attitude levels, there is a disparity between Levels 2–3 and Levels 4–5, which immediately suggests not just that distinct approaches to employee engagement would be needed, but that the messaging would have to be crafted to match the relationship that employees have with security and security policy.

The strongest security attitude in the organisation overall ([Figure 6.2](#)) is portrayed predominantly by *Hierarchists* and *Fatalists*. The *Hierarchists* are akin to the idealised 'security champion' that Gabriel and Furnell (2011), Furnell and Rajendran (2012) and Posey et al. (2014) have described. They follow the rules and have security skills to support them (being mostly at attitude Levels 4 and 5). The limited representation of *Egalitarian* behaviours would suggest that individuals respond to security challenges in isolation (perhaps because of, or as the cause of, the aforementioned barrier to working with policy). The 'champions' the organisation may need most may then be team leaders or others who can bring people together and motivate them through social activities and interactions. These individuals do not need to have a high security attitude level: most employees are at a high attitude level already ([Table 6.5](#)).

Behaviour type	Level 2	Level 3	Level 4	Level 5
Individualist	1.6%	4.8%	10.7%	14.1%
Egalitarian	0.5%	0.7%	5.6%	6.7%
Hierarchist	1.2%	3.3%	16.8%	34.9%
Fatalist	2.0%	4.4%	10.9%	21.9%

**Table 6.5:** Absolute percentages of Behaviour Types compared with Attitude Levels

We also examine a number of specific business divisions/units in the organisation. This allows us to compare security practice in different working environments in a

large and complex organisation. Business divisions in a large organisation may differ in security culture to the point where security champions need a different set of skills and strengths to support protection of the overall organisation.

Behaviour type	Level 2	Level 3	Level 4	Level 5	Total
Individualist	2.01%	8.72%	20.81%	25.50%	42.95%
Egalitarian	3.33%	0.00%	15.00%	26.67%	36.67%
Hierarchist	0.78%	1.17%	8.98%	20.70%	24.61%
Fatalist	2.48%	4.46%	10.40%	17.82%	32.18%


**Table 6.6:** Response rates to free-text response questions by Behaviour Type and Attitude Levels

To support analysis, we also refer to free-text responses for the scenario-based questions (where this was optional for participants, as described in [Section 6.8.1](#)). The response rates for optional comments are captured in [Table 6.6](#). We see that response rates generally increase with security attitude level for all behaviour types. That *Hierarchists* and *Fatalists* make fewer comments, which is in keeping with their behaviours: not questioning rules, either because they adhere strictly to the policies or because they consider security to be ‘somebody else’s job’. When discussing specific business divisions in the following sections, we refer to free-text responses that illustrate the qualities of different kinds of approaches to organisational security. Alongside each quote the participant’s behaviour type and attitude level are included (as classified by their responses to scenario questions). Quotations allude to aspects of security that already have their champion or heroic deeds which save an otherwise unworkable situation, in turn illustrating the benefits that employee input can bring to the security culture of the organisation.

To support analysis, we focus on the four largest divisions in the organisation: Sales and Services (292 participants, [Figure 6.3](#)), Operations (152, [Figure 6.4](#)), Business (33, [Figure 6.5](#)), and Finance & Professional services (47, [Figure 6.6](#)).

### 6.9.1 Sales & Services division

[Figure 6.3](#) shows the distribution of behaviour types for each attitude level in the Sales & Services division. The starkest difference compared to the organisation as a whole (see [Figure 6.2](#)) is that approximately 63% (248 employees) of all participants from this department exhibit as *Hierarchists* (vs. 40% for the whole organisation, statistically significant with Fisher’s exact test with  $p < 0.01$ ). The free-text responses for this group included additional comments regarding a scenario where,




**Figure 6.3: Kiviat diagram for the the Sales & Service division** Maturity level 2 (red), 3 (blue), 4 (yellow), 5 (green) compared to the distribution of behaviour types.

due to IT limitations, the scenario’s protagonist is unable to securely send data to a client. As an example:

*‘The employee is put in a no-win situation. If the business permit flexible working then the only allowable option here is for the data not to be sent.’* (Hierarchist, Level 3 / ‘Ad-hoc’)

Here the employee weighs up options, leaning toward adherence to policy without compromising security. A *Hierarchist* approach is for the most part the standard security behaviour in this division. The second largest group represented at attitude Level 5 (‘Active Approach’) are *Fatalists*, with a share of 20% (52 employees), with little *Individualist* behaviour.

Considering the high representation of *Hierarchists* in this division, it may be that prescribed security behaviours may align with the context in this division, in that it has the most outward-facing customer interaction of all of the divisions, and predictable processes may be beneficial for managing those interactions. There is an extremely low representation of *Egalitarians* and *Individualists* in this division; for one this means that we cannot be sure whether security rules can be followed without impacting business opportunities. Missed business opportunities are noted elsewhere as a potential cost of being constrained by organisational security controls (Beautement, Sasse and Wonham 2008). *Egalitarians* and *Individualists* may adapt security procedures in such situations so as not to impact service.



**Figure 6.4: Kiviat diagram for the Operations division** Maturity level 2 (red), 3 (blue), 4 (yellow), 5 (green) compared to the distribution of behaviour types.

### 6.9.2 Operations division

A contrasting picture is found in the Operations division, as shown in [Figure 6.4](#). Here 55% (65 employees) of the employees at attitude Level 5 are *Fatalists*. If *Fatalists* see their own actions as irrelevant to the preservation of security, it may be that there is a separation from the larger organisation’s security function. This may be a perceived separation, or that employees perceive security as being addressed elsewhere by someone else. That these employees are at Level 5 (‘Active Approach’) implies that their security understanding, which may or may not stand alongside *policy-compliant* security, is superb. The high proportion of *Fatalists* implies a disconnect; either the role of security policy in the activities of the division is not clear and visible, or there are insufficient efforts by the security function to engage employees. The following quote, although referring to policy, also mentions ‘best practice’, which is how IT operates in reality:

*‘This question needs to be contextualised around how important the information is and how important the consequent information security policy/policy level applied is. [...] Existing best practice for teams [is] to share logins to certain systems precisely because individual logins might not be working. [...] I’m not sure what significance [a] Password Manager has, because I’m not aware of anyone in [Operations] using that facility. Most of the tools are not even [tested and approved for internal use], much less supported by something as silly as a Password Manager.’*  
(*Fatalist*, Level 3 / ‘Ad-hoc’)

The individual is resigned to working with an IT system that does not support business processes. The idea that a password manager could be beneficial to the efficiency of an employee is considered laughable, because of the perceived state of the organisation's systems. The employee's comments are useful to security managers simply for referring to IT as a larger element of the organisation, which security ought to be aligned with. Similarly, the *Fatalists* in this division may have 'seen it all', and accepted that personal involvement in maintaining secure operations can in some instances prove futile. This is emphasised by the high attitude Levels of *Fatalists* here: these employees understand the consequences of their actions, and conclude that even the most effective approaches to security still have the capacity to fail in practice.

Another employee voices their opinion in a less exasperated manner, when considering the expectation of having to share passwords:

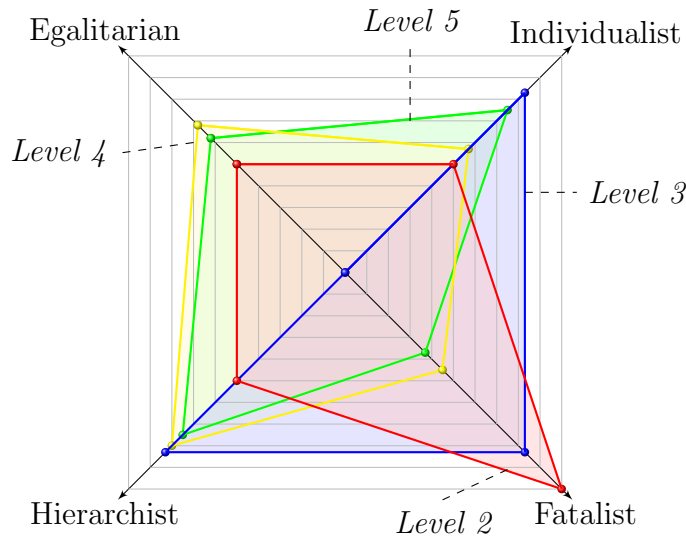
*'Assuming he can change the password straight after—that's not too bad.'* (scores equally as an *Individualist & Fatalist*,  
Level 4 / 'Policy Compliant')

This *Individualist/Fatalist* may not necessarily be considering policy, but nonetheless they are attempting to maintain some level of security. That *Individualists* also have a sizable representation in this division, albeit at the lower levels, further implies that the division's internal security culture is driven by the role of security in highly-skilled technology-related roles. *Individualists* and *Fatalists* who fit security to their role may be more naturally able to articulate the relevance of security to the goals they are trying to achieve. Where there is a lack of *Egalitarians*, the security function may compensate by talking to staff in the division, arranging security surgeries etc., to learn from the collective experiences of employees.

### 6.9.3 Business division

A more diverse security culture can be seen in the Business division. There is an approximately equal mix of *Egalitarians*, *Hierarchists* and *Individualists* present here. This alone suggests that a 'one-size-fits-all' approach to engaging with employees would not reach everyone in the division. The existence of diverse organisational (sub)cultures can conversely be a source of resilience during times of change (Schein 2010); this division may have valuable insights to offer security managers through open dialogue, rather than needing their guidance.

The division may have a good security posture that accounts for new and unexpected security dilemmas, if they can be addressed in a timely way. The free-text responses from this division were varied, showing security compromise as well as



**Figure 6.5: Kiviat diagram for the Business division** Maturity level 2 (red), 3 (blue), 4 (yellow), 5 (green) compared to the distribution of behaviour types.

policy enforcement. When faced with the prospect of having to share credentials in order to get work done, one respondent comments:

*‘Assuming the colleagues are from the same team and have the same clearance then they are equally trustworthy.’*  
(Fatalist, Level 2 / ‘Technically Controlled’)

Another respondent stands up for the policy, declaring that the actions offered to address the survey question’s dilemma are not sufficient. When faced with the prospect of transferring data over an insecure connection, s/he states:

*‘Would liked to have seen this option as a choice: [additional option] Report the [connection] problem and sit back until its fixed. Ignoring the fact that the work is crucial.’* (Fatalist, Level 4 / ‘Policy Compliant’)

*Individualist* responses to such a query are more balanced. When faced with insecure choices for transferring restricted data, one respondent shows a highly mature attitude to policy (as on the maturity scale) while at the same time risking actions that may be judged negatively by security managers:

*‘It depends on the level of security on the [bring your own device] laptop—if it’s password protected and has encryption that is more acceptable. Online services such as Dropbox should not generally be used for confidential information, particularly if not [approved for use].’*  
(Individualist, Level 4 / ‘Policy Compliant’)




*Egalitarians* are by default social creatures; they thrive in groups to solve problems. This division stands out as it contains the highest proportion of *Egalitarians* from all divisions considered. Their social leaning may well help to engage others. A typical comment from an *Egalitarian* person, when encountering an unlocked and unattended workstation, is:

‘[I would] send an email from the user of the unlocked machine to the team, offering to buy ice creams for everyone.’

(*Egalitarian*, Level 5 / ‘Active Approach’)

Although there is a lot of variation in behaviour types in this division, it could be useful to engage employees here in a number of different ways to capture all of their experiences (as demonstrated by the quotes above), certainly *before* any attempts to reinforce policy from the top down. The large percentage of attitude Level 2 and 3 *Fatalists* may however benefit from having a clear, workable process that employees can follow; those at higher levels can inform how that can be achieved in a way that does not hinder reaching business goals.

#### 6.9.4 Finance & Professional Services division



**Figure 6.6: Kiviat diagram for the Finance & Prof. Services division** Maturity level 2 (red), 3 (blue), 4 (yellow), 5 (green) compared to the distribution of behaviour types.

Considering the Finance & Professional Services division, the data shown in [Figure 6.6](#) illustrates that the predominant behaviour type is *Individualist*: approximately 40% (26 employees) of attitude Levels 3, 4 and 5 display this type. In this division, there is a discrete switch from *Fatalist* to *Individualist* from attitude Level

3 (‘Ad-hoc Knowledge and Application’) upwards, this is statistically significant, with Pearson’s  $r = -0.91$ ,  $n = 45$ ,  $p < 0.01$ . Both approaches are *individualised*, but differ in that the former experience prescribed ‘Inequality’ and the latter act toward ‘Equality’ (Adams 2003). It might be that those employees whose jobs are constrained by IT are relying on the organisation to ensure security for all, whereas those with an understanding of policy feel that they own how it is applied.

A top-down, prescriptive approach to engaging employees may fail to achieve results here, where *Individualists* at higher attitude levels may have an increasingly clearer understanding of how security policy fits with business goals. Indeed, communicating more security-related information to those at lower levels of security attitude may set employees up for *cognitive overload* and *embarrassment* (Beautelement, Sasse and Wonham 2008).

### 6.9.5 Discussion and limitations

We found that separate divisions of the organisation exhibited different distributions of both behaviour type and attitude level, immediately indicating that a ‘one-size-fits-all’ approach to messaging around security would achieve mixed results. By combining behaviour types and attitude levels (as defined in the maturity model summarised in Section 6.4.5.2), we were able to reason about differences in the security experiences of distinct groups as well as the dynamic between employees and the security provisions available to them.

Similarly, employees’ free-text responses indicate that they can offer insights about the security challenges they face and how to craft solutions for their working environment: ‘champions’ who can improve security can then be found in different places. For example, staff in the Operations division tended toward *Fatalist* behaviour, where security would be something taken care of ‘by somebody else’. Given that this included *Fatalists* with high attitude levels, any enhancements to security here may look at ensuring timely support from IT security representatives to leverage the high *risk understanding* seen in the division to spot risks as they arise and mobilise support. Conversely, in the Finance & Professional Services division more *Hierarchists* would be desired to improve the understanding of policy. Security improvement must leverage the *existing* security culture; the recruitment of security champions needs to be tailored to specific groups and divisions.

Regarding the potential for those exhibiting different behaviour types to inform policy improvements, free-text responses indicated that all four behaviour types have the potential to provide additional comments for use by security managers. For instance, *Egalitarians* were under-represented in the number of survey re-

sponses (Table 6.6), but like the other behaviour types provided increasing amounts of feedback as security maturity levels increased, with comparable response rates (Table 6.5). This indicates that if the organisation simply acted to move employees to a different behaviour type (e.g., *Hierarchist*) that much of the fuel for developing policy that works for the business would be lost (for instance, information about how to improve policy so that *Hierarchists* have a more effective policy to adhere to). An interesting finding is that the framing of an employee's feedback has the potential in itself to indicate their security behaviour (and I will be making use of this in Section 6.10). That is to say, that given the right prompts, a few sentences from an employee provides information to inform policy but can also allude to the security culture that the employee is experiencing.

Acknowledging that different approaches to security-related behaviour can exist in the same organisation, or even the same organisation division, has the potential to be an effective path toward lasting behaviour change. In considering the development of positive security culture which can persist over time, Pfleeger, Sasse and Furnham (2014) refer to the moral foundations theory of Haidt (2012), suggesting that knowing the moral profiles for individuals and groups can be a way to understand where support is needed, providing awareness and training which leverage each person's moral values using a range of options (e.g., warning messages, reminders, reading resources). Where here we use behavioural inclinations to characterise the relationship between an employee and the security infrastructure around them, group-specific packages of interventions which can be applied in this way can be used to reach employees more effectively, but also moderate the amount of messaging that employees receive.

The sample of the workforce studied here has few *Egalitarians* (see Table 6.5). It implies that there are few opportunities or little interest in socialising security. Developing a two-way *security dialogue* between security managers and employees can promote the alignment of security policy and process with the working environment (Ashenden and Lawrence 2016). Identifying individual and team trust dynamics could identify the most effective ways to engage in that dialogue, as security challenges may be resolved at a local team level as much as by following prescribed policy directives (Kirlappos and Sasse 2015). The organisation studied here is a large organisation with thousands of staff; having a distinct security function in a large, complex work environment may indirectly result in a *Fatalist* approach of security being 'not my job' (as hinted at for the Operations division), so there would be value in examining diverse organisations to determine the effect of organisation size and the viability of engagement activities.

A limitation of the analysis conducted here is that the data collected is a *snap-*

*shot* of the security experience at one point in time, and does not account for how security processes have developed over time up to that point, or any disruptive events which may have occurred within the organisation (such as a merger, change in applicable regulations, large-scale IT renewal, etc.). However, if events were disruptive enough to result in a change in the experience of security (e.g., adoption of another company's policies), employees may be inclined to comment on the change from their perspective when engaged in an activity such as the survey described here. Similarly, collection of survey responses over time can build a picture of security culture development, not least to understand the impact of security awareness initiatives.

### 6.9.6 Recommendations for practitioners and researchers

For the policy owner, our work reaffirms the need for security policy to be relevant to work activities and for employees to understand the risks that relate directly to those activities. Furthermore, policy should not be seen as an immovable object; when circumstances in the organisation change, the security policy should be revisited. The best sources of information for the policy owner are the employees: regular, direct, two-way interaction with individuals from all departments will enable policy to remain aligned with reality. The correct policy might not be known, but colleagues can be engaged to identify where it falls short and identify the underlying causes of non-compliance. A policy is only workable if employees are involved in shaping it, and employees can accept that they have to 'do the right thing'.

Our recommendation to security awareness professionals and researchers is to target awareness content to specific security behaviour-types and attitude-levels. Any one mode of engagement will resonate best with a different portion of the employee population. Encouraging meaningful responses from employees may require a combination of different approaches, such as surveys, workshops, and individual interactions. The tailoring of advice does not need to be guesswork. Engaging with groups of employees in the right way can immediately inform a picture of the security culture across the organisation. Ultimately, an understanding of relevant risks and work-related motivations could make *targeted* interventions much more successful. If security doesn't work for people, it doesn't work.

### 6.9.7 Conclusion

In this section I have analysed the survey responses for Company B by combining behaviour types and maturity levels to characterise the quality of security policy for groups of employees. I focus on the organisation as a whole, and four business

divisions in detail, including Operations and Business divisions. The analysis of 189 optional free-text comments linked to the survey further suggests that those who follow policy can contribute to effective security, but so can those who question policy, socialise solutions, or would otherwise expect security to be part of their productive work if it was important. These various security cultures all have the potential to help improve security for the whole organisation, where here we have identified a range of security heroes for security managers to engage.

The ideal type of security hero engagement for the organisation studied here would be a composite of different approaches: it cannot be identified by one specific set of traits, but rather is entirely dependent on the social context it manifests in. Individuals may act alone or together, with policy in mind or in isolation from it. The methodology presented here, combining security attitude and security behaviour, is a useful tool for investigating the interplay between policy and action. Attempting to narrowly define and promote the characteristics of a security champion is counterproductive, as it simplifies the challenge of involving employees in the process of improving organisational security.

## 6.10 Validation of Surveys

So far in this chapter, we have relied on being able to cluster the participants into predefined groups based on their survey responses. In this final section of this chapter, I aim to establish the usefulness of framing survey questions around active security controls and problems experienced by employees, by assessing the validity of the clustering. We introduce measures for the appropriateness of the survey scenarios for each organisation and the quality of candidate answer options. We use these scores to articulate the methodological improvements between the two surveys.

I develop a methodology to verify the clustering of participants, where 516 (A) and 195 (B) free-text responses are coded by two annotators. Inter-annotator metrics are adopted to identify agreement. Further, we analyse 5196 (Company A) and 1824 (Company B) appropriateness and severity scores to measure the appropriateness and quality of the questions.

We explore the capacity to utilise additional types of questions to reflect on the survey design without further effort by the researchers. If the participants are given an opportunity to indicate the applicability of the scenarios to their environment, we can tailor the results not just to specific user groups, but also reflect on how a survey engages with diverse groups and their security needs.

Participants rank questions in B as more appropriate than in A, although the variations in the severity of the answer options available to participants is higher in B than in A. We find that the scenarios presented in B are more recognisable to the participants, suggesting that the survey design has indeed improved. The annotators mostly agree strongly on their codings with Krippendorff's  $\alpha > 0.7$ . A number of clusterings should be questioned, although  $\alpha$  improves for reliable questions by 0.15 from A to B.


To be able to draw valid conclusions from survey responses, the train of analysis needs to be verifiable. Our approach allows us to further validate the clustering of responses by utilising free-text responses. Further, we establish the relevance and appropriateness of the scenarios for individual organisations. While much prior research draws on survey instruments from research before it, this is then often applied in a different context; in these cases adding metrics of appropriateness and severity to the survey design can ensure that results relate to the security experiences of employees.

### 6.10.1 Employee types

In each of the two surveys we attempt to position participating employees on two dimensions. In A these are Attitude and Behaviour types, whereas in B these have evolved to Maturity levels and Behaviour types. For the definitions of the types please see [Table 6.7](#) or [Sections 6.4.4](#) and [6.4.5](#). Foremost, these two dimensions can be examined individually and in combination, across age groups, business divisions, and physical locations, to target interventions which reduce friction between security and productivity in the workplace.

The attitude types in A focus on individuals' interaction with security apparatus. In B, these have evolved to a scale of Maturity Levels, which are ranked levels of individuals' interaction with the organisation's policy (such that interventions would act to improve employees' working interactions with centralised security policy and security provisions). In both A and B, the participants were also asked to assign an appropriateness score for each answer option on a 5-point Likert scale, ranging from *not acceptable at all* to *very acceptable*.

The behaviour types in A are a measure of the individuals' likelihood to trade-off security for productivity. This evolved to the more abstract concepts in B where the answers are now mapped to four distinct behaviour types as defined by Adams (2003), to better represent the role of teams and organisational culture in individual security behaviours. Additionally, participants were asked to assign a severity score to each answer option of the behaviour type questions, as well as give a general

<hr/> 1 Discount suspicions, cause no bother, passive, 2 Report suspicions but take no direct action, 3 Take direct action through official channels, 4 Take direct personal action against the threat. <hr/> (a) Attitude Types for A, as described in <a href="#">Section 6.4.4.1</a>
<hr/> 1 Prepared to perform insecure acts to maximise productivity, 2 Show a minor priority for work over security when the two conflict, 3 Passive, expects others to take the initiative to ensure security, 4 Tries to remain secure wherever possible. <hr/> (b) Behaviour Types for A, as described in <a href="#">Section 6.4.4.2</a>

<hr/> 1 Is not engaged with security in any capacity, 2 Follows security policy only when forced to do so by external controls, 3 Understands that a policy exists and follows it by rote, 4 Has internalised the intent of the policy and adopts good security practises even when not specifically required to, 5 Champions security to others and challenges breaches in their environment. <hr/> (c) Maturity Levels for B, as described in <a href="#">Section 6.4.5.2</a>
<hr/> <p><b>Individualists</b> rely on themselves for solutions to problems,  <b>Egalitarians</b> rely on social or group solutions to problems,  <b>Hierarchists</b> rely on existing systems or technologies for solutions to problems,  <b>Fatalists</b> take a ‘naive’ approach to solving problems, feeling that their actions are not significant in creating outcomes.</p> <hr/> (d) Behaviour Types for B, as described in <a href="#">Section 6.4.5.1</a>

**Table 6.7:** The dimensions by which survey responses are measured in Company A & B

indicator as to how acceptable to the business it would be for the participant not to finish the task described in each scenario.

### 6.10.2 Survey design

In each organisation, surveys were crafted for participants based upon their department to improve relevance, as described earlier in this chapter. Question design attempted to offer the participants a number of options which would all be regarded as equally appropriate, based upon the themes identified from the preceding interviews with employees. The participants were asked to rank the four options in order of their preferences. Additionally, participants were allowed to offer additional comments, which included the following example:

*‘Shamal needs to find out who manages the common drive now, and whether the company authorises use of Dropbox and personal USB sticks, before using any of those options.’*

As part of the work described here, two annotators coded the volunteered comments for the types (without reference to the alignment of responses to types already defined for each question). For example, the quotation above could be coded as a Hierarchist’s point of view, as the individual falls back to existing structures for solutions to the problem.

*‘This scenario could easily be avoided by providing sufficient space on the common drives.’*

Conversely, this statement has been coded as a Fatalist. The employee is frustrated that the natural solution to the problem is outside their reach.

### 6.10.3 Methodology

The methodology laid out in this section establishes three metrics to measure the quality of the survey design and its external validity retrospectively. The quality of the survey includes how engaged participants are in considering a scenario, and how relevant a scenario and its options are to their own experiences. If an organisation is committed to measuring how well its security provisions support the effective completion of business tasks towards identifying and removing frictions, decision-makers would have a natural interest in having a realistic picture of the current experiences of employees.

#### 6.10.3.1 Appropriateness and applicability

For each of the answer options to Attitude and Maturity questions (see [Tables 6.7\(a\)](#) and [6.7\(c\)](#)) participants were asked to specify the acceptability of that answer on a 5-point Likert scale ranging from ‘Not acceptable at all’ to ‘Very acceptable’. There



are of course biases present here, namely that given a participant's type they may see some options as more acceptable than others. Indeed there is a statistically significant (at  $p < 0.001$ ) correlation in our survey response data between the ranking of options and their associated behaviour types with Kendall's  $\tau$  of 0.62 (as described in [Table 6.4\(a\)](#)). Yet the ideal scenario design would leave the participants with four objectively equally acceptable options, and allow the participant to freely rank the option. Hence a high appropriateness score is desired.

For each of the answer options to Behaviour type questions (see [Tables 6.7\(b\)](#) and [6.7\(d\)](#)) participants were asked to specify the severity of each option as well as the *acceptability of failing to complete the task* for each scenario on a 5-point Likert scale. Again, the severity scores are statistically significantly (at  $p < 0.001$ ) correlated with the ranking of the answers (Kendall's  $\tau = -0.20$ , [Table 6.4\(b\)](#)), with less severe answers being ranked as more preferable. The severity scores of the different answers should be ranked equally by the participants (as questions are designed with no one 'right' answer), resulting in a low standard deviation throughout the questions. The ideal mean of the standard deviation of severity of options is 0, which would imply that all options given to the participants are perceived as equally severe.

The *acceptability of failing to complete the task* metric would ideally be identically distributed for all questions in order to allow for inter-question comparison. This is a metric that is difficult to establish through prior analysis. If participants think that for a scenario it is more acceptable for it not to be completed given the given consequences as in the scenario and its options, the participants do not fully commit to their choices of behaviour types, as in no scenario there is an option to do nothing (and in turn avoid side effects from the chosen solution).

### 6.10.3.2 Validation of ranked types by free-text responses

In the survey design for both A and B, participants are asked to rank four answer options according to their preferences. Participants are also invited to provide additional comments on the questions. We find that there are two common types of responses: those that further confirm a respondent's answer, or elicit suggestions and solutions that are not included in the question and the associated options. We code these according to the applicable mapping (Attitude or Behaviour) in each organisation as listed in [Table 6.7](#), e.g., for each free-text response the annotators have to choose from one of four options. While this is opportunistic (not all participants provided additional comments), we can validate the mappings by calculating inter-annotator agreement metrics as described in the following sections.

### 6.10.3.3 Inter-annotator agreement

Coder B	Coder A			
	T1	T2	T3	T4
T1	4	3	0	0
T2	0	36	2	15
T3	0	2	59	1
T4	0	0	1	32

**Table 6.8: Inter-coder confusion matrix** Confusion matrix for Question 1 in A between the coders' assignment of types to the free-text responses

The calculation of the inter-annotator agreement between the two coders is straightforward. We first calculate a confusion matrix for each question (an example is shown in Table 6.8), and then calculate Krippendorff's chance corrected inter-annotator agreement metric  $\alpha$ . Krippendorff's  $\alpha$  ranges from  $-1$  to  $+1$ , where  $0$  corresponds to chance agreement and  $+1$  to perfect agreement. As the attitude types and maturity levels in Tables 6.7(a) and 6.7(c) are on a ranked scale, we weight the disagreement linearly. For the other two types described in Tables 6.7(b) and 6.7(d) the agreement is binary.

### 6.10.3.4 Validating the mapping

Rank	Coding type			
	T1	T2	T3	T4
1	2	6	84	3
2	5	26	22	10
3	2	19	12	34
4	2	43	6	34

**Table 6.9: Coder-mapping confusion matrix** Confusion matrix for Question 1 in A between participants assigned ranks to the potential answers and the types assigned to the participants by the coders based on the coding of the free-text responses

We validate the mapping of the survey answer options by treating the participants as another annotator and calculating the inter-annotator metric  $\alpha$ . However, the participants rank their options, but the coders annotate separate (but not independent) text statements. For example, a participant may provide a ranking of Type 3 > Type 2 > Type 4 > Type 1 for a specific question, and from the coders we may see Coder X: Type 3, Coder Y: Type 2.

In this case the standard agreement table approach (Gwet 2014) for  $> 2$  annotators cannot be used. Yet Krippendorff's  $\alpha$  naturally extends to non-square weight matrices. In our case, this leads to a confusion matrix such as in Table 6.9. Here we tabulate the frequency that a coder has annotated a statement with a specific type with the rank that the participant gave that type. Perfect validation would therefore imply that all types chosen by the coders have rank 4; i.e. the only non-zero entries are in the bottom row of the confusion matrix. Given this matrix we can execute the calculation of Krippendorff's  $\alpha$  with a weights matrix that treats numbers in the bottom row as perfect agreement, and linearly increases disagreement for lower ranked options.

### 6.10.3.5 Estimating confidence in $\alpha$

In order to calculate the confidence in the calculated value of  $\alpha$  we rely on  $\alpha$ 's standard deviation. As an analytic expression is not available, we bootstrap the calculation of  $\alpha$ . In the following sections the confidence intervals are calculated using 1000-fold bootstrapping.

### 6.10.4 Results

In this section we present the application of the metrics defined in Section 6.10.3 to both Company A and B. There are four tables to consider in this section; Tables 6.10 and 6.11 for the analysis of the secondary coding of the free-text responses, Table 6.12 for the analysis of appropriateness scores on attitude/maturity questions, and Table 6.13 for the analysis of the severity metrics on behaviour type questions.

#### 6.10.4.1 Analysis of clustering

Table 6.9 is an example confusion matrix calculated based on the methodology presented in Section 6.10.3.4. The column headers list the four possible types assigned to the free-text responses by the coders. If a free-text response by the coders was judged to be type 1, but the participant ranked the answer corresponding to type 1 as rank 2, this would increment the number in row 2, column 1. Perfect agreement would be represented by the type assigned through coding of the free-text responses always being ranked highest (rank 4) by the participants. This would be a confusion matrix of non-zero entries in the bottom row only.

The strong disagreement between the coders and the assigned rank in question 1 can be identified by the strong mismatch in type 3: of the 124 statements assigned to type 3 by the coders, 84 were ranked least likely (rank 1) by the participants.

This implies that the answer option assigned to type 3 (*‘Request that those with access share their (main log-in) account details and passwords with those without to allow them access to the information’*) does not match behaviour type 3 (as defined in Section 6.10.1) (*‘Passive, expects others to take the initiative to ensure security’*).

Interestingly, this disagreement is not reflected in the coding of the free-text responses themselves. Table 6.8 shows the confusion matrix for Question 1 for the two coders. There is virtually no disagreement for types 1, 2 and 3; but some disagreement for type 4, where 15 statements assigned to type 4 by coder A were considered to be type 2 by coder B. The internal validity for the coding of free-text responses for Question 1a can be accepted based on Krippendorff’s  $\alpha$  of  $0.77 \pm 0.00$  as shown in Table 6.10, but we are unable to validate the mapping of answer options to types.

Question	#	Mapping $\alpha$	Coder’s $\alpha$
Q4	40	$0.21 \pm 0.02$	$0.29 \pm 0.02$
Q5	34	$0.35 \pm 0.02$	$0.02 \pm 0.03$
Q6	2	$-0.33 \pm 0.76$	$0.00 \pm 0.67$
Q8	29	$0.30 \pm 0.04$	$0.94 \pm 0.02$
Q10	37	$0.23 \pm 0.03$	$0.73 \pm 0.02$
Q1	155	$-0.03 \pm 0.01$	$0.77 \pm 0.00$
Q2	137	$0.43 \pm 0.01$	$0.91 \pm 0.00$
Q3	12	$0.33 \pm 0.08$	$0.38 \pm 0.10$
Q7	25	$0.24 \pm 0.03$	$0.13 \pm 0.04$
Q9	45	$0.13 \pm 0.02$	$0.76 \pm 0.02$

**Table 6.10:** Krippendorff’s  $\alpha$  measures for compA with 95% confidence intervals.

Question	#	Mapping $\alpha$	Coder’s $\alpha$
QID	33	$0.27 \pm 0.02$	$0.85 \pm 0.02$
QCDP	53	$0.31 \pm 0.01$	$0.38 \pm 0.02$
QT	22	$0.24 \pm 0.04$	$0.34 \pm 0.05$
QSD	27	$0.53 \pm 0.02$	$0.47 \pm 0.04$
QRM	12	$0.27 \pm 0.07$	$0.42 \pm 0.07$
QVPN	23	$-0.09 \pm 0.03$	$0.37 \pm 0.04$
QFS	18	$0.19 \pm 0.05$	$0.46 \pm 0.05$
QCC	7	$0.38 \pm 0.13$	$0.75 \pm 0.21$

**Table 6.11:** Krippendorff’s  $\alpha$  measures for compB with 95% confidence intervals.

Tables 6.10 and 6.11 list the number of free-text responses coded and Krippendorff’s  $\alpha$  for both the validation of the mapping as well as the coders agreement.

#### 6.10.4.2 Suitable values for $\alpha$

Before discussing this data further we must delineate the boundaries for which we consider Krippendorff's  $\alpha$  to be reliable. From a statistical perspective we can conduct a t-test where the null hypothesis is  $\alpha = 0$ , i.e. the data is equivalent to chance. This t-test is represented in our tables through the use of 95% confidence intervals. Indeed all rows that are statistically significant at the 95% confidence interval are also significant at the 99% confidence interval. However the literature (Gwet 2014) is clear that primary data is only sufficiently reliable for further analysis at  $\alpha > 0.667$ .

It is clear that most of the coder's agreement values in A satisfy this criteria. There are a number of exceptions: *Q5*, *Q3* and *Q7*. The inter-coder agreement is not as strong in B, where only *QID* satisfies this criteria. When focusing on the validation of the mapping/clustering however, none of the scenarios satisfies this stringent criteria.

Considering the difficulty the coders have to establish agreement on the free-text responses in B, the low mapping  $\alpha$  values are not surprising: the coding is a difficult task (given the brevity of comments and potential lack of contextual information). Yet rather than discarding the results at this stage, it may be more important to identify the scenarios which are indistinguishable from random data: scenarios *Q1* in A and *QVPN* in B. Apart from these two scenarios, our data allows the focus of further investigations and policy decisions to be guided by data with known uncertainty.

#### 6.10.4.3 Appropriateness

Table 6.12 shows the appropriateness scores the participants have given the answer options for specific questions. The scores vary from 0 (not appropriate) to 1 (very appropriate). The mean appropriateness score is more varied in A than in B, although it is close to 0.5 for all questions, indicating that the average answer option is balanced. This is desirable as it offers participants the option to swing to both extremes as necessary. The appropriateness score given by participants to their highest ranked choice is very high, confirming the participant's stance that they view their preferred choice as most appropriate.

#### 6.10.4.4 Severity

Table 6.13 compares the distribution of severity scores and *acceptability of failing the task* scores across the different scenarios and organisations. There are a number of

Question	#	Mean		1st choice	
		mean	std	mean	std
Company A					
Q4	374	0.626	0.120	0.923	0.195
Q5	820	0.570	0.110	0.925	0.161
Q6	137	0.427	0.138	0.821	0.321
Q8	364	0.529	0.085	0.983	0.084
Q10	903	0.483	0.082	0.917	0.185
Company B					
QID	152	0.488	0.122	0.778	0.316
QCDP	456	0.508	0.108	0.893	0.220
QT	164	0.499	0.095	0.873	0.252
QSD	292	0.546	0.118	0.939	0.181

**Table 6.12: Appropriateness scores for each attitude question.** The higher the score, the more appropriate. As each answer option is assigned an appropriateness score, the mean represents the mean appropriateness score of all answer options irrespective of that answer’s ranking. The *1st choice* only considers the appropriateness assigned by the participants to their top choice.

Question	#	Failing		Std of Severity	
		mean	std	mean	std
Company A					
Q1	903	0.281	0.307	0.270	0.128
Q2	893	0.270	0.296	0.239	0.123
Q3	137	0.394	0.340	0.271	0.122
Q7	291	0.458	0.393	0.296	0.144
Q9	374	0.668	0.449	0.274	0.123
Company B					
QRM	152	0.196	0.312	0.377	0.101
QVPN	152	0.439	0.370	0.323	0.120
QFS	164	0.430	0.318	0.297	0.114
QCC	292	0.240	0.410	0.182	0.163

**Table 6.13: Acceptability of failing to complete the task and severity scores analysis.** The higher the score, the more acceptable is failing to complete the task outlined in the scenario. The severity scores are an analysis between the answer options for behaviour type scenarios.

variations: scenarios in A are considered less acceptable to be left undone, however the standard variations of the severity scores across the different scenario options are higher. According to [Table 6.13](#) the scenarios in B are therefore believed by

the participants to be more applicable to their environment (particularly *QRM* and *QCC*), but the answer options are more balanced in severity in A, implying that options represent potential solutions that may be seen in everyday work in A compared to the more contrived answer options in B.

### 6.10.5 Conclusions

In this section we have described a methodology for post-hoc assessment of the quality of situated security behaviour survey designs. We utilise free-text responses and reflective metrics to measure the surveys' external validity. We have demonstrated this approach on two surveys in two large organisations, drawing on 711 free-text responses and over 7000 reflective scores in the process. This has allowed us to quantify the evolution of our scenario-based surveys through clearly-defined and repeatable metrics, and partially validate the mapping from survey responses to constructs. This knowledge will allow security managers to tailor future improvements to their organisation's security policy and behavioural interventions more accurately to the local working environment, relative to the demonstrable strengths and weaknesses of the survey design.

## 6.11 Discussion

This research supports a process of continuous improvement to organisational security, by providing measures for (i) typical workaround to regular frictions with security in the workplace (by analysing the perceived suitability of solutions derived from interviews), and (ii) how the interactions employees have with security apparatus can be designed to minimise the demand on their 'compliance budget' (Beautement, Sasse and Wonham 2008). Employee's willingness to expend effort for the security of not only themselves but those around them can be explored by articulating embodied security cultures which may arise in any number of situations in the workplace where security controls can be applied. Both the survey results and the free-text responses can inform targeted interventions as part of incremental improvement, an approach advocated by Renaud and Goucher (2014). Unfortunately in striving for internal validity for security behaviour constructs it is easy to overlook the need to establish the applicability of the results to the real world, that is, to measure the quality of *engagement* with employees (where tensions can arise with local demands on effort and capacity). The related works discussed in [Section 2.1](#) demonstrate this well.

Security managers ought to identify the non-divisible security behaviours in their own organisations, and equally deploy information security surveys that shine light on previously unseen *workaround* or *compromise* behaviours by engaging employees. To do this, available options (and ideally, additional feedback from users) must point to clear responses to security-related challenges that employees see as acceptable given the pressures they perceive in a particular situation. Where respondents imply confusion about what is being asked of them in a scenario-based survey question, or indeed, see two or more behaviours as one and the same, this implies that more can be done to clearly separate candidate behaviours. In turn, this can be achieved if security managers act to grow their understanding of how security manifests for employees who have other competing demands for their attention (see also Ashenden and Lawrence (2013), Herley (2014) and Parkin et al. (2016)).

Our proposed survey methodology and validity measures address these challenges. This is achieved both internally (by way of inter-annotator agreement), and externally (by way of appropriateness and severity scores). We are able to highlight strengths and shortcomings in the survey design which not only inform the design of realistic scenarios by researchers, but also inform the investment in security by policy managers when designing interventions. Organisational environments are complex, and researchers cannot assume that they have a full understanding of security behaviours prior to deploying a survey. This research helps to identify these known unknowns. Security practitioners considering potential investments may do well to understand the quality of the data they base their decisions upon (Hubbard 2014).

### 6.12 Conclusion

In this chapter I have discussed the Productive Security methodology and the analysis of two scenario-based survey procedures at two large companies with 1299 and 608 responses (Sections 6.4 to 6.7). These responses define the security culture in the organisation, and I re-frame the idea of a security champion for Company B (Section 6.8). Finally, myself and Simon Parkin both code 516 (A) and 195 (B) free-text responses in order to validate the scenario-based surveys.

We find that tailoring our diagnostic tools to the operating context and working practices of the organisation provides meaningful results. Security awareness material can similarly be crafted to resonate with the experiences of employees in weaving security into their productive tasks. Tsohou, Karyda and Kokolakis (2015) discuss ways of interpreting cognitive and cultural biases, such as those described in the behaviour-type scenarios, to produce effective security awareness material. Aware-



ness should be a two-way street: security specialists should use the understanding of what drives individuals' behaviour to engage with those individuals and be receptive and find collaborative solutions to conflicts between security and business processes.

This chapter has delivered a grounded understanding into the security culture of an organisation, based purely on self-reported responses. This mixed-method approach highlights the versatility of free-text responses, as anecdotal, yet frequent comments have been re-purposed through coding. The following chapter contrasts the methodology here by primarily relying on log data to study behaviour.



# Password Project

*‘I wasted time, and now doth time waste me.’*

(W. Shakespeare 1595)

While in the previous chapter I focussed on studying security through self-reported surveys, for the study presented in this chapter I had access to the log data of UCL’s password change system. In particular, UCL had decided to adopt a new password policy, and I was able to opportunistically study the impact of the new password policy with 100,000 staff and students. This research was carried out in collaboration with Simon Parkin and Angela Sasse, and this chapter is based on a paper presented at USENIX Security (Becker, Parkin and Sasse 2018).

The goal of the IT staff who conceived the policy was to encourage stronger passwords by varying password lifetime according to password strength. Strength was measured through Shannon entropy (acknowledged to be a poor measure of password strength by the academic community, but still widely used in practice). When users change their password, a password meter informs them of the lifetime of their new password, which may vary from 100 days (50 bits of entropy) to 350 days (120 bits of entropy).

I analyse data of nearly 200,000 password changes and 115,000 resets of passwords that were forgotten/expired over a period of 14 months. The new policy took over 100 days to gain traction, but after that, average entropy rose steadily. After another 12 months, the average password lifetime increased from 146 days (63 bits) to 170 days (70 bits).

I also found that passwords with more than 300 days of lifetime are 4 times as likely to be reset as passwords of 100 days of lifetime. Users who reset their

password more than once per year (27% of users) choose passwords with over 10 days fewer lifetime, and while they also respond to the policy, maintain this deficit.

It appears that linking password lifetime to strength at the point of password creation is a viable strategy for encouraging users to choose stronger passwords (at least when measured by Shannon entropy).

## 7.1 Introduction

A new password management system was deployed at UCL in Autumn 2016, replacing the system used by all university staff and students. Users are being moved from fixed-length passwords (8 characters, complexity 3, expiring after 150 days) to a variable-length password scheme requiring complexity 3 (3 out of 4 character groups, see [Figure 7.3](#)) with an expiry dependent on the length and strength of the password (see [Section 7.3.1](#)). The introduction of a ‘Longer Password Longer Life’ scheme is novel, especially at a large organisation such as UCL.

The new mechanism was created by an experienced systems administrator who was motivated by the high cost of helpdesks for resetting passwords and user complaints, but unaware of the research literature or the NCSC Password Guidance (National Cyber Security Centre (NCSC) [2016a](#)). The intent was to change to a ‘3-words-strung-together’ scheme, but the UCL committee responsible for security policies mandated that existing complexity 3 requirements must remain.

The new policy allows users to select any password of character length 8 or more with an estimated information entropy (Shannon entropy, a poor measure of cracking resistance, but still widely deployed) of at least 50 bits (see [Section 7.3.4](#) for the policy specifics). The new system retains the expectation that users will harden their accounts with strong passwords, but in a twist provides a reward of longer password lifetime for selecting stronger passwords. A password with an estimated entropy of 50 bits has a lifetime of 100 days, and every additional bit of entropy increases the lifetime by approximately 3 days, up to 350 days for 120 bits of entropy.

We then use the term *password strength* here as the number of days a password lives for before being expired, as this is a measure of account strength that is visible to both the users and managers of the system.

The research questions examined in this chapter are:

- RQ1** What effect does the password policy of variable expiration have on a user’s choice of password?
- RQ2** Are there identifiable groups of users with analytically different responses to the new password rules and introduction of the new policy?

**RQ3** What can be discerned about the impact of a policy intervention at a large institution from system logs?

We believe that this research constitutes the largest analysis of password data from a single institution with over 100,000 enrolled users in the system, who change their passwords nearly 200,000 times and reset (forgotten or expired) their passwords 115,000 times over a period of 14 months. Our approach is novel as we analyse routine change and intentional reset events together, to understand individual users' journeys through adoption and continued use of the new system. This approach leverages the working relationship with the system managers, who allowed continuing access to the anonymised log data and kept us informed on events outside of the system which could impact use and hence the logs themselves (such as university-wide events).

We begin the remainder of the chapter with an overview of the literature related to this project in [Section 7.2](#). After an introduction to our methodology in [Section 7.3](#), we describe and compare the general statistics of our dataset to prior studies on large password analysis ([Section 7.4](#)). This is followed by an analysis of the password change data in particular, answering our research questions in [Section 7.4.4](#). We draw on 93 interviews with staff and students for anecdotal user feedback in [Section 7.4.7](#). We then discuss the impact of the results in [Section 7.5](#) and close with conclusions and recommendations in [Section 7.6](#).

## 7.2 Related literature

The expiration of passwords for machine accounts has had a long history. Tracing back to 1979, expiration was a tool to stop users sharing accounts on the first university computers (Morris and Thompson 1979). This was not a need borne of security: it was a management mandate to allow for proper accounting of computation time. However the notion has been appropriated to serve security, spread by various international government guidelines that have since prescribed the expiration of passwords (Brand and Makey 1985; Burr, Dodson and Polk 2004). Various justifications for password expiration have been found: the longer a password is 'alive', the higher the chance of compromise and the need to reset passwords (due to sustained attacks or inevitable leakage), or, that expiration limits the portability of a compromised password, as old passwords may be replicated on other services for convenience (Bonneau et al. 2012; Cheswick 2013; Herley and Oorschot 2012).

These myths have been thoroughly debunked. The security benefits of password expiration are marginal at best (Chiasson and Oorschot 2015; Zhang, Monroe and Reiter 2010). Users regularly choose new passwords that are very similar to a previ-

ous password (through for instance incremental changes to a number in a sequence of passwords) (Zeuschwitz, Luca and Hussmann 2013; Zhang, Monroe and Reiter 2010). Further, passwords of sufficient strength can be combined with background protections to be strong enough in most scenarios: a password which can resist  $10^6$  guesses is all but uncrackable in an online attack scenario (Florêncio, Herley and Van Oorschot 2016), if combined with sensible throttling (Florêncio, Herley and Van Oorschot 2016; Weir et al. 2010). To defend against the offline attacks a password is required to withstand  $10^{14}$  guesses.

This body of research has informed practical advice, and a change of guidelines. Both the National Institute of Standards and Technology (NIST, Grassi, Garcia and Fenton (2017)) and the National Cyber Security Centre (NCSC) (2016a) now prescribe that passwords should not expire unless there is evidence of compromise.

A holistic view of password policy management is required in practice. For example, a user's choice to re-use passwords across separate accounts is rational when there are simply too many passwords to remember (Inglesant and Sasse 2010). Users may apply strategies to group accounts by perceived importance and assign a password to each group (Florêncio, Herley and Van Oorschot 2014b).

The literature on password research is divided into the following sections: we start with a discussion of password strength estimation, then focus on the user's role in password management and password studies.

### 7.2.1 Password strength estimation

Traditionally, password strength has been measured as the entropy of a password through a calculation involving a password's length and the different number of character classes it uses (Kelley et al. 2012) (Shannon entropy, which is also the estimation technique our institution uses, albeit with a few modifications as described in Section 7.3.4). These estimates are however not representative of the cracking effort, as passwords are not actually chosen randomly (de Carnavalet and Mannan 2014). This has led to the creation of strength meters inspired by password-cracking, which estimate the number of attempts required for a password to be guessed. The current state of the art is *zxcvbn* (Wheeler 2016), which algorithmically accurately estimates the strength of weak ( $< 10^4$  guesses) passwords with only 234kB of data. For stronger passwords the strength estimation error of *zxcvbn* increases, but it is still a better estimator of cracking resistance than information entropy. To accurately estimate the strength of stronger passwords, significantly more storage and processing power is required, however this is infeasible for real-time feedback (Ur et al. 2015).

### 7.2.2 The role of users in password security

A primary question that is easily ignored when conducting password research is the attacker's modus operandi, and consequent interactions with the state of security defenses. The main attack vectors of interest are online and offline attack. An online attacker performs attacks over a wire, while the offline attacker has access to the physical system. While an online attack can be rate limited, blacklisted, and actively monitored (Alsaleh, Mannan and Van Oorschot 2012; Pinkas and Sander 2002; Van Oorschot and Stubblebine 2006), none of these defenses are possible against an offline attack. This implies that the defensive requirements on the password are very different (Florêncio, Herley and Coskun 2007; Florêncio, Herley and Van Oorschot 2014a). For passwords to be resistant to offline attacks they realistically need to be able to withstand  $10^{14}$  guesses. In the context of an organisation it is not sufficient for the mean password strength to achieve this level: an attacker is often satisfied when compromising any one account with access to an asset of value, hence every password needs to withstand such an attack, which is infeasible (Herley 2009). When the entire system is under attack, the defense should be centered on the system too, rather than offloading it to all the users, for example through Ersatzpasswords (Almeshekah et al. 2015).

As researchers have identified the need to raise the minimum strength of passwords, a large number of studies have focused on helping and educating the user in choosing stronger passwords. Users have been subjected to immediate feedback and suggestions before submitting their password choices with varying degrees of success (Segreti et al. 2017; Shay et al. 2015). Research has attempted to improve users' ability to remember passwords, for example by allowing much longer composite passwords (Shay et al. 2014), memory aids (Yan et al. 2004), or training (Charoen, Raman and Olfman 2008). Perhaps unsurprisingly, positive attitudes towards security correlate with stronger passwords (Choong and Theofanos 2015). Such interventions are often measured over a relatively short timeframe; a wide-reaching intervention such as a password system overhaul may require time. We then leverage the opportunity to measure behaviour through password change events over time (where this would be impacted by users' capacity to remember passwords and use longer passwords in practice).

### 7.2.3 Studying passwords in the wild

A considerable amount of password research has been conducted in a lab setting. This allows for great internal validity through the ability to control the environment and measure specific properties of users choices and behaviours around passwords.

However, Fahl et al. found that only about half of passwords gathered in a lab study are comparable to users' real-world passwords (2013). This problem is not specific to password studies, a large number of lab-based studies in security suffer from a lack of ecological validity. However, studying security perceptions in the real-world comes with its own issues (Krol et al. 2016). Fortunately there are a number of password studies that are conducted in live environments.

The first scientific dissemination of password data was conducted on leaked password datasets (Dell'Amico, Michiardi and Roudier 2010; Weir et al. 2010). More recently Bonneau pushed the scientific principles of conducting password research by legitimately and rigorously analysing passwords of 70 million Yahoo! users (2012). The flurry of data breaches at large online services have fuelled research by providing extremely large datasets. Yet in all of these cases the user is often a customer of the organisation, with two consequences: service password policies tend to bow to the need for accessibility, as services that make access difficult don't have as many customers (Florêncio and Herley 2010). Users may not assign much value to these accounts, unless their personal data/money is stored there.

Apart from our research in [Chapter 7](#), the only other comparable study of password behaviour in a work environment with high value passwords to study is by Mazurek et al. (2013). Here the entire plaintext password database of over 25,000 accounts was available to the researchers (although considerable security precautions were taken to limit access to the plaintext passwords). The authors discover significant correlations between a number of demographic and behavioural factors and password strength, and we will be comparing our demographic findings to this research primarily.

Related to passwords, Parkin et al. studied a static password expiration policy of 100 days in a university, contrasting the analysis of helpdesk-related system events over a period of 30 months to findings from a small set of 20 interviews with system users (2015). Users appreciated the need for security and strong passwords, but their attempts to create strong passwords were frustrated by usability issues not directly apparent from system events (such as an inability to know in advance what the system would accept as a valid password).

Zhang, Monroe and Reiter studied 31,075 passwords belonging to 7,936 university accounts in order to analyse the dependency between consecutive passwords (2010). We contrast their main results to our data in [Section 7.4](#).



#### 7.2.4 Password policy

A comprehensive overview of the last 30 years of password policy research is given by Zhang-Kennedy, Chiasson and Oorschot (2016). Ever since ‘*Users Are Not the Enemy*’ there has been a sustained effort to design security policies for the user, taking into account their strengths and limitations (Adams and Sasse 1999). Strength aspects such as length and composition, as well as management aspects such as change-it-often, do-not-reuse, do-not-write-down and do-not-share-with-anyone have been either entirely revised or are at least strongly challenged (Brostoff and Sasse 2003; Burr, Dodson and Polk 2004; Florêncio, Herley and Van Oorschot 2016; Grassi, Garcia and Fenton 2017; National Cyber Security Centre (NCSC) 2016a).

User capability, user inclusion in their own and others’ security, and a holistic approach to defensive security then together serve as indicators for identifying a *sustainable, workable*, and ultimately *secure* password system.

### 7.3 Methodology

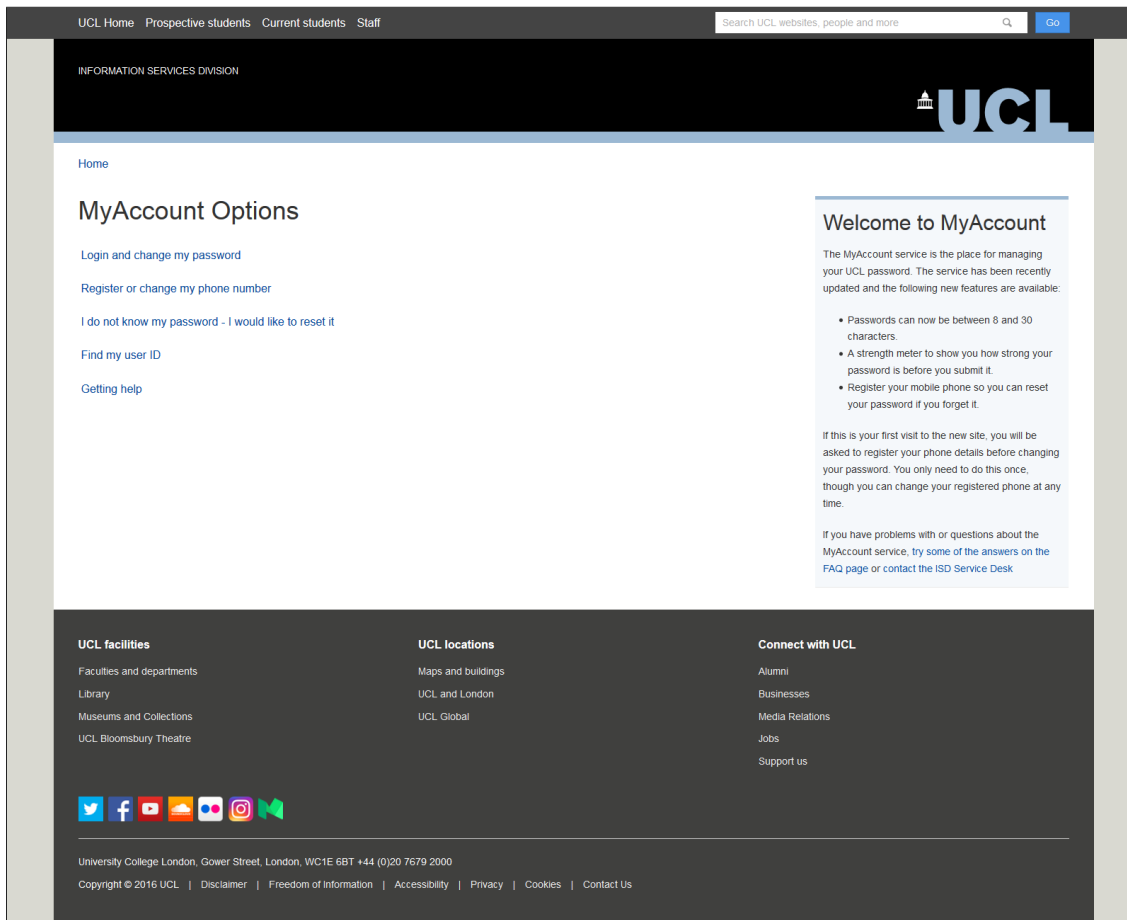
Here we describe the methodology for analysing the logs of the password change system at UCL. We were not involved in the design of the policy or the choice of password strength estimator. We were approached by the IT services department who were eager to collaborate on exploring the scientific value of their policy design and its impact on the system’s users. This led to a productive working relationship for this project, which helped us to reason about the results and discuss possible causes for data patterns outside of the password system itself. This is especially important given the complexities not only of the data and the systems to which the data applies, but also the institution, being that it has tens of thousands of account holders with varying levels and modes of interaction with the system.

The main contribution of this work is a scientific analysis of the effect of the policy. The analysis is informed by consideration of the cost of the policy to users.

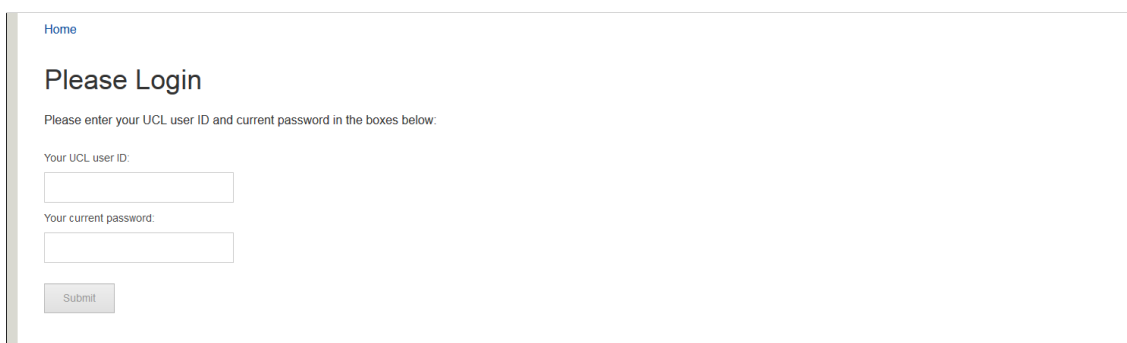
#### 7.3.1 The interface

The system starts at <https://myaccount.ucl.ac.uk/> as shown on Figure 7.1. The first option listed opens the log-in website shown in Figure 7.2, and after successfully logging in, forwards you to Figure 7.3. In order to register a new phone, the user also has to log-in. So after filling out Figure 7.2, the user is redirected to the page shown in Figure 7.4. If the user has forgotten their password, clicking on the third link opens the page shown in Figure 7.5. Lastly, there is also an option to retrieve the user’s userid.

## 7: PASSWORD PROJECT



**Figure 7.1:** Screenshot of the myaccount home page at <https://myaccount.ucl.ac.uk/>.



**Figure 7.2:** Screenshot of the login page at <https://myaccount.ucl.ac.uk/changepw>. The header and footer has been cropped.

The login page is shown in [Figure 7.2](#). After successfully entering a username and password the system loads the requested page.

[Figure 7.3](#) shows the page shown when the user requests a password change. When typing a new password, the password strength meter will automatically change colours and provide written strength feedback. The new password has to be typed

Home | Logged in: Ingolf Becker | Logout

## Change your password

Current password:

New password:

Confirm new password:

Password strength:  
**Weak**

### Get the most from your Password:

You can set a password of between 8 and 30 characters. The stronger it is, the longer you can keep it. One suggestion is to use 4 or more words. There are still a few rules and you will need to include 3 out of the following:

- Lowercase characters
- Uppercase characters
- Numbers
- Symbols, i.e. ! % ^ \* ( ) \_ + - = \* \* ' : < > , ? / @ \$ & [ ] { }

You can not use the symbols: ~ : | £ \

Avoid accented characters, characters from non-English scripts, symbols (other than those listed above), Emoji, etc. (This is because some of our older systems can't work with those types of character.)

Your password can not contain your user ID, forename or surname and you can not use a password you have previously used.

**Note:** If you are using some older UCL Administrative Systems, your password should additionally NOT use: the @ (at) symbol or start with either a space or " (double quote). If your password contains these characters and you have problems logging into some systems, try resetting your password to one that doesn't use them.

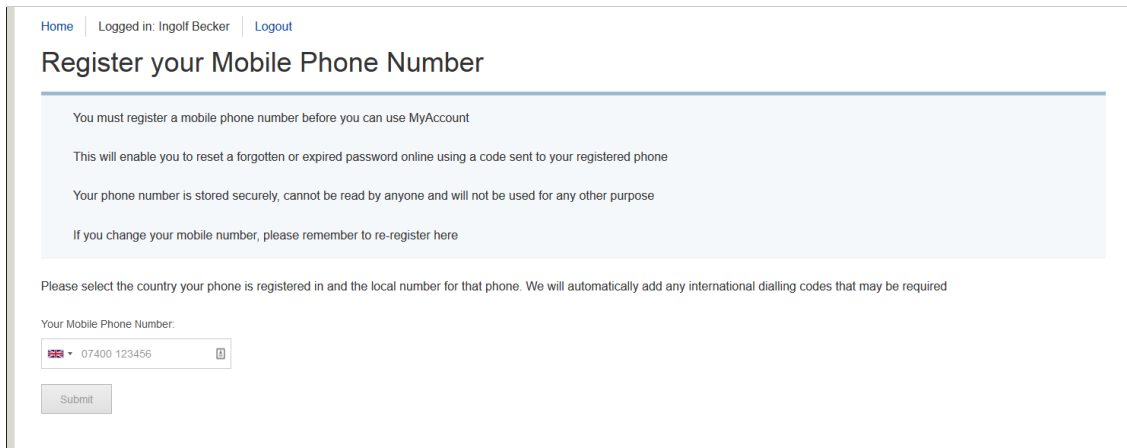
**Figure 7.3:** Screenshot of the change password website at <https://myaccount.ucl.ac.uk/changepw>. The header and footer has been cropped.

twice. Below the second password entry box are a password strength meter and a text field that displays the new password's lifetime in days, for example **This password can be used for 175 days**. Both meter and days of password lifetime update on any change to the first new password form field. For passwords of < 50 bits of entropy the strength meter states *Too weak* and the password cannot be submitted. Passwords of lifetime 100 to 163 days are stated to be of *Medium* strength (yellow strength bar). Between 164 and 223 days a password is considered to be *Strong* (green bar), and beyond that the password is classed as *Very strong* (dark green bar).

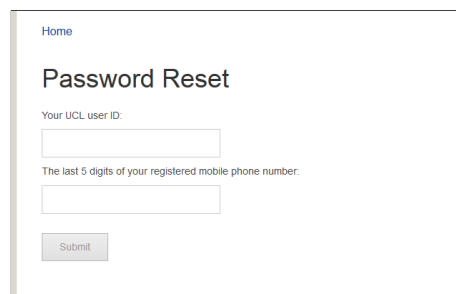
**Figure 7.4** allows the user to enter their mobile phone number in order to register for the password reset self service system. Upon clicking submit, the system sends a confirmation code that the user has to enter on the subsequently loaded webpage (not shown here).

**Figure 7.5** assuming the user has registered with the self-service password reset system, the user can use this page to begin the password reset process. Upon entering the user's userid and the last 5 digits of their mobile phone number, the system sends a text message with a password reset token to the user's mobile phone. The user then enters this code as well as their new choice of password on the subsequently loaded screen (similar to **Figure 7.3**, but with the 'current password'

## 7: PASSWORD PROJECT




**Figure 7.4:** Screenshot of the register your phone website at <https://myaccount.ucl.ac.uk/registerphone>. The header and footer has been cropped.



**Figure 7.5:** Screenshot of the reset your password website at <https://myaccount.ucl.ac.uk/resetpw>. The header and footer has been cropped.

input box replaced with an input box for ‘reset code’).

### 7.3.2 Empirical evaluation



**Figure 7.6:** Empirical analysis of the password strength validation function.

Figure 7.6(a) shows a clear linear relationship between the password expiration

and the (reported) password entropy for entropies between 50 (100 days) and 120 (350 days). For passwords with entropy below 50 the password rejects the system, even if it confirms to the written guidelines. For passwords with entropy above 120 the expiry is fixed at 350 days.

**Figure 7.6(b)** plots the relationship between the length of the password and its expiry. It is clear that length of a password does directly relate to the expiry of a password. For example, the password `9W,79R02,702` is 12 characters long and satisfies the states password requirements discussed, but with an entropy of 48 it does not qualify.

In order to achieve the maximum password lifetime of 350 days, a password needs to be at least 20 characters long. Yet longer passwords do not necessarily achieve this maximum lifetime: `58239858204928e36640904538?97` is 29 characters long but with entropy 86 only achieves 229 days lifetime.

### 7.3.3 The dataset

We received access to the password change and reset logs, which consisted of timestamps, anonymised user IDs, action performed (i.e., change/reset/etc), the integer password lifetime of the new password (100–350), as well as some coarse demographics information for the 100,000 users. We received IRB approval for our approach to log analysis, alongside in-person interviews with a subset of system users (see [Section 7.4.7](#)) (UCL Ethics ID 5336/007). Regarding the dataset, we had no individually-identifying information (an arrangement made with the system owners at point of data access), as well as only a single number for the user’s password strength (i.e., not the password itself or any element of it). The password log data was stored on encrypted drives, and regular extensions to the dataset over time were transferred and stored securely.

The policy came into effect in October ’16 and users began using the new system from that date when next requiring to change or reset their password. As the previous policy’s expiration was set to 150 days, all active passwords will have been transferred to the new policy by April ’17 (so that in effect it was a soft transition). Although we continue to have access to new data, we are confident that 14 months of complete data is sufficient, for the following reasons:

- The dataset includes at least one academic year’s worth of data and regular events in an academic year, such as school closures and holidays;
- All currently active passwords were set on the new system;
- There are approximately six months of system events for the annual intake of new students (academic year starts in September to October, as seen for

instance in [Figure 7.8](#)), who were never exposed to the previous policy.

### 7.3.4 Calculation of entropy

The minimum password requirements involve a complex combination of a number of fixed rules. Passwords are initially checked against static requirements. Passwords are required to: include at least one character from three of four possible character types (lowercase character, uppercase character, number, and symbol); be between 8 and 30 characters long, and; not contain the user's username or parts of their real name. The entropy of a password is then calculated by estimating the information entropy of the password by multiplying the size of the character class of each of the characters ([Alistratov 2010](#)). A number of factors decrease the entropy: repeated characters; lexicographically subsequent characters as well as the presence of a substring of the password in a dictionary of size 306,000. Common character substitutions are also checked against the dictionary.

### 7.3.5 Uses of a password

Studying adoption and use of the system over time is important, where understanding new authentication systems in terms of how easy they are to learn is critical ([Bonneau et al. 2012](#)). The password studied should be the only password staff and students require to access necessary services for work or study respectively. UCL uses one password for all of their services. This includes access to timetabling, e-learning resources, university e-mail, logging on to physical desktop machines, and WiFi. The frequency of use of this password is expected to vary naturally for different user types, who use different services, and access them from different machines (the most simple differentiation being a device they manage themselves or a fixed-place common-access machine). While users may resort to password managers to store their password for use in browsers, students (Undergraduate, Postgraduate and Medical) accessing the machines in university computer rooms will still have to type the password. Similarly, administrative staff work on a university computer and therefore have to regularly type the password to log in to and unlock their machines. Research staff and students however may have the flexibility to type their password very infrequently, especially if (a) they are using devices which they themselves manage and which no other user would have access to, and (b) they can complete their work or study activities with minimal or ad-hoc access to services managed through the single-sign-on system. Ad-hoc access may be governed by the nature of the work done by distinct specialised groups, hence we are also interested in adoption and use differentiated by faculty/department. Users may then

balance the convenience of accessing a system with the security of the mechanism that facilitates access to that system (Beautement, Sasse and Wonham 2008).

### 7.3.6 UCL's threat model

Given the universal use of the password for staff and students alike, compromised accounts pose a significant threat to the organisation. Staff accounts have access to significant amounts of personal information as well as inside knowledge of grades and future assessments. All accounts have access to exclusive material, and universities have previously been threatened with access restrictions to academic publishers due to account misuse. From an ideal security perspective, UCL would deploy 2-factor authentication to reduce the reliance on passwords alone for user authentication. However its implementation is not trivial due to an inconceivably diverse system of software and hardware environments and use cases present across over 100 departments, institutes and research centres. There is a steady push to simplify these systems, however password changes still appear to be propagated through scheduled synchronisation tasks. This makes a principled approach to secure authentication difficult, and given the risks UCL faces, the organisation's leadership has pushed for a policy that encourages stronger passwords for each individual user.

### 7.3.7 Perceived value of a password

Individuals in organisations will strive to protect their account if they perceive and understand a need to keep their organisation secure (Adams and Sasse 1999). The UK's Universities and Colleges Information Systems Association (UCISA) distinguishes between the information security roles and competencies for distinct groups in universities (UCISA 2015). Assuming that system users are aware of responsibilities like those described in the guide, they may have distinct attitudes towards the security of their accounts, and the associated passwords. Researchers may for instance have access to sensitive data, whereas administrators and teaching staff alike may manage staff and student records. Students may have access to their own information, but also the university's IT infrastructure; postgraduate students might have access to research data.

By considering factors which may influence the perceived value of a user's password, the scope of RQ2 is refined. Given both the frequency of use and the perceived value of accounts, we expect students to have weaker passwords than other groups, and researchers to have stronger passwords. We also expect administrative staff to value their account security while balancing any increases to password strength (delaying password change) with lower time cost per system authentication event.

Regular enactment of security tasks over a working day may push users in an organisation to find ways to reduce the burden of security that relates to their primary productive work (Beautement, Sasse and Wonham 2008). We test these hypotheses in [Section 7.4.4](#).

### 7.3.8 User interviews

In addition to the password log analysis, 93 users of university systems were interviewed between February and March '17 (53 students and 40 staff). Users who had changed their password in the prior 2–3 months, or who had just received a reminder to change their password, were invited for interview. This framing allowed for the possibility that participants would not know that there was a new password policy.

The study was advertised via staff and student newsletters, and flyers positioned around the main university campus. Interviews were approximately 30 minutes in duration, and included discussion of: services accessed through university login; perceptions of passwords and security in relation to university-related tasks, and; participants views of the university's password system. A computer displaying the interface of the new system supported the interview (as described in [Section 7.3.1](#)). Participants were provided with a £15 voucher for completing the interview.

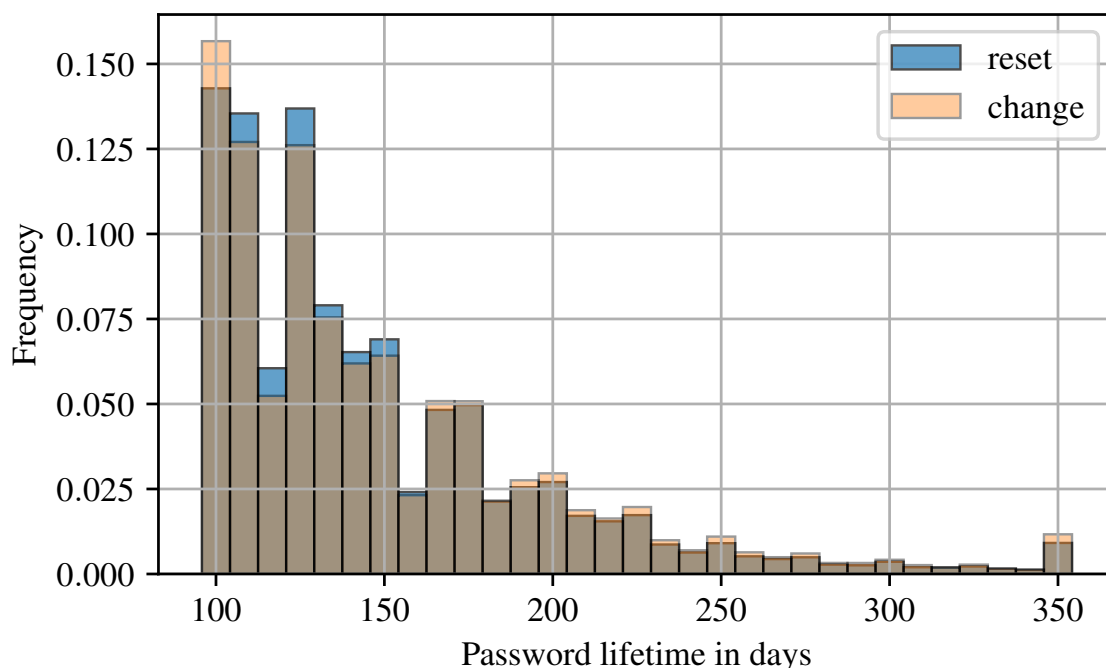
The average participant age of staff members and students were 34.6 and 22.8 respectively. Student participants had been at the organisation on average for approximately two years (including many who had joined the university just before the new system was deployed); staff participants had used the university systems for on average of approximately five years. Participants represented a range of schools and divisions (including administrative functions).

## 7.4 Results

In this section we describe the properties of user passwords found in the data, as well as characterise the adoption and usage behaviour for the new system across the user population and specific groups. We put our results in the context of existing research and highlight the impact of the policy on user behaviour.

[Figure 7.7](#) describes the distribution of strength of all passwords observed in the university. The two distributions of password resets (when a password has been forgotten or it has expired) and changes (when the user still knows the previous password) are virtually identical. The histogram is strongly skewed to the left





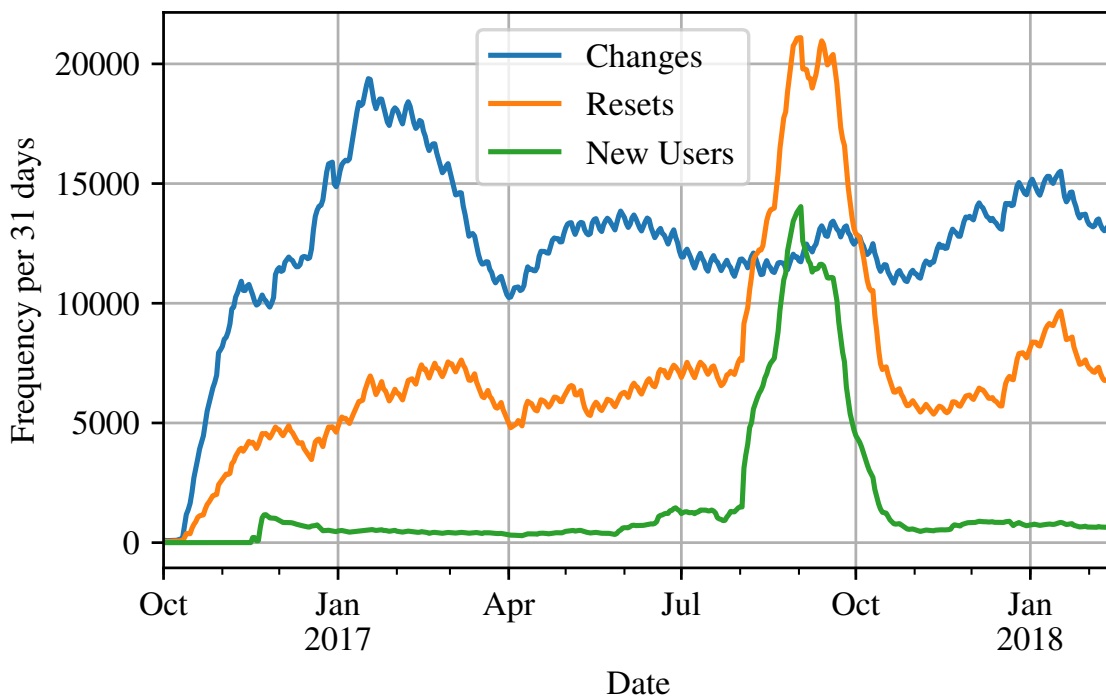
**Figure 7.7: Normalised frequency of password lifetime** The mean frequency is 147.74 and 146.60 days for changes/resets respectively.

and decays rapidly, apart from approximately 1% of passwords that achieve the maximum strength of 350 days.

It is interesting to compare this distribution to the password strength distribution of Mazurek et al.’s study performed at Carnegie Mellon University (CMU) (Mazurek et al. 2013, Figure 7, page 11). Their measured password strengths approximate a uniform distribution between  $10^9$  (100 days) and  $10^{14}$  (225 days) guesses, and only 42% of passwords are guessed in  $10^{14}$  guesses. Their estimated mean password entropy is 36.8 bits, compared to 69.64 bits here.

There are two systematic explanations for these stark differences. First, the mean password entropy reported by Mazurek et al. is calculated by state-of-the-art brute-forcing, compared to an information theoretic approach chosen by our IT department that only weakly correlates to actual password strength. Thus, our entropy estimates are likely large over-estimations (Wheeler 2016, Fig. 8). Secondly, the entropy estimate in our analysis is the same estimate used for providing feedback to the user in the form of the password meter (principally the fullness of the bar), and the weakest allowed password has an entropy of 50 bits. This explains the high concentration of passwords with 100 days lifetime, compared to the study performed at CMU; where policy and strength meter are not linked to the measured guessing strength.

The same explanations also apply to the differences between our analysis and




**Figure 7.8: Number of password changes and resets over time** 31-day moving average of the number of password changes and resets, as well as the number of new users joining the university and using the system for the first time. The legend is in order of final values.

Bonneau’s analysis of cracking attempts of the Yahoo! password dataset (Bonneau 2012, Figure 6 in particular). Their identified cumulative distribution is aligned with our data, although Bonneau achieves a 50% success rate with  $10^6$  guesses.

#### 7.4.1 Noteworthy events during the study

As with any study of an active real-world system, there are external events that have an effect on the system being studied. As we cannot control for these events, they should be acknowledged in the analysis. Further, external events can be leveraged to understand if there are particular kinds of events which can influence the adoption and use of an authentication system at a large organisation. Figure 7.8 highlights three families of events.

From the deployment of the new system in October ’16 the userbase of the new system slowly grows as users change or reset their passwords (where this forces them to use the new system and hence appear in the dataset). Secondly, there is a peak of password resets in Jan-Feb 2017, which corresponds to the expiration of all passwords of users who joined the university in September ’16 and had a fixed lifetime of 150 days. We expected that the rate of resets would decrease once



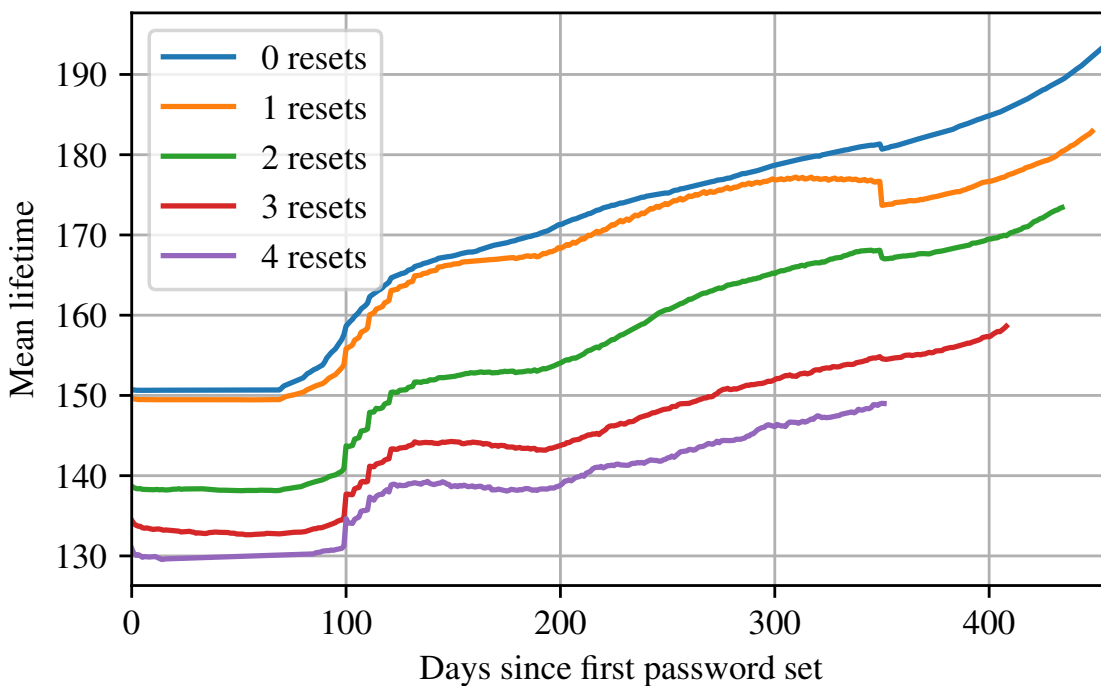
**Figure 7.9: Distribution of the number of changes and resets the users in the dataset have made.** Mean frequency is 2.41 and 1.08 for changes and resets respectively. 66% of users have reset their password at least once.

users became familiar with the new system. This did not happen, indicating that familiarity with the system does not reduce the need to reset. The third event of note refers to the peak of new user being onboarded to the system in September '17 in time for the new academic year, where over 10,000 new students joined the university. This also causes the simultaneous peak in the number of changes, as setting an initial password is classified as a change.

### 7.4.2 Password change behaviour

The effect of the password policy on changes and resets is shown in [Figures 7.9](#) and [7.10](#). In the full period studied, more users (66%) had to reset their password than not: on average, a user had to reset their password 1.08 times. Users may have to reset their passwords for two reasons: if they have forgotten their original password, or if their password has expired. The cost of a reset is significantly higher than a change, as it requires either physical presence at the institution's help desk or using a phone-based reset system. Over the period studied, the mean number of password changes and resets per user is 3.5. This is investigated further in [Section 7.4.3](#).


There is a strong positive correlation between each user's previous password



**Figure 7.10: Average password lifetime of unexpired passwords by number of password resets.** After 100 days the weakest passwords expire and users choose stronger passwords, which accounts for the steep rise. This pattern repeats after another 100 days. At 350 days users change their previously strongest passwords to one that is as strong or weaker password, causing a pronounced dip in the average password expiration.

strength and the likelihood of that same user resetting their password before expiration (i.e., forgetting the password, Spearman’s  $\rho = 0.95$ ,  $p < 10^{-15}$ ). A user with a password lifetime of more than 300 days is four times as likely to forget their password than a user with a password with a 100 day lifetime. The minimum reset frequency per day of actual password lifetime is achieved with passwords which have a 100 day lifetime. Most resets however occur shortly after passwords have been set, and not after a user has been using a password for 100 days. Having a relatively strong password on the system then incurs the additional cost of potentially needing to reset that password. This may not only negate the advantages of having a strong password in the first place, but results like these can also inform predictive helpdesk/support provisioning (Parkin et al. 2010), i.e., if users are encouraged to maintain stronger passwords, they may require more helpdesk support to reset passwords.

This is in contrast with Figure 7.10: the more password resets a user will have had, the weaker their password choice. While the average password lifetime of all groups is increasing as the users renew their password, the division between users with 0 or 1 reset and users with more resets remains pronounced, separated by at

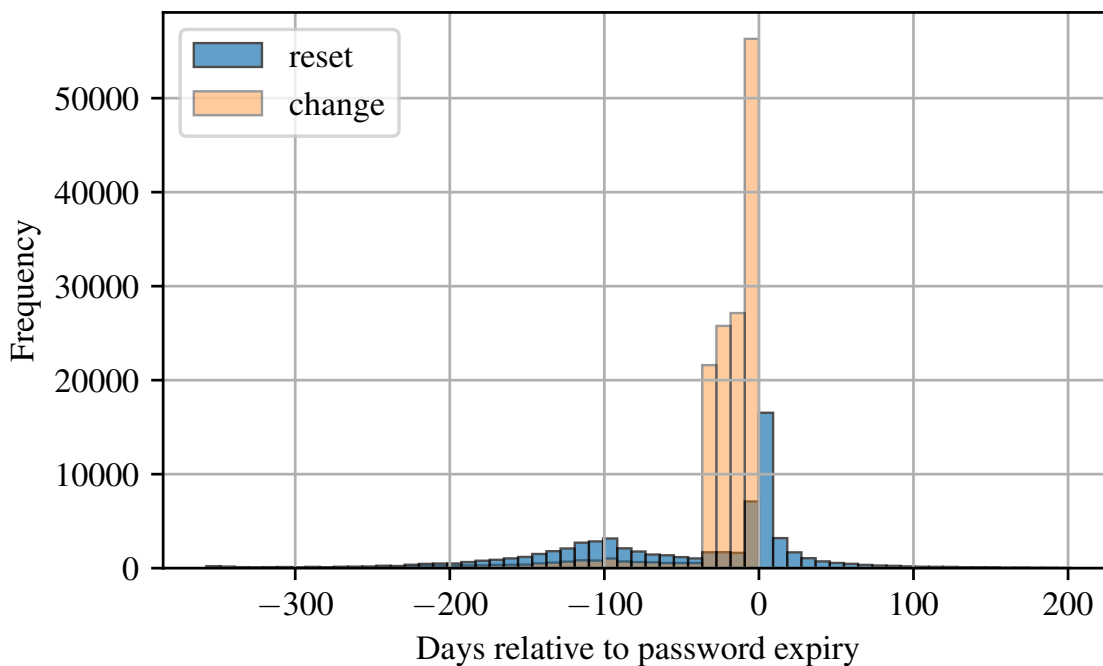


**Figure 7.11: Distribution of the change in the password lifetime after the password change/reset** Mean change is 11.97 and 4.55 days for changes and resets respectively.

least 10 days of lifetime. This analysis suggests that one reset per year does not affect the system’s performance, but two or more resets do (which applies to 27% of users). While system owners should obviously try to minimise the number of resets required, it appears one reset per year per user is an acceptable upper bound.

The answer to our first research question is alluded to in the mean password strength change of 12.73 days (as shown in Figure 7.11). This shows positive increases in password strength on consecutive password changes and resets on average. One common finding in password expiration research is that when forced to change one’s password, the new password will be similar to the old one. Figure 7.11 indicates that this effect may also be present in our dataset: 20% of changed passwords have identical expiration as their previous password, and 36% vary within 3 bits of entropy.

These figures vary slightly during the period of time analysed here, with a gradual increase to 28% in February (3 months after the change in policy) but returning to 20% in June and remaining constant from then on. Prior literature has examined this behaviour: Adams and Sasse found that 50% of their participants varied some element of their password when creating new passwords. Zhang, Monroe and Reiter study behaviours at greater scale, by analysing 7,700 accounts and developing an efficient transformation algorithm to test for related passwords. The authors are

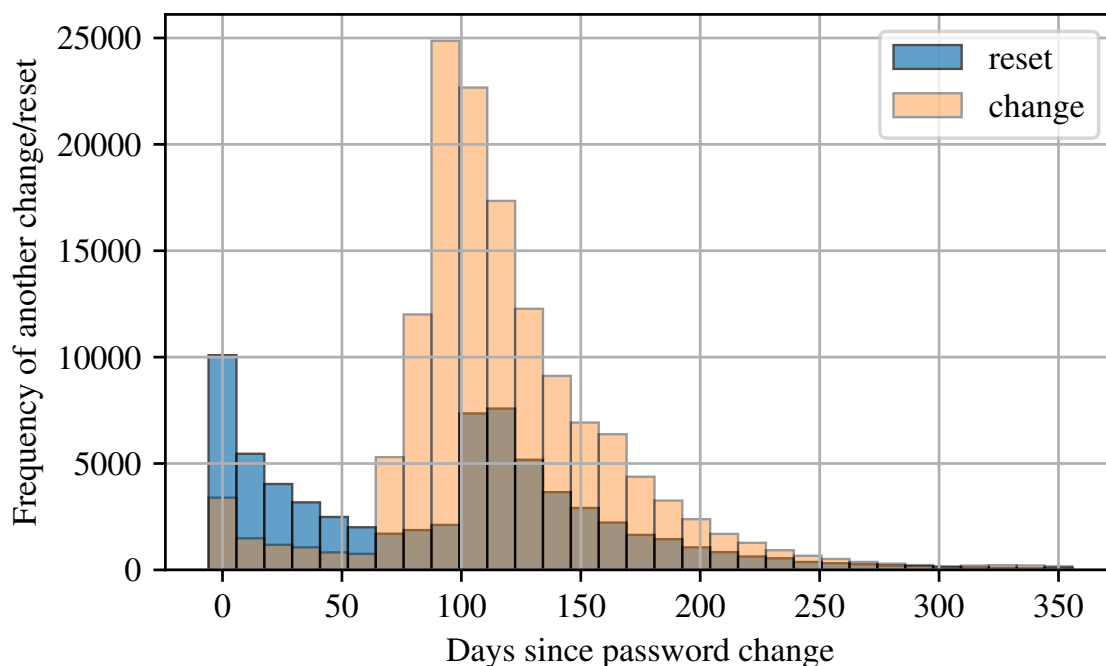


**Figure 7.12:** The frequency of password changes by the number of days relative to password expiration (day 0). The mean time for changes is  $-22.18$  days and the mean time for resets is  $-52.09$  days.

then able to break 17% of their accounts within 5 guesses, and 41% within 3 sec of CPU time ( $\approx 10^7$  guesses, our estimate) (Zhang, Monroe and Reiter 2010). While we cannot determine the true dependence between current and prior passwords in our dataset, the strength proxy (through Figure 7.11) may suggest a similar proportion of related passwords.

### 7.4.3 Time dependence of subsequent changes/resets on prior lifetimes

Users are sent an email reminding them of their password’s impending expiration 30, 20, 10, 4 and 1 day(s) in advance. The effect of the reminder is shown in Figure 7.12 with a bin size of 10 days. 10% of users act upon the reminder on average within 24 hours and subsequently change their password. Each following expiration warning causes an immediate increase in change rates, with the largest peak on the day of expiration, where another 13% of users change their password. This is followed by users resetting their passwords immediately after expiration, presumably after having been denied access to university resources. The general effect of these frequent reminders for the organisation is that the average user changes their password 22 days before expiration, essentially reducing the lifetime of their password voluntarily. This indicates that users in this institution change or reset passwords in response to reminders, and seldom voluntarily. This might be the case for users

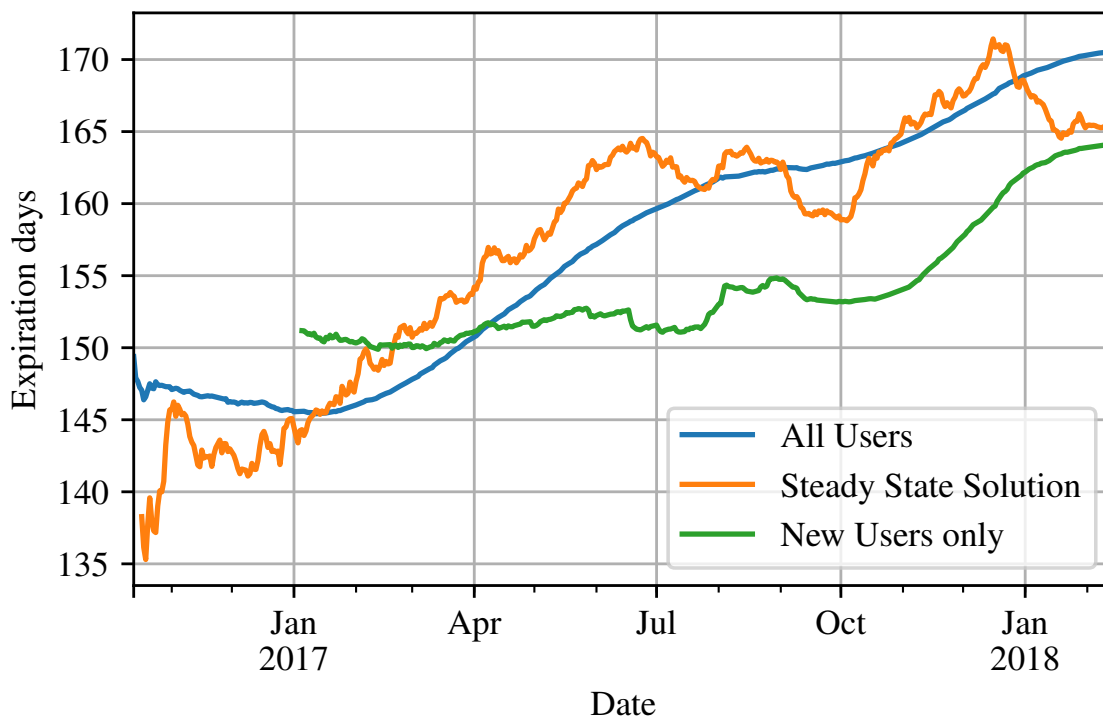


**Figure 7.13:** The distribution of the time between consecutive password changes. The mean time for changes is 117.16 days and the mean time for resets is 90.48 days.

changing their password before even receiving the first 30-day advance warning of expiration, as can be seen in [Figure 7.12](#).

[Figure 7.13](#) is an analysis of the same time series as [Figure 7.12](#), but anchored at the time of password creation rather than expiration. The main observation here is the strong concentration of password resets in the immediate proximity of password creation: users often forget their newly set password. The passwords created by reset within the first 48 hours after changing a password have a mean password strength of 6.9 days less than their previous password. This suggests that some users choose a weaker password due to forgetting the previous one (where in fact some users may be choosing weaker and weaker passwords in a cascade). The change rate initially decays before exhibiting the shape of a gamma distribution starting at 70 days at the time of the first expiration warning email for passwords of 100-day strength, peaking at just before day 100, when a large number of passwords expire.

These results imply that users reset their passwords primarily for two reasons: failure to recall the password, and the forced expiration of the password by the system. This is in line with personal password behaviours observed elsewhere ([Inglesant and Sasse 2010](#)). These drivers are in contrast to instances where users would reset their password for primarily security reasons (such as believing that their password has been compromised).



**Figure 7.14: 31-day moving average of the mean password strength of all users and new users.** The ‘steady state solution’ estimates the average strength of passwords in the system if users were to continue making their passwords stronger (or weaker) consistently with how they did so in the current measurement window. The legend is in order of final values.

#### 7.4.4 Password change time series

In this section we study the password strength measure over time. The results answer two of our research questions: ‘What effect does the password policy of variable expiration have on user’s passwords? Given the freedom, how will users choose?’ (RQ1), and ‘Are there contextual circumstances of groups of users which may influence their choice of password strength?’ (RQ2). In [Figures 7.8](#) and [7.14](#) to [7.16](#) we apply the same 31-day moving window to smooth out fluctuations due to weekly patterns (e.g., weekends, when most users are not actively using the system).

[Figure 7.14](#) shows the evolution of the university’s mean password strength over time. Initially we observe a small drop in strength between November ’16 and February ’17 (after the adoption of the policy), as users become accustomed to the new system. After this, the mean strength increases from 145.5 days to 170.1 days, an increase by 6.9 bits of entropy. This strongly suggests that users have adapted slowly to the new password policy, and eventually make use of their ability to increase password lifetime by strengthening their passwords.

The ‘steady state solution’ is an approximation of the attractor of the password



change distribution. It is calculated by performing a linear regression on users' previous ( $x$ ) and new password lifetimes ( $y$ ). The solution of this linear regression for  $y = x$  identifies the attractor. Users with previous passwords weaker than this attractor tend to reduce the lifetime of their new password, and vice versa.

The evolution of the mean password strength is underpinned by cyclical behaviours. A quarter of users have a password lifetime of less than 110 days (see [Figure 7.7](#)), and have to change their passwords on average every 80 days (see [Figure 7.12](#)), but every time they do, they increase their average password strength. This manifests twice in [Figure 7.14](#): at the start of the deployment of the new system where there are no existing users (the increase in password strength is delayed until February '17); and again with the enrolment of over 10,000 new users who set their first password around September '17 (see [Figure 7.8](#)), in time for the start of the new academic year. As this large number of users have all set their initial passwords in a short time frame, their first regular password change occurs from November '17 onwards. Their change behaviour also causes the temporary plateau around September '17 and the subsequent increase of the mean password strength of all users, which is a statistically significant increase (paired t-test,  $t(10892) = -47.19, p = 0$ ).

The 'steady state solution' gives us insights into the password changing trend over time: for example, if users had continued to choose new passwords in the same manner as they did in April '17, the mean password lifetime of the university would settle at 156 days. However, as the steady state solution continues to increase, it appears that the users are still responding to the policy. The artifacts of the cyclical changes are also evident in the trend.


The relatively small drop in the steady state solution after January '18 aligns with an increase in password resets at this time (see [Figure 7.8](#)). This could be due to users having forgotten their passwords after returning from the Christmas break. As new users have yet to catch up to the password strength of existing users, it is likely that the mean password strength in the university will increase further.

As we do not have data for the users' password strength before the adoption of the new password change system and policy, we are unable to do a rigorous before-after comparison of strength data that takes into account all factors that may have contributed to this change, for example the old system did not give any feedback on their password strength. This implies that interface design for the password creation/reset process may also have a part to play in users increasing their password strength (where a subset of users migrating between the old and new systems provided feedback in [Section 7.4.7](#)).

As the new users have not had experience of the previous system, and as there have been no other initiatives by the university to encourage stronger passwords, we

consider the increase in users' average password expiration likely to be a consequence of the policy, answering RQ1. It appears to have taken around 150 days for the effect of the policy to start to achieve its aims.

#### 7.4.5 Password change time series by school



**Figure 7.15: 31-day moving average password expiration for selected schools over time.** The legend is in descending order of the final expiration values.

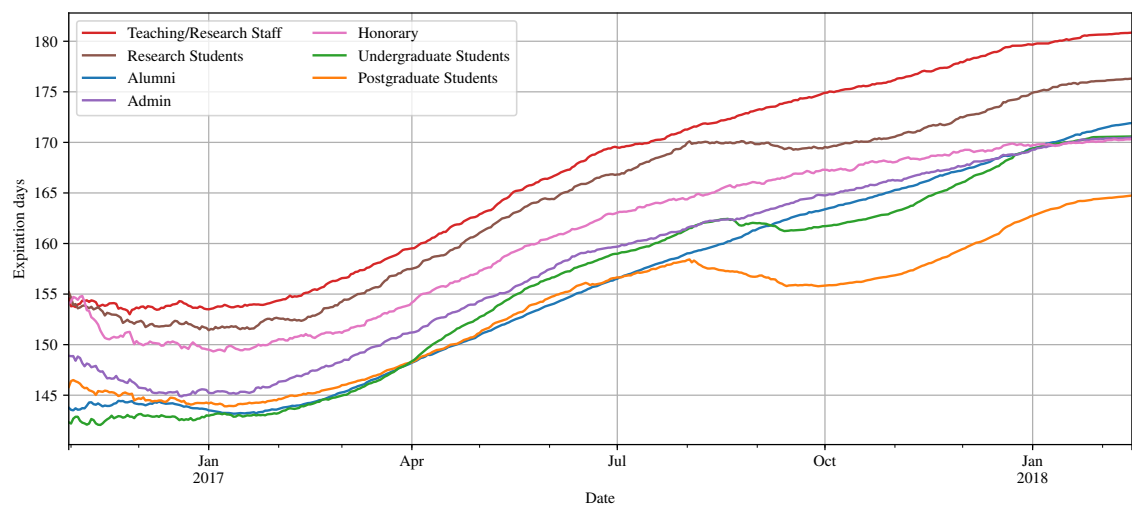
We are fortunate to have some coarse demographic information for each user recorded in the data. Figure 7.15 compares the evolution of password strength for selected schools. The users of each school have together made at least 11,000 password changes; we calculated bootstrapped, bias-corrected and accelerated (Davison and Hinkley 1997) confidence intervals for each of the schools. The 95% confidence intervals were within 1% of the mean for all schools in Figure 7.15 from January 2017 onwards. We have hence omitted the confidence intervals. For brevity, we omitted a number of smaller schools closely aligned with the university mean.

Throughout all schools there is a statistically significant positive increase in password strength (in-sample t-test,  $p = 0$ ). The school of Education displays the lowest increase of 18 days, while Maths and Physics increased their password strength by 27 days. The differences between schools are also pronounced, with passwords in Engineering being 13.4 days (4 bits) stronger than in the school of Education. It is of note that the university's Education school has been part of the university for only a few years. A joint linear regression of the password strength changes of all faculties predicting the password strength was conducted. Each school contributed statistically significantly, explaining 82% of variance ( $R^2 = 0.816$ ,  $F(6, 49201) = 36320$ ,  $p < 10^{-10}$ ).

In previous research, only Mazurek et al. compare different university units for their respective password strength. Their password cracking algorithm managed to predict in  $3.8 \times 10^{14}$  guesses the passwords of 38% of computer science accounts and 61% of business school accounts. They then performed a Cox regression on password survival times, reporting a 1.83 times chance of password compromise for business school passwords than for computer science.

In a naive model,  $3.8 \times 10^{14}$  guesses could be estimated as fully eliciting 48.43 bits. Given that the weakest allowed password in our university has an entropy of 50 bits, we expect 2.59% of Engineering accounts and 2.92% of School of Education accounts to be compromised after  $3.8 \times 10^{14}$  guesses. If we increase the attacker's brute force capacity to 60 bits ( $10^{18}$  guesses), the expected proportion of accounts which may be compromised increases to 36% and 44% respectively. In either case School of Education passwords are 1.13 and 1.22 times as likely as Engineering passwords to be guessed.

#### 7.4.6 Password change time series by relationship



**Figure 7.16: 31-day moving average password expiration for various relationships with the university over time.** The legend is in descending order of the final expiration values.

In addition to an analysis by school/faculty, we are also able to differentiate between the different roles of individuals within the university. The evolution of the respective user group's password strength can be found in Figure 7.16. Relationships with less than 5,000 / 2% of the total password changes/resets have been omitted. As for the previous graph, all user groups show an upward trend in their password strength over time. There are also significant variations between the groups, with

Teaching/Research staff exhibiting password strengths 21 days stronger than Post-graduate students. A linear regression predicting the password strength depending on the relationship types was carried out. Each type of relationship contributed statistically significantly, explaining 89% of variance ( $R^2 = 0.893$ ,  $F(13, 12559) = 7957$ ,  $p < 10^{-10}$ ).

The differences are in line with the hypotheses in [Section 7.3](#): there appears to be both a positive correlation between password strength and likely value attached to the account (see [Section 7.3.7](#)), and a negative correlation between password strength and frequency of use. For example, Teaching/Research staff are likely to value their account security highly (using their accounts to access research and teaching data, which undergraduate students for instance would not. We observe that this group has the highest average password strength.

Administrative staff may value their account security highly too, but they also have a high frequency of use of the password, which may act to moderate their password strength. An interesting group to investigate in further research are the Alumni. These users are very different to the rest of the population: their account usage is low, so a long password expiration time will help minimise the frequency of password changes/resets; being potentially remote to the university, they may perceive the potential cost of a forgotten password as being much higher.

The results presented in this section answer our initial research questions: users have responded to the freedom of choosing their password lifetime slowly, but have in time increased their password lifetime considerably. The user population has needed time to adapt to the change in authentication protocols; 14 months after the intervention, the password strength of all user groups has yet to plateau. We have identified differences in how users react to the policy change, by analysing the evolution of password strength between different subgroups (role and division). Other work has demonstrated that security preparedness and perceptions can differ between roles and divisions in a large organisation ([Beautelement et al. 2016](#)).

#### 7.4.7 User feedback

Here we present a preliminary summary and discussion of field notes taken by interviewers (see [Section 7.3.8](#)). Feedback from the 93 interview participants informs the view of factors which may influence decisions around the construction and use of passwords on the studied system. We discuss general observations, with representative participant quotes. Participant identifiers signify E### (Employee/Staff) or S### (Student).

A few participants reported changing their password-related habits in response

to the new system. This included beginning to store the password in a password manager, or as with E19, making a written note:

*‘Well, normally I just memorise it. This time around I did actually write it down when I changed it last week. Because it was so much longer than normal. Because previously they were eight characters. Now I think my password is like twelve characters. And it had to be that long to get the security up too. Because of now they rate it like low, medium, strong securities. So I had to keep adding characters to get it to say strong. So it is longer than I normally have.’* (E19)

Many participants appreciated the flexibility of the new password policy. Some had however used the new system but not explored the differences between it and the old system; the differences between systems—and policies—were not immediately apparent to all participants. With the introduction of the new system, participants were split as to whether they believed passwords should be expired or remain valid indefinitely.

There was a general even split among interview participants as to whether they saw a link between password age and password strength. The data supports this, as a year after deployment users’ average password strength has yet to settle (as notable in [Figure 7.14](#)). This could potentially be as much about discovering the features of the new system as it is about skillfully using it. For those who were aware of it, some did see it as an incentive to make a stronger password, such as E20:

*‘If they say if you make a stronger password you can keep it for longer, maybe it would help. [...] It wasn’t clear that it was contingent on the strength of your password. I don’t know if it is.’* (E20)

Conversely, E25 found it difficult to create a valid password that was not labelled ‘Weak’:

*‘I probably tried about 6–7 passwords before I got to the one that it would accept [...] It [the password meter] just kept not getting past the failed point [...]’* (E25)

Others would consider password length alongside the need to type the password many times, and as a result would aim for a ‘Medium’-strength password of around 8–12 characters. E30:

*‘Or trying to find a better password that would work. It does get harder because I had to change it so many times [...] trying to think of a password. In a way it is not good that you are supposed to change a*

*password. You run out of ideas of what to use. It's good that they are aware of your security but it does get a bit stressful.'* (E30)

A number of participants commented that although they had created a longer password than before, they immediately reset their password as they found it too complicated to type, such as S17:

*'Even though I could remember it wasn't practically very helpful if you have to put in you know twenty characters. It's not great. So then I changed it to something that was shorter and last a little less time I just could remember that.'* (S17)

This aligns with our findings in [Figure 7.10](#) and [Section 7.4.2](#), and also with Mazurek et al.'s engagement with system users (Mazurek et al. 2013): those finding longer passwords unworkable will act to find a solution which is workable, abandoning the potential for longer lifetimes.

The summatory findings indicate that there may be a number of factors influencing password choice which are not represented in the dataset. The analysis in [Section 7.4](#) was based on the available *data*, and the available *data fields*. Future collaboration will explore how the design of password system logs can be augmented to provide a more directly holistic view.

## 7.5 Discussion

There are hidden costs of the change in policy that should be considered. The intervention took time to gain traction, and it may have been that this time could have been shortened in some way. In some cases, users were voluntarily changing their passwords to a weaker combination of characters, taking time to learn how to *skillfully* choose stronger passwords (i.e., sustain stronger passwords over successive change events). The analysis informing [Figure 7.10](#) uncovered that over 27% of users have had to reset their passwords more than once per year, and that these users have passwords with much shorter expiration. It could be that system usability hinders the adoption of the policy for a proportion of users.

As noted by Adams and Sasse (1999), most users in an organisation will want to behave securely, where insecure behaviour arises as they try to manage excessive demands in their workplace (where security would be just one of those demands). That the changes across different departments and user groups follow relatively similar patterns suggests that there was a collective change in password use, perhaps due to a collective culture towards security or influence from how peers are seen to behave.

From a security perspective, the implications of our results are clear. In the current format of the policy, the weakest possible password is strong enough to withstand an online attack (need to withstand  $10^6$  guesses); the increase in strength has not been pronounced enough to protect against offline attacks (Florêncio, Herley and Van Oorschot 2014a). Rather than improving robustness to a wider range of attacks, the intervention has identified each user's individual threshold for trading off password complexity for password lifetime. It is a combination of the subjective cost optimisation of the individual's time (time spent both resetting and authenticating), acceptance of the perceived effort in managing a complex password, and their perceived value of their account. As different individuals interact differently with the university, this optimisation varies across user groups, as in Figures 7.15 and 7.16.

From a cost-benefit analysis, the policy has increased cost through increased individual effort cost and organisational support cost due to resets. The benefits for users rely on their perceptions: our user interviews found that the possibility of longer lifetimes was welcomed, and perceived this as an improvement considering their previous experiences of organisational password policies.

Here we have considered the different contexts in which users interact with the password policy. A further hidden cost arises from the interruption of the primary task from expiration of passwords, the reminder emails, and the planning of when to next change one's password (as one might be about to travel or go on leave, for instance). In studying the use of passwords and support of users in a large organisation, Brostoff (2005) identified a range of 'costs' related to the expiry of passwords, such as designing new passwords, re-design of a candidate password if the system does not permit it, and amending any recall aids such as written notes. Brostoff's results also suggest that users may confuse prior and current passwords, where having had expired passwords then contributes to the daily cost of entering a current password correctly. The extra reward perceived for a stronger password must be greater than the cumulative additional time (i.e., perceived effort) required to correctly enter the password when it is needed. This is to say nothing of the frustration that may be caused in recalling and entering passwords, and the batching of tasks that may occur to reduce the regularity of password entry events (Steves et al. 2014). A similar approach to the work described in Steves et al. (2014), of asking users to complete diaries or otherwise report on their experience of using the system, may more clearly identify the workload caused by the authentication system.

### 7.5.1 Limitations

Our main limitation stems from studying passwords ‘in the wild’: our study did not have a control group. This means we are unable to observe if users would choose stronger passwords without the presence of the greater lifetime incentive. However, the existing literature (Zhang, Monroe and Reiter 2010; Zhang-Kennedy, Chiasson and Oorschot 2016) suggests that users choose new passwords that are similar to previous ones, rather than continuously act themselves to improve the strength of their password.

We did not have log data for users prior to deployment of the new system. However, new users who were unaware of the old system behaved similarly to the existing population, suggesting that effects are due to the new policy rather than the change in systems.

## 7.6 Conclusion

Here we evaluated the impact of a new password policy upon 100,000 users at our university. In what is a novel policy designed by system managers, users were able to choose passwords with lifetime varying from 100 (50 bits of entropy) to 350 days (120 bits of entropy).

While the security community is moving away from prescribing password expiration, we have found that users ‘play the game’ and adapt their passwords in order to receive longer lifetimes. Results show that the intervention took over 100 days to gain traction, and that users took over 12 months to move from a lower-than-initial average 146-day (63 bits) to a higher 170-day (70 bits) password lifetime. The policy had both apparent and potential costs for individual users: 66% of users had to reset—as opposed to routinely change—their passwords, often multiple times. The average user had 3.5 passwords over the duration of the study. Users who are forced to reset their password more than once a year compensate by choosing significantly weaker passwords. Depending on the implementation of the reset procedure, both the actual and user-perceived cost may be high.

The analysis has revealed different levels of engagement with the policy. Had the system been monitored more directly for the impact upon users, the high reset rate and varied degrees of adoption amongst different user groups could have been seen as *early indicators* of the need for further support. It should also be noted that the policy intervention described in this chapter gave users a choice in balancing delayed expiration and cost of managing a stronger password, rather than forcing the policy on them (Inglesant and Sasse 2010). We continue to work with the system



managers to analyse new log data, and to explore how user needs and challenges can be anticipated.

### 7.6.1 Policy interventions

One take-away here is that conclusions about the impact of an intervention should not be drawn based on immediate improvement or lack thereof. Other studies of the impact of behaviour change caused by security policies—in particular, lab studies—should measure interventions at meaningful intervals over a suitably long period of time, where arguably this would be a continuous activity.

When designing a new intervention, practitioners should consider how to measure the effectiveness of a change and the associated impact on users. After an intervention is deployed it may benefit from being monitored and *calibrated*, towards reducing problems and reward secure behaviour, where dynamic policy that reacts to users is far from being a common capability.

We have found that users will generally change their password in response to password expiry warnings and reminders; warning users too early effectively reduces the password lifetime. This potentially confuses the boundaries and meaning behind what password expiry is for, and what password expiry warnings are intended to achieve. Similarly, some of the cost of password resets can be avoided by allowing expired passwords to be changed, rather than going through a reset procedure.

Considering our findings on password resets and voluntary password changes, *a reward* of a longer password lifetime is not the same as an *optimal reward*; this opens up avenues of research to find optimally secure and workable defenses. In an ideal scenario we envision a deployment of a policy linking password expiration with password strength only if the weakest acceptable password is below the  $10^6$  guesses threshold identified by Florêncio, Herley and Van Oorschot (2014a). Passwords would then expire in line with the expected online guessing resistance of the password; if a password is stronger than the online guessing threshold it should not unconditionally expire.



## Conclusions

Individuals perform a number of security tasks in their daily lives, both private and professional. These security tasks are supposed to support some primary activity, but often appear to conflict with objectives. This thesis set out to build methods and tools that measure the relationship between humans and security, because only when one is able to establish the status quo can the effect of change be quantified.

The literature review in [Chapter 2](#) describes existing measurements, but prior works often frame security (or compliance with security policy as a poor proxy) as the goal itself, and do not take individual's circumstances or the environment itself into account. Security can only work if it takes peoples' capabilities and limitations into account. The different perceptions of security policy designers, security awareness professionals and individuals as to what constitutes reasonable security lead to the dysfunctional eco system we live in.

In this thesis I have studied the interactions between individuals and security. It is separated into three distinct parts: 1) measuring the understanding of security; 2) measuring security in situ; 3) measuring the response to a security intervention.

My research includes the study of effective communication of security aspects in an end-user friendly manner ([Section 5.1](#)) and in an end-user unfriendly manner ([Section 4.5](#)). Essentially, I found that security communication is ignored at best, and misleading and non-workable in real life at worst ([Section 4.4](#)). Even when we deployed a tested methodology to rigorously construct new usable security advice, the performance was still poor.

The reverse communication channel, security professionals' view on communicating security to employees, was also studied ([Section 5.2](#)). The current approach

to security awareness appears to be *just do something*, without any measurement of impact. At least professionals do voice a strong demand for being able to evaluate the impact of their interventions.

To do so, the second part of my thesis focusses on measuring the security culture of the organisation. We find that context-driven scenario surveys have the capacity to differentiate between security cultures in the organisation. This enables security awareness and behaviour change professionals to tailor their interventions to the circumstances of the organisations, and to evaluate the impact by re-measuring over time. Yet in the two companies we collaborated with, we were unable to observe the impact of an intervention.

However, a change in password policy at UCL did allow me to study a change in security culture. Through the log data of the password change system I analysed the responses of all users to the incentive of longer password lifetime in return for stronger passwords over a period of 18 months. We found that the pain of having to change one's password frequently is greater than having to live with a complex password. The response took 3 months to materialise, and has yet to settle: even after 18 months, users have yet to reach the equilibrium point between the two inconveniences. The increase in mean password strength does not actually increase security at UCL as the weakest permissible password should already be sufficiently strong to protect the organisation.

There is cognitive, as well as physical, effort associated with each security task. The demands of a task may be excessive or perceived as ill-fitting, promoting the development of *coping strategies* (Sasse, Brostoff and Weirich 2001). This avoidance of effort may in fact be rational when limited personal gains are attached to the security task as perceived by the individual (Herley 2009): the research on bank's terms and conditions, the scenarios of the Productive Security research, and the password policy research all expose classic examples where users deploy coping strategies to game the system.

In organisations, security mechanisms and policies are provided to inform secure behaviours. A person might expend security effort for the benefit of others around them and the organisation as a whole (Beautement, Sasse and Wonham 2008). Cumulative effort from burdensome or repetitive security tasks can push an individual past a point where it is then harder for the organisation to encourage a return to good security behaviours: the password project observed that individuals with more than one password reset tend to choose significantly weaker passwords. *Fatalists* in Company B were very security mature (Section 6.9), but resigned to the inadequacy of security procedures and provisions.

Aside from the immediate cognitive and physical demands of security, the con-

---

sequences of committing effort to both successful and failed security activities can promote avoidance of those same associated personal costs. These costs can include potential embarrassment in the presence of others (should failure occur), and missed opportunities if the security effort associated with a task is excessive

I set out to investigate what constitutes *reasonable security*. This thesis discussed factors acting on security from various angles that need to be compromised: organisational factors such as business processes and business interests, systemic factors through regulations, as well as individual factors such as ability, conditions of actual use and usefulness. This creates conflict: regulation and policies may specify requirements that have a negative impact on business processes, yet may improve usefulness to individuals, and vice versa. The two types of policies studied in this thesis both specified what users are allowed to do. However, as the content of bank's T&Cs were inapplicable users ignored them or deployed workarounds, UCL's password policy at least engaged users and allowed them to choose a level of security that they perceived as acceptable.

Automation of security effort in technology may seem like a virtuous solution, taking the burden away from the individual. However, if the automation does not, for instance, support rule-based performance, the completion of trained-for problems is made more difficult; the individual may assess a situation as requiring input on their part, but the act of putting cognition into action is removed from them. Conscious effort will not change the result. We have observed this as part of the Security Metaphors work in [Section 5.1](#). The security automation of end-to-end encryption protocols has advanced so far that users mental models are incompatible with reality and cannot easily be improved. The properties of security systems have to remain understandable, otherwise one is unable to take grounded decisions.

Security is wholly dependent on culture. For banks, this culture is societal, as virtually everyone is a bank's customer. Similarly, communication apps are supposed to be used by everyone. In the Productive Security work, the culture is dependent on the employers of Companies A & B, and a large number of people interacting with a university such as UCL. Security should consider whether the policies and procedures are within the individual's abilities and are acceptable given the circumstances of each and every single person in these groups.

For the productive security work and the password policy study at UCL, we were able to analyse the security culture for specific demographics (age, department, relationship with the organisation, etc). This allows us to understand if specific groups are struggling, and identify how to improve the system to make it suitable for everyone.

The economical argument is also interesting. When a policy has an impact on the

processes of a large number of individuals, it ought to be cheaper and more reliable to change the system than the users. In the payment card case, it appears that banks have been unable to agree on a standard set of rules. The banks define the eco system and have full knowledge over behaviours and fraudulent actors: surely they are better placed to mitigate the risk of fraud than their customers?

The same argument holds true for UCL: password expiration has been thoroughly debunked. Both the National Institute of Standards and Technology (NIST, Grassi, Garcia and Fenton 2017) and the National Cyber Security Centre (NCSC) (2016a) now prescribe that passwords should not expire unless there is evidence of compromise. Again, the system owners are much better placed to defend the system than the users: they can monitor every single login attempt and have clear visibility of ongoing attacks.

While Whitten and Tygar focus on the the design principles of the user facing application, Adams and Sasse focus on the demands on human ability that are required by the passwords (Adams and Sasse 1999). Up to this point, users were seen as a barrier to security, and human limitations were to be blamed for security failures. Instead of focusing on the limitations, successful security solutions need to be designed for human capabilities instead, and need to be compatible with organisational and work procedures. If organisations optimised their security policies from a altruistic security perspective, these policies would never be considered. Yet it could be that the security eco-system is working optimally, but individual's time and effort are just an *acceptable collateral*.

## 8.1 Recommendations for researchers

There are no well-established ways for measuring security in organisations. There are many existing constructs that attempt to capture some aspect of security (primarily on compliance), but often their validity and applicability should be questioned (Chapter 3). This thesis discusses some additional measuring tools and argues why they are advantageous; but regardless of one's choice of measurement this thesis calls for future research to take validation seriously and respect the importance of the context when studying behaviours.

The central methodological finding of this thesis is the usefulness of open-ended free-text responses: rather than relying on closed or multiple choice questions, open-ended free text responses allow participants to qualify their understanding and supply opportunity for post-hoc validation to researchers. I have used them in the banking T&C research, where free text responses allow unbiased responses (Section 4.5.2), and in the Productive Security work (Chapter 6).

As we have seen in [Chapter 3](#), most existing survey studies use value-monoistic constructs with Likert scale responses. It appears that this methodology is too limited in capturing the interdependencies of security behaviours. Scenario-based surveys offer more space to explore the multi-dimensional conflicts, and personal preferences are more respected when the survey does not offer *obviously correct* answer options. Alternatives that future research could explore here are Q-methodology based surveys as well as paired comparisons that force participants to rank their responses.

In [Chapter 7](#), we measured the impact of a password policy change on a large set of users. After 14 months, the average password strength had increased considerably, but the conclusions gained from the research could have been different had we not had access to continuous metrics. Indeed, after 5 months users' mean password strength had decreased (see [Figure 7.14](#)). This highlights the importance of measuring the impact of an intervention multiple times, over significant time frames.

## 8.2 Recommendations for practitioners

Part of the motivation for this thesis is the demand for more robust measurement tools from information security awareness practitioners. [Chapters 6](#) and [7](#) offer two suitable avenues for measuring security compliance from the context of the user. The holistic approach described should re-focus information security awareness campaigns by responding to the actual issues that user face, and support them in their primary business tasks. Rather than seeking a 'global' measure of effectiveness of an awareness campaign, data collection surrounding the affected systems (in this case, password change logs) provides evidence of user responses.

Security policies should be consistent, applicable and realistic to their relevant context of use. The varying context of use and existing security cultures may require policies to become dynamic and flexible to support all users. This requires listening to users and involving them in security decision making.

## 8.3 Limitations

The main limitation of this thesis stems from the difficulty of collaborating with an organisation over a long period of time, allowing one to not only study the culture, but also orchestrate, observe and measure an intervention all in the same organisation. We studied the security culture in Companies A & B and successfully reported the findings back to the companies, but were unable to sustain the collaboration through organisational changes. Often our contact with an organisation hinges on

one individual. The collaboration at UCL started with observing the change in policy. Since then I have been ‘handed over’ twice now; and the main engineer and architect of the system have moved on. While I hope there is the will to improve the current policy, it will require building new working relationships.

The limitations of the individual studies have been discussed in their respective chapters. From a methodological perspective, the main draw back of my studies is the ad-hoc collection of often self-reported data. Obviously, the benefit of working with real organisation is the opportunity to obtain in situ data, which has many advantages, particularly when related to security and privacy issues. Studying actual people and organisations often comes at the cost of proper sampling and rigorous validations of methods. In this thesis I acknowledge these shortcomings and design additional post hoc validations steps where possible.

## 8.4 Future work

The individual chapters of this thesis consider policy, policy understanding, actual behaviour and behaviour change in the context of an intervention. We have learned that the security culture of an organisation is diverse, and that improving the overall security posture requires a concerted effort to properly re-examine primary tasks and their security requirements. Organisations need to value the time of individuals more highly; in an environment of existing security fatigue where even knowledgeable individuals are resigned to ineffectual security (as we found in Companies A & B), *more* is not an option. A holistic approach to security is required. Future work could start by rethinking organisational security in a bottom-up manner from the existing business processes and contrast the resulting design to existing practices.

The individual studies all have potential follow-ups. In the European Union, Payment Service Directive II is being implemented, and banks have recently updated their T&Cs to integrate the demand for open APIs. Strong authentication is now a requirement, and there have been changes to liability as well, giving rise to ample research questions.

The original Productive Security work (Section 6.3) has had a large number of follow up works (my work in Sections 6.8 and 6.9, and a number of publications by Simon Parkin, Iacovos Kirlappos, Adam Beutement and Odette Beris). From a methodological perspective, future work should integrate the validation described in Section 6.9 immediately into the survey as well as contrast the measurements with some of the measures discussed in Chapter 3.

Future research should embrace people and organisations and study security in the context that matters: the real world.



# Bibliography

- Abu-Salma, Ruba, Kat Krol, Simon Parkin, Victoria Koh, Kevin Kwan, Jazib Mahboob, Zahra Traboulsi and M. Angela Sasse. 2017. *'The Security Blanket of the Chat World: An Analytic Evaluation and a User Study of Telegram'*. In *EuroUSEC*. Paris, France: Internet Society. ISBN: 1-891562-48-7. doi:[10.14722/eurousec.2017.23006](https://doi.org/10.14722/eurousec.2017.23006).
- Abu-Salma, Ruba, M. Angela Sasse, Joseph Bonneau, Anastasia Danilova, Alena Naiakshina and Matthew Smith. 2017. *'Obstacles to the Adoption of Secure Communication Tools'*. In *IEEE Symposium on Security and Privacy*. San Jose, California, US: IEEE Computer Society. doi:[10.1109/SP.2017.65](https://doi.org/10.1109/SP.2017.65).
- Adams, Anne, and Martina Angela Sasse. 1999. *'Users Are Not the Enemy'*. *Communications of the ACM*, Communications of the ACM, 42 (12): 40–46. doi:[10.1145/322796.322806](https://doi.org/10.1145/322796.322806).
- Adams, John. 2003. *'Risk and Morality: Three Framing Devices'*. *Risk and morality*: 87–106.
- Albrechtsen, Eirik, and Jan Hovden. 2009. *'The Information Security Digital Divide between Information Security Managers and Users'*. *Computers & Security* 28 (6): 476–490. doi:[10.1016/j.cose.2009.01.003](https://doi.org/10.1016/j.cose.2009.01.003).
- Alistratov, Oleg. 2010. *Data::Password::Entropy*. Accessed 8 February 2018. <http://search.cpan.org/~zero/Data-Password-Entropy-0.08/>.
- Almeshekah, Mohammed H., Christopher N. Gutierrez, Mikhail J. Atallah and Eugene H. Spafford. 2015. *'ErsatzPasswords: Ending Password Cracking and Detecting Password Leakage'*. In *Proceedings of the 31st Annual Computer Security Applications Conference*, 311–320. ACSAC 2015. New York, NY, USA: ACM. ISBN: 978-1-4503-3682-6. doi:[10.1145/2818000.2818015](https://doi.org/10.1145/2818000.2818015).
- Alsaleh, Mansour, Mohammad Mannan and Paul C. Van Oorschot. 2012. *'Revisiting Defenses against Large-Scale Online Password Guessing Attacks'*. *IEEE Transactions on Dependable and Secure Computing* 9 (1): 128–141. ISSN: 1545-5971. doi:[10.1109/TDSC.2011.24](https://doi.org/10.1109/TDSC.2011.24).
- Alty, James L., Roger P. Knott, Ben Anderson and Michael Smyth. 2000. *'A Framework for Engineering Metaphor at the User Interface'*. *Interacting with Computers* 13 (2): 301–322. ISSN: 0953-5438. doi:[10.1016/S0953-5438\(00\)00047-3](https://doi.org/10.1016/S0953-5438(00)00047-3).

- Anderson, Michael C., and James H. Neely. 1996. 'Interference and Inhibition in Memory Retrieval'. In *Memory. Handbook of Perception and Cognition*, 2nd ed., 237–313. Academic Press.
- Anderson, Ross. 2007. 'Closing the Phishing Hole—Fraud, Risk and Nonbanks'. In *Federal Reserve Bank of Kansas City, Conference on Nonbanks in the Payments System*. Accessed 27 April 2017. <https://www.cl.cam.ac.uk/~rja14/Papers/nonbanks.pdf>.
- Anderson, Ross, Mike Bond and Steven J. Murdoch. 2006. 'Chip and Spin'. *Computer Security Journal* 22 (2): 1–6. Accessed 27 April 2017. <https://pdfs.semanticscholar.org/69e1/cee578535bd77e0b261e24f639a70b625eb3.pdf>.
- Anderson, Ross, and Tyler Moore. 2006. 'The Economics of Information Security'. *Science* 314 (5799): 610–613. doi:10.1126/science.1130992.
- Ashenden, Debi, and Darren Lawrence. 2013. 'Can We Sell Security like Soap?: A New Approach to Behaviour Change'. In *Proceedings of the 2013 Workshop on New Security Paradigms Workshop*, 87–94. ACM.
- . 2016. 'Security Dialogues: Building Better Relationships between Security and Business'. *IEEE Security & Privacy* 14 (3): 82–87.
- Bai, Wei, Moses Namara, Yichen Qian, Patrick Gage Kelley, Michelle L Mazurek and Doowon Kim. 2016. 'An Inconvenient Trust: User Attitudes toward Security and Usability Tradeoffs for Key-Directory Encryption Systems'. In *SOUPS*. ACM.
- Barclays PLC. 2002. *Response to Consultation on a Possible Legal Framework for the Single Payment Area in the Internal Market*. Accessed 6 May 2017. [http://ec.europa.eu/internal\\_market/payments/docs/framework/framework-workingdoc-contrib/barclays\\_en.pdf](http://ec.europa.eu/internal_market/payments/docs/framework/framework-workingdoc-contrib/barclays_en.pdf).
- Beautement, Adam, Ingolf Becker, Simon Parkin, Kat Krol and M. Angela Sasse. 2016. 'Productive Security: A Scalable Methodology for Analysing Employee Security Behaviours'. In *Twelfth Symposium on Usable Privacy and Security (SOUPS 2016)*. Denver, CO: USENIX Association. ISBN: 978-1-931971-31-7. <https://www.usenix.org/conference/soups2016/technical-sessions/presentation/beautement>.
- Beautement, Adam, Robert Coles, Jonathan Griffin, Christos Andronis, Brian Monahan, David Pym, Martina Angela Sasse and Mike Wonham. 2009. 'Modelling the Human and Technological Costs and Benefits of USB Memory Stick Security'. *Managing Information Risk and the Economics of Security*: 141–163. doi:10.1007/978-0-387-09762-6\_7.
- Beautement, Adam, Martina Angela Sasse and Mike Wonham. 2008. 'The Compliance Budget: Managing Security Behaviour in Organisations'. In *Proceedings of the 2008 Workshop on New Security Paradigms*, 47–58. NSPW '08. ISBN: 978-1-60558-341-9. doi:10.1145/1595676.1595684.

- Becker, Ingolf, Alice Hutchings, Ruba Abu-Salma, Ross Anderson, Nicolas Bohm, Steven J. Murdoch, M. Angela Sasse and Gianluca Stringhini. 2016. 'International Comparison of Bank Fraud Reimbursement: Customer Perceptions and Contractual Terms'. In *Workshop on the Economics of Information Security (WEIS '16)*. Berkeley, US. [http://weis2016.econinfosec.org/wp-content/uploads/sites/2/2016/05/WEIS\\_2016\\_paper\\_78-1.pdf](http://weis2016.econinfosec.org/wp-content/uploads/sites/2/2016/05/WEIS_2016_paper_78-1.pdf).
- . 2017. 'International Comparison of Bank Fraud Reimbursement: Customer Perceptions and Contractual Terms'. *Journal of Cybersecurity* 3 (2): 109–125. doi:10.1093/cybsec/tyx011.
- Becker, Ingolf, Simon Parkin and M. Angela Sasse. 2017a. 'Finding Security Champions in Blends of Organisational Culture'. In *Proc. EuroUSEC '17*, 11. Paris, France: Internet Society. ISBN: 1-891562-48-7. [https://www.internetsociety.org/sites/default/files/eurousec2017\\_07\\_Becker\\_paper.pdf](https://www.internetsociety.org/sites/default/files/eurousec2017_07_Becker_paper.pdf).
- . 2017b. 'Measuring the Success of Context-Aware Security Behaviour Surveys'. In *Learning from Authoritative Security Experiment Results (LASER)*. Arlington, VA: USENIX Association. [https://www.usenix.org/system/files/conference/laser2017/laser2017\\_becker.pdf](https://www.usenix.org/system/files/conference/laser2017/laser2017_becker.pdf).
- . 2018. 'Rewarding Users for Stronger Passwords: Linking Password Lifetime to Strength'. In *USENIX Security Symposium*. Baltimore, Maryland, USA: USENIX Association. <https://www.usenix.org/conference/usenixsecurity18/presentation/becker>.
- Bergeron, François, Louis Raymond, Suzanne Rivard and Marie-France Gara. 1995. 'Determinants of EIS Use: Testing a Behavioral Model'. *Decision Support Systems* 14 (2): 131–146. ISSN: 0167-9236. doi:10.1016/0167-9236(94)00007-F.
- Beris, Odette, Adam Beautement and M. Angela Sasse. 2015. 'Employee Rule Breakers, Excuse Makers and Security Champions: Mapping the Risk Perceptions and Emotions That Drive Security Behaviors'. In *New Security Paradigms. NSPW '15*. Twente, Netherlands: ACM. doi:10.1145/2841113.2841119.
- Blythe, John. 2015. 'Information Security in the Workplace: A Mixed-Methods Approach to Understanding and Improving Security Behaviours'. PhD Thesis, Northumbria University. [http://nrl.northumbria.ac.uk/30328/1/blythe.john\\_phd.pdf](http://nrl.northumbria.ac.uk/30328/1/blythe.john_phd.pdf).
- Blythe, John M., Lynne M. Coventry and Linda Little. 2015. 'Unpacking Security Policy Compliance: The Motivators and Barriers of Employees' Security Behaviors'. In *Eleventh Symposium On Usable Privacy and Security (SOUPS 2015)*, 103–122. Ottawa: USENIX Association. ISBN: 978-1-931971-24-9, accessed 14 July 2017. <https://www.usenix.org/conference/soups2015/proceedings/presentation/blythe>.
- Bohm, Nicholas, Ian Brown and Brian Gladman. 2000. 'Electronic Commerce: Who Carries the Risk of Fraud'. *Journal of Information Law and Technology* 3:00–3. Accessed 27 April 2017. <https://pdfs.semanticscholar.org/e653/bf1e51e2ba870261b324a9d1f7441de5b2fd.pdf>.

- Böhme, Rainer, and Stefan Köpsell. 2010. *'Trained to Accept?: A Field Experiment on Consent Dialogs'*. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 2403–2406. CHI '10. New York, NY, USA: ACM. ISBN: 978-1-60558-929-9. doi:[10.1145/1753326.1753689](https://doi.org/10.1145/1753326.1753689).
- Bonneau, Joseph. 2012. *'The Science of Guessing: Analyzing an Anonymized Corpus of 70 Million Passwords'*. In *Proceedings of the 2012 IEEE Symposium on Security and Privacy*, 538–552. SP '12. Washington, DC, USA: IEEE Computer Society. ISBN: 978-0-7695-4681-0. doi:[10.1109/SP.2012.49](https://doi.org/10.1109/SP.2012.49).
- Bonneau, Joseph, Cormac Herley, Paul C. van Oorschot and Frank Stajano. 2012. *'The Quest to Replace Passwords: A Framework for Comparative Evaluation of Web Authentication Schemes'*, 553–567. IEEE. ISBN: 978-1-4673-1244-8 978-0-7695-4681-0. doi:[10.1109/SP.2012.44](https://doi.org/10.1109/SP.2012.44).
- Bonneau, Joseph, Sören Preibusch and Ross Anderson. 2012. *'A Birthday Present Every Eleven Wallets? The Security of Customer-Chosen Banking PINs'*. In *Financial Cryptography and Data Security*, 25–40. LNCS. Springer. ISBN: 978-3-642-32945-6. doi:[10.1007/978-3-642-32946-3\\_3](https://doi.org/10.1007/978-3-642-32946-3_3).
- Boudreau, Marie-Claude, David Gefen and Detmar W. Straub. 2001. *'Validation in Information Systems Research: A State-of-the-Art Assessment'*. *MIS Quarterly* 25 (1): 1–16. ISSN: 0276-7783. doi:[10.2307/3250956](https://doi.org/10.2307/3250956). JSTOR: [3250956](https://www.jstor.org/stable/3250956).
- Bowerman, Mark. 2007. *Radio Interview with APACS Spokesperson*. BBC Radio Merseyside.
- Brand, Sheila, and J Makey. 1985. *Department of Defense Password Management Guideline CSC-STD-002-85*. Department of Defense Computer Security Center.
- Brostoff, Sacha. 2005. *'Improving Password System Effectiveness'*. Doctoral Thesis, University of London. Accessed 1 November 2017. <http://discovery.ucl.ac.uk/1445330/>.
- Brostoff, Sacha, and Martina Angela Sasse. 2003. *'"Ten Strikes and You're out": Increasing the Number of Login Attempts Can Improve Password Usability'*. In *Proc. CHI Workshop on HCI and Security Systems*. Ft. Lauderdale, FL, USA. Accessed 23 June 2015. <http://discovery.ucl.ac.uk/19826/>.
- Bulgurcu, Burcu, Hasan Cavusoglu and Izak Benbasat. 2010. *'Information Security Policy Compliance: An Empirical Study of Rationality-Based Beliefs and Information Security Awareness'*. *MIS quarterly* 34 (3): 523–548. Accessed 26 January 2016. <https://dl.acm.org/citation.cfm?id=2017477>.
- Burkell, Jacquelyn A., Alexandre Fortier, Lorraine Wong and Jennifer Lynn Simpson. 2013. *The View From Here: User-Centered Perspectives on Social Network Privacy*, Library and Information Science Publications Paper 25. Western University. <http://ir.lib.uwo.ca/fimspub/25>.
- Burr, William E., Donna F. Dodson and W. Timothy Polk. 2004. *Electronic Authentication Guideline NIST SP 800-63v1.0.1*. Gaithersburg, MD: National Institute of Standards and Technology. doi:[10.6028/NIST.SP.800-63v1.0.1](https://doi.org/10.6028/NIST.SP.800-63v1.0.1).

- Camp, L Jean. 2009. 'Mental Models of Privacy and Security'. *IEEE Technology and society magazine* 28 (3). doi:10.1109/MTS.2009.934142.
- Chang, Hsin Hsin, and Shuang-Shii Chuang. 2011. 'Social Capital and Individual Motivations on Knowledge Sharing: Participant Involvement as a Moderator'. *Information & Management* 48 (1): 9–18. ISSN: 0378-7206. doi:10.1016/j.im.2010.11.001.
- Chapple, Constance L., Julia A. McQuillan and Terceira A. Berdahl. 2005. 'Gender, Social Bonds, and Delinquency: A Comparison of Boys' and Girls' Models'. *Social Science Research* 34 (2): 357–383. ISSN: 0049-089X. doi:10.1016/j.ssresearch.2004.04.003.
- Charoen, Danuvasin, Murali Raman and Lorne Olfman. 2008. 'Improving End User Behaviour in Password Utilization: An Action Research Initiative'. *Systemic Practice and Action Research* 21 (1): 55–72. ISSN: 1094-429X, 1573-9295. doi:10.1007/s11213-007-9082-4.
- Cheney, Julia, Robert Hunt, Vyacheslav Mikhed, Dubravka Ritter and Michael Vogan. 2014. *Consumer Use of Fraud Alerts and Credit Freezes: An Empirical Analysis*. Discussion Paper, Payment Cards Center. Philadelphia, PA: Federal Reserve Bank of Philadelphia.
- Cheng, Lijiao, Ying Li, Wenli Li, Eric Holm and Qingguo Zhai. 2013. 'Understanding the Violation of IS Security Policy in Organizations: An Integrated Model Based on Social Control and Deterrence Theory'. *Computers & Security* 39:447–459. ISSN: 0167-4048. doi:10.1016/j.cose.2013.09.009.
- Cheswick, William. 2013. 'Rethinking Passwords'. *Commun. ACM* 56 (2): 40–44. ISSN: 0001-0782. doi:10.1145/2408776.2408790.
- Cheung, Waiman, Man Kit Chang and Vincent S Lai. 2000. 'Prediction of Internet and World Wide Web Usage at Work: A Test of an Extended Triandis Model'. *Decision Support Systems* 30 (1): 83–100. ISSN: 0167-9236. doi:10.1016/S0167-9236(00)00125-1.
- Chiasson, Sonia, and P. C. van Oorschot. 2015. 'Quantifying the Security Advantage of Password Expiration Policies'. *Designs, Codes and Cryptography* 77 (2-3): 401–408. ISSN: 0925-1022, 1573-7586. doi:10.1007/s10623-015-0071-9.
- Choong, Yee-Yin, and Mary Theofanos. 2015. 'What 4,500+ People Can Tell You – Employees' Attitudes Toward Organizational Password Policy Do Matter'. In *Human Aspects of Information Security, Privacy, and Trust*, 299–310. Lecture Notes in Computer Science. Springer, Cham. ISBN: 978-3-319-20375-1 978-3-319-20376-8. doi:10.1007/978-3-319-20376-8\_27.
- Clark, Louise, and M. Angela Sasse. 1997. 'Conceptual Design Reconsidered: The Case of the Internet Session Directory Tool'. In *People and Computers XII: Proceedings of HCI*, 97:67–84. Springer. doi:10.1007/978-1-4471-3601-9\_5.

- Compeau, Deborah R., and Christopher A. Higgins. 1995. 'Computer Self-Efficacy: Development of a Measure and Initial Test'. *MIS Quarterly* 19 (2): 189–211. ISSN: 0276-7783. doi:10.2307/249688. JSTOR: 249688.
- Connolly, Lena, Michael Lang and J Doug Tygar. 2015. 'Investigation of Employee Security Behaviour: A Grounded Theory Approach'. In *IFIP International Information Security Conference*, 455:283–296. IFIPAICT. Springer. doi:10.1007/978-3-319-18467-8\_19.
- Cranor, Lorrie Faith, and Simson Garfinkel. 2005. *Security and Usability: Designing Secure Systems That People Can Use*. O'Reilly. ISBN: 0-596-00827-9.
- Cronbach, Lee J. 1949. *Essentials of Psychological Testing*. Essentials of Psychological Testing. Oxford, England: Harper.
- Cronbach, Lee J., and Paul E. Meehl. 1955. 'Construct Validity in Psychological Tests.' *Psychological bulletin* 52 (4): 281.
- Crossler, Robert E, Allen C Johnston, Paul Benjamin Lowry, Qing Hu, Merrill Warkentin and Richard Baskerville. 2013. 'Future Directions for Behavioral Information Security Research'. *Computers & Security* 32:90–101. doi:10.1016/j.cose.2012.09.010.
- D'Arcy, John, Tejaswini Herath and Mindy K Shoss. 2014. 'Understanding Employee Responses to Stressful Information Security Requirements: A Coping Perspective'. *Journal of Management Information Systems* 31 (2): 285–318. doi:10.2753/MIS0742-1222310210.
- D'Arcy, John, Anat Hovav and Dennis Galletta. 2009. 'User Awareness of Security Countermeasures and Its Impact on Information Systems Misuse: A Deterrence Approach'. *Information Systems Research* 20 (1): 79–98. ISSN: 1047-7047. doi:10.1287/isre.1070.0160.
- Davison, Anthony Christopher, and David Victor Hinkley. 1997. *Bootstrap Methods and Their Application*. Vol. 1. Cambridge university press. ISBN: 978-0-511-80284-3. doi:10.1017/CB09780511802843.
- De Carnavalet, Xavier de Carné, and Mohammad Mannan. 2014. 'From Very Weak to Very Strong: Analyzing Password-Strength Meters.' In *NDSS*, 14:23–26. [http://www.ndss-symposium.org/wp-content/uploads/sites/25/2017/09/06\\_3\\_1.pdf](http://www.ndss-symposium.org/wp-content/uploads/sites/25/2017/09/06_3_1.pdf).
- Dekker, Sidney. 2014. *The Field Guide to Understanding 'Human Error'*. Ashgate Publishing, Ltd.
- Dell'Amico, M., P. Michiardi and Y. Roudier. 2010. 'Password Strength: An Empirical Analysis'. In *2010 Proceedings IEEE INFOCOM*, 1–9. doi:10.1109/INFOCOM.2010.5461951.

- Demjaha, Albesë, Jonathan M. Spring, Ingolf Becker, Simon Parkin and M. Angela Sasse. 2018. ‘*Metaphors Considered Harmful? An Exploratory Study of the Effectiveness of Functional Metaphors for End-to-End Encryption*’. In *Proc. USEC*. San Diego, CA: Internet Society. doi:[10.14722/usec.2018.23015](https://doi.org/10.14722/usec.2018.23015).
- Dhillon, Gurpreet, and James Backhouse. 2001. ‘*Current Directions in IS Security Research: Towards Socio-organizational Perspectives*’. *Information Systems Journal* 11 (2): 127–153. ISSN: 1365-2575. doi:[10.1046/j.1365-2575.2001.00099.x](https://doi.org/10.1046/j.1365-2575.2001.00099.x).
- diSessa, Andrea. 1986. ‘*Models of Computation*’. *User Centered System Design: New Perspectives on Human Computer Interaction*. Hillsdale, NJ: Lawrence Erlbaum.
- Dodier-Lazaro, Steve, Ruba Abu-Salma, Ingolf Becker and M. Angela Sasse. 2017. ‘*From Paternalistic to User-Centred Security: Putting Users First with Value-Sensitive Design*’. In *CHI 2017 Workshop on Values in Computing*. ACM. [http://www.valuesincomputing.org/wp-content/uploads/2017/03/dodier\\_paternalistic\\_vic2017.pdf](http://www.valuesincomputing.org/wp-content/uploads/2017/03/dodier_paternalistic_vic2017.pdf).
- Dodier-Lazaro, Steve, Ingolf Becker, Jens Krinke and M. Angela Sasse. 2017. ‘*No Good Reason to Remove Features: Expert Users Value Useful Apps over Secure Ones*’. In *International Conference on Human Aspects of Information Security, Privacy, and Trust*, 25–44. Springer, Cham. doi:[10.1007/978-3-319-58460-7\\_3](https://doi.org/10.1007/978-3-319-58460-7_3).
- Drucker, Peter F. 2004. ‘*What Makes an Effective Executive*’. *Harvard Business Review* 82 (6).
- Egelman, Serge, Marian Harbach and Eyal Peer. 2016. ‘*Behavior Ever Follows Intention?: A Validation of the Security Behavior Intentions Scale (SeBIS)*’. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, 5257–5261. ACM. doi:[10.1145/2858036.2858265](https://doi.org/10.1145/2858036.2858265).
- Fahl, Sascha, Marian Harbach, Yasemin Acar and Matthew Smith. 2013. ‘*On the Ecological Validity of a Password Study*’. In *Proceedings of the Ninth Symposium on Usable Privacy and Security*, 13:1–13:13. SOUPS ’13. New York, NY, USA: ACM. ISBN: 978-1-4503-2319-2. doi:[10.1145/2501604.2501617](https://doi.org/10.1145/2501604.2501617).
- Fahl, Sascha, Marian Harbach, Thomas Muders, Matthew Smith and Uwe Sander. 2012. ‘*Helping Johnny 2.0 to Encrypt His Facebook Conversations*’. In *SOUPS*. doi:[10.1145/2335356.2335371](https://doi.org/10.1145/2335356.2335371).
- Feledi, Daniel, and Stefan Fenz. 2012. ‘*Challenges of Web-Based Information Security Knowledge Sharing*’. In *2012 Seventh International Conference on Availability, Reliability and Security*, 514–521. doi:[10.1109/ARES.2012.59](https://doi.org/10.1109/ARES.2012.59).
- Fereday, Jennifer, and Eimear Muir-Cochrane. 2006. ‘*Demonstrating Rigor Using Thematic Analysis: A Hybrid Approach of Inductive and Deductive Coding and Theme Development*’. *International Journal of Qualitative Methods* 5 (1): 80–92. ISSN: 1609-4069. doi:[10.1177/160940690600500107](https://doi.org/10.1177/160940690600500107).

- Financial Fraud Action UK. 2017. *Year-End 2016 Fraud Update: Payment Cards, Remote Banking and Cheque*. Bulletin. London, UK.
- Financial Ombudsman Service. 2014a. *Case 116/02*. Ombudsman News March/April 2014. <http://www.financial-ombudsman.org.uk/publications/ombudsman-news/%20116/116-disputed-transactions.html>.
- . 2014b. *Case 116/08*. Ombudsman News March/April 2014. <http://www.financial-ombudsman.org.uk/publications/ombudsman-news/%20116/116-disputed-transactions.html>.
- . 2014c. *Case 116/09*. Ombudsman News March/April 2014. <http://www.financial-ombudsman.org.uk/publications/ombudsman-news/%20116/116-disputed-transactions.html>.
- Florêncio, Dinei, and Cormac Herley. 2010. ‘Where Do Security Policies Come From?’ In *Proceedings of the Sixth Symposium on Usable Privacy and Security*, 10:1–10:14. SOUPS ’10. New York, NY, USA: ACM. ISBN: 978-1-4503-0264-7. doi:10.1145/1837110.1837124.
- Florêncio, Dinei, Cormac Herley and Baris Coskun. 2007. ‘Do Strong Web Passwords Accomplish Anything?’ In *Proceedings of the 2Nd USENIX Workshop on Hot Topics in Security*, 10:1–10:6. HOTSEC’07. Berkeley, CA, USA: USENIX Association. Accessed 5 October 2017. <http://dl.acm.org/citation.cfm?id=1361419.1361429>.
- Florêncio, Dinei, Cormac Herley and Paul C. Van Oorschot. 2014a. ‘An Administrator’s Guide to Internet Password Research’. In *Proc. USENIX LISA*. Accessed 8 October 2014. <http://people.scs.carleton.ca/~paulv/papers/WhatsaSysadminToDo.pdf>.
- . 2014b. ‘Password Portfolios and the Finite-Effort User: Sustainably Managing Large Numbers of Accounts’. In *Proc. USENIX Security*, 575–590. San Diego, CA: USENIX Association. ISBN: 978-1-931971-15-7, accessed 14 September 2015. <https://www.usenix.org/conference/usenixsecurity14/technical-sessions/presentation/florenccio>.
- . 2016. ‘Pushing on String: The ‘Don’t Care’ Region of Password Strength’. *Commun. ACM* 59 (11): 66–74. ISSN: 0001-0782. doi:10.1145/2934663.
- Fowler, Floyd. 2009. *Survey Research Methods*. 4th Edition. 2455 Teller Road, Thousand Oaks California 91320 United States: SAGE Publications, Inc. ISBN: 978-1-4129-5841-7 978-1-4522-3018-4. doi:10.4135/9781452230184.
- Furnell, Steven, and Anish Rajendran. 2012. ‘Understanding the Influences on Information Security Behaviour’. *Computer Fraud & Security* 2012 (3): 12–15.
- Gabriel, Trevor, and Steven Furnell. 2011. ‘Selecting Security Champions’. *Computer Fraud & Security* 2011 (8): 8–12. ISSN: 1361-3723. doi:10.1016/S1361-3723(11)70082-3.



- Gaw, Shirley, Edward W. Felten and Patricia Fernandez-Kelly. 2006. 'Secrecy, Flagging, and Paranoia: Adoption Criteria in Encrypted Email'. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 591–600. CHI '06. New York, NY, USA: ACM. ISBN: 978-1-59593-372-0. doi:10.1145/1124772.1124862.
- Good, Nathaniel S., Jens Grossklags, Deirdre K. Mulligan and Joseph A. Konstan. 2007. 'Noticing Notice: A Large-Scale Experiment on the Timing of Software License Agreements'. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 607–616. CHI '07. New York, NY, USA: ACM. ISBN: 978-1-59593-593-9. doi:10.1145/1240624.1240720.
- Grassi, Paul A, Michael E Garcia and James L Fenton. 2017. *Digital Identity Guidelines: Revision 3* NIST SP 800-63-3. Gaithersburg, MD: National Institute of Standards and Technology. doi:10.6028/NIST.SP.800-63-3.
- Guo, Ken H., Yufei Yuan, Norman P. Archer and Catherine E. Connelly. 2011. 'Understanding Nonmalicious Security Violations in the Workplace: A Composite Behavior Model'. *Journal of Management Information Systems* 28 (2): 203–236. ISSN: 0742-1222. doi:10.2753/MIS0742-1222280208.
- Gwet, Kilem L. 2014. *Handbook of Inter-Rater Reliability: The Definitive Guide to Measuring the Extent of Agreement among Raters*. Advanced Analytics, LLC. Accessed 11 July 2017. <http://www.agreestat.com/book4/>.
- Hache, Ana Carolina Blanco, and Nicholas Ryder. 2011. 'Tis the Season to (Be Jolly?) Wise-up to Online Fraudsters. Criminals on the Web Lurking to Scam Shoppers This Christmas:1 a Critical Analysis of the United Kingdom's Legislative Provisions and Policies to Tackle Online Fraud'. *Information & Communications Technology Law* 20 (1): 35–56. ISSN: 1360-0834. doi:10.1080/13600834.2011.557537.
- Haidt, Jonathan. 2012. *The Righteous Mind: Why Good People Are Divided by Politics and Religion*. Vintage.
- Halasz, Frank, and Thomas P. Moran. 1982. 'Analogy Considered Harmful'. In *Proceedings of the 1982 Conference on Human Factors in Computing Systems*, 383–386. CHI '82. New York, NY, USA: ACM. doi:10.1145/800049.801816.
- Heine, Steven J., Darrin R. Lehman, Kaiping Peng and Joe Greenholtz. 2002. 'What's Wrong with Cross-Cultural Comparisons of Subjective Likert Scales?: The Reference-Group Effect'. *Journal of personality and social psychology* 82 (6): 903. doi:10.1037/0022-3514.82.6.903.
- Herath, Tejaswini, and H. Raghav Rao. 2009a. 'Encouraging Information Security Behaviors in Organizations: Role of Penalties, Pressures and Perceived Effectiveness'. *Decision Support Systems* 47 (2): 154–165. ISSN: 0167-9236. doi:10.1016/j.dss.2009.02.005.

- Herath, Tejaswini, and H. Raghav Rao. 2009b. 'Protection Motivation and Deterrence: A Framework for Security Policy Compliance in Organisations'. *European Journal of Information Systems* 18 (2): 106–125. ISSN: 0960-085X, 1476-9344. doi:10.1057/ejis.2009.6.
- Herley, Cormac. 2009. 'So Long, and No Thanks for the Externalities: The Rational Rejection of Security Advice by Users'. In *Proceedings of the 2009 Workshop on New Security Paradigms Workshop*, 133–144. NSPW '09. Accessed 10 November 2012. <http://dl.acm.org/citation.cfm?id=1719050>.
- . 2014. 'More Is Not the Answer'. *IEEE Security Privacy* 12 (1): 14–19. ISSN: 1540-7993. doi:10.1109/MSP.2013.134.
- Herley, Cormac, and P. Van Oorschot. 2012. 'A Research Agenda Acknowledging the Persistence of Passwords'. *IEEE Security Privacy* 10 (1): 28–36. ISSN: 1540-7993. doi:10.1109/MSP.2011.150.
- Herzberg, Amir, and Hemi Leibowitz. 2016. 'Can Johnny Finally Encrypt? Evaluating E2E Encryption in Popular IM Applications'. In *ACM Workshop on Socio-Technical Aspects in Security and Trust (STAST)*. doi:10.1145/3046055.3046059.
- Hsu, Chin-Lung, and Judy Chuan-Chuan Lin. 2008. 'Acceptance of Blog Usage: The Roles of Technology Acceptance, Social Influence and Knowledge Sharing Motivation'. *Information & Management* 45 (1): 65–74. ISSN: 0378-7206. doi:10.1016/j.im.2007.11.001.
- Hsu, Jack Shih-Chieh, Sheng-Pao Shih, Yu Wen Hung and Paul Benjamin Lowry. 2015. 'The Role of Extra-Role Behaviors and Social Controls in Information Security Policy Effectiveness'. *Information Systems Research* 26 (2): 282–300. doi:10.1287/isre.2015.0569.
- Hsu, Meng-Hsiang, Teresa L. Ju, Chia-Hui Yen and Chun-Ming Chang. 2007. 'Knowledge Sharing Behavior in Virtual Communities: The Relationship between Trust, Self-Efficacy, and Outcome Expectations'. *International Journal of Human-Computer Studies* 65 (2): 153–169. ISSN: 1071-5819. doi:10.1016/j.ijhcs.2006.09.003.
- Huang, Chi-Cheng. 2009. 'Knowledge Sharing and Group Cohesiveness on Performance: An Empirical Study of Technology R&D Teams in Taiwan'. *Technovation* 29 (11): 786–797. ISSN: 0166-4972. doi:10.1016/j.technovation.2009.04.003.
- Hubbard, Douglas W. 2014. *How to Measure Anything: Finding the Value of Intangibles in Business*. John Wiley & Sons.
- Hui, C. Harry, and Harry C. Triandis. 1985. 'Measurement in Cross-Cultural Psychology: A Review and Comparison of Strategies'. *Journal of cross-cultural psychology* 16 (2): 131–152. doi:10.1177/0022002185016002001.

- Ifinedo, Princely. 2012. 'Understanding Information Systems Security Policy Compliance: An Integration of the Theory of Planned Behavior and the Protection Motivation Theory'. *Computers & Security* 31 (1): 83–95. ISSN: 0167-4048. doi:10.1016/j.cose.2011.10.007.
- . 2014. 'Information Systems Security Policy Compliance: An Empirical Study of the Effects of Socialisation, Influence, and Cognition'. *Information & Management* 51 (1): 69–79. ISSN: 0378-7206. doi:10.1016/j.im.2013.10.001.
- Inglesant, Philip G., and Martina Angela Sasse. 2010. 'The True Cost of Unusable Password Policies: Password Use in the Wild'. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 383–392. CHI '10. New York, NY, USA: ACM. ISBN: 978-1-60558-929-9. doi:10.1145/1753326.1753384.
- Jaeger, Lennart. 2018. 'Information Security Awareness: Literature Review and Integrative Framework'. *Hawaii International Conference on System Sciences 2018 (HICSS-51)*. [https://aisel.aisnet.org/hicss-51/os/information\\_security/2](https://aisel.aisnet.org/hicss-51/os/information_security/2).
- Jansen, Jurjen, and Rutger Leukfeldt. 2016. 'Phishing and Malware Attacks on Online Banking Customers in the Netherlands: A Qualitative Analysis of Factors Leading to Victimization'. *International Journal of Cyber Criminology; Thirunelveli* 10 (1): 79–91. doi:10.5281/zenodo.58523.
- Jeon, Su-Hwan, Young-Gul Kim and Joon Koh. 2011. 'Individual, Social, and Organizational Contexts for Active Knowledge Sharing in Communities of Practice'. *Expert Systems with Applications* 38 (10): 12423–12431. ISSN: 0957-4174. doi:10.1016/j.eswa.2011.04.023.
- Johnston, Allen C., and Merrill Warkentin. 2010. 'The Influence of Perceived Source Credibility on End User Attitudes and Intentions to Comply with Recommended IT Actions'. *Journal of Organizational and End User Computing* 22 (3): 1–21. doi:10.4018/joeuc.2010070101.
- Johnston, Allen C, Merrill Warkentin, Maranda McBride and Lemuria Carter. 2016. 'Dispositional and Situational Factors: Influences on Information Security Policy Violations'. *European Journal of Information Systems* 25 (3): 231–251. doi:10.1057/ejis.2015.15.
- Jones, Elaine G., and Margarita Kay. 1992. 'Instrumentation in Cross-Cultural Research'. *Nursing research* 41 (3): 186–188. Accessed 28 April 2017. [http://journals.lww.com/nursingresearchonline/Citation/1992/05000/Instrumentation\\_in\\_Cross\\_Cultural\\_Research.12.aspx](http://journals.lww.com/nursingresearchonline/Citation/1992/05000/Instrumentation_in_Cross_Cultural_Research.12.aspx).
- Karlsson, Fredrik, Martin Karlsson and Joachim Åström. 2017. 'Measuring Employees' Compliance – the Importance of Value Pluralism'. *Information and Computer Security* 25 (3): 279–299. ISSN: 2056-4961. doi:10.1108/ICS-11-2016-0084.

- Katz, Jonathan, and Yehuda Lindell. 2014. *Introduction to Modern Cryptography*. CRC press.
- Kelley, Patrick Gage, Saranga Komanduri, Michelle L. Mazurek, Richard Shay, Timothy Vidas, Lujo Bauer, Nicolas Christin, Lorrie Faith Cranor and Julio Lopez. 2012. ‘*Guess Again (and Again and Again): Measuring Password Strength by Simulating Password-Cracking Algorithms*’. In *2012 IEEE Symposium on Security and Privacy*, 523–537. doi:[10.1109/SP.2012.38](https://doi.org/10.1109/SP.2012.38).
- Kelman, Alistair. 2009. ‘*Job v Halifax PLC (Not Reported) Case Number 7BQ00307*’. In *Digital Evidence and Electronic Signature Law Review*, vol. 6.
- Al-Khaldi, Muhammad A., and R. S. Olusegun Wallace. 1999. ‘*The Influence of Attitudes on Personal Computer Utilization among Knowledge Workers: The Case of Saudi Arabia*’. *Information & Management* 36 (4): 185–204. ISSN: 0378-7206. doi:[10.1016/S0378-7206\(99\)00017-8](https://doi.org/10.1016/S0378-7206(99)00017-8).
- Kirlappos, Iacovos. 2016. ‘*Learning from "Shadow Security": Understanding Non-Compliant Behaviours to Improve Information Security Management*’. Doctoral, UCL (University College London). Accessed 15 July 2017. <http://discovery.ucl.ac.uk/1521997/>.
- Kirlappos, Iacovos, Adam Beutement and Martina Angela Sasse. 2013. ‘*"Comply or Die" Is Dead: Long Live Security-Aware Principal Agents*’. In *Workshop on Usable Security*. Okinawa, Japan. doi:[10.1007/978-3-642-41320-9\\_5](https://doi.org/10.1007/978-3-642-41320-9_5).
- Kirlappos, Iacovos, Simon Parkin and M. Angela Sasse. 2015. ‘*Shadow Security as a Tool for the Learning Organization*’. *ACM SIGCAS Computers and Society* 45 (1): 29–37. Accessed 5 January 2016. <http://dl.acm.org/citation.cfm?id=2738216>.
- Kirlappos, Iacovos, Simon Parkin and Martina Angela Sasse. 2014. ‘*Learning from "Shadow Security": Why Understanding Non-Compliance Provides the Basis for Effective Security*’. In *Workshop on Usable Security (USEC 2014)*. San Diego, California. doi:[10.14722/usec.2014.23007](https://doi.org/10.14722/usec.2014.23007).
- Kirlappos, Iacovos, and Martina Angela Sasse. 2015. ‘*Fixing Security Together*’. In *Proceedings of USEC 2015*. San Diego, California: Internet Society. doi:[10.14722/usec.2015.23013](https://doi.org/10.14722/usec.2015.23013).
- Krol, Kat, Simon Parkin and M. Angela Sasse. 2016. ‘*Better the Devil You Know: A User Study of Two CAPTCHAs and a Possible Replacement Technology*’. Accessed 11 January 2017. <https://www.internetsociety.org/sites/default/files/blogs-media/better-the-devil-you-know-user-study-of-two-captchas-a-possible-replacement-technology.pdf>.
- Krol, Kat, E. Philippou, Emiliano De Cristofaro and M. Angela Sasse. 2015. ‘*"They Brought in the Horrible Key Ring Thing!" Analysing the Usability of Two-Factor Authentication in UK Online Banking*’. In *The 2015 Network and Distributed System Security (NDSS) Symposium: USEC Workshop*. doi:[10.14722/usec.2015.23001](https://doi.org/10.14722/usec.2015.23001).

- Krol, Kat, Jonathan M. Spring, Simon Parkin and M. Angela Sasse. 2016. 'Towards Robust Experimental Design for User Studies in Security and Privacy'. In *Learning from Authoritative Security Experiment Results (LASER) Workshop*, 21–31. San Jose, CA. Accessed 9 June 2017. <https://www.usenix.org/system/files/conference/laser2016/laser2016-paper-krol.pdf>.
- Lawshe, C. H. 1975. 'A Quantitative Approach to Content Validity'. *Personnel Psychology* 28 (4): 563–575. ISSN: 1744-6570. doi:10.1111/j.1744-6570.1975.tb01393.x.
- Lee, Sang M., Sang-Gun Lee and Sangjin Yoo. 2004. 'An Integrative Model of Computer Abuse Based on Social Control and General Deterrence Theories'. *Information & Management* 41 (6): 707–718. ISSN: 0378-7206. doi:10.1016/j.im.2003.08.008.
- Li, Han, Jie Zhang and Rathindra Sarathy. 2010. 'Understanding Compliance with Internet Use Policy from the Perspective of Rational Choice Theory'. *Decision Support Systems* 48 (4): 635–645. ISSN: 0167-9236. doi:10.1016/j.dss.2009.12.005.
- Likert, Rensis. 1932. 'A Technique for the Measurement of Attitudes.' *Archives of psychology*.
- Lin, Ming-Ji James, Shiu-Wan Hung and Chih-Jou Chen. 2009. 'Fostering the Determinants of Knowledge Sharing in Professional Virtual Communities'. *Computers in Human Behavior, Including the Special Issue: The Use of Support Devices in Electronic Learning Environments*, 25 (4): 929–939. ISSN: 0747-5632. doi:10.1016/j.chb.2009.03.008.
- Masson, Michael EJ, and Mary Anne Waldron. 1994. 'Comprehension of Legal Contracts by Non-Experts: Effectiveness of Plain Language Redrafting'. *Applied Cognitive Psychology* 8 (1): 67–85. doi:10.1002/acp.2350080107.
- Mazurek, Michelle L., Saranga Komanduri, Timothy Vidas, Lujo Bauer, Nicolas Christin, Lorrie Faith Cranor, Patrick Gage Kelley, Richard Shay and Blase Ur. 2013. 'Measuring Password Guessability for an Entire University'. In *Proceedings of the 2013 ACM SIGSAC Conference on Computer & Communications Security*, 173–186. CCS '13. New York, NY, USA: ACM. ISBN: 978-1-4503-2477-9. doi:10.1145/2508859.2516726.
- Morgan, Mary S. 2014. 'Resituating Knowledge: Generic Strategies and Case Studies'. *Philosophy of Science* 81 (5): 1012–1024. doi:10.1086/677888.
- Morris, Robert, and Ken Thompson. 1979. 'Password Security: A Case History'. *Commun. ACM* 22 (11): 594–597. ISSN: 0001-0782. doi:10.1145/359168.359172.

- Murdoch, Steven J., Ingolf Becker, Ruba Abu-Salma, Ross Anderson, Nicholas Bohm, Alice Hutchings, M. Angela Sasse and Gianluca Stringhini. 2016. 'Are Payment Card Contracts Unfair?' In *Financial Cryptography and Data Security*, 600–608. Lecture Notes in Computer Science. Springer, Berlin, Heidelberg. ISBN: 978-3-662-54969-8. doi:[10.1007/978-3-662-54970-4\\_35](https://doi.org/10.1007/978-3-662-54970-4_35).
- Murdoch, Steven J., Saar Drimer, Ross Anderson and Mike Bond. 2010. 'Chip and PIN Is Broken'. In *IEEE Symposium on Security and Privacy*, 433–446. Oakland, CA, USA: IEEE. doi:[10.1109/SP.2010.33](https://doi.org/10.1109/SP.2010.33).
- National Cyber Security Centre (NCSC). 2016a. *Password Guidance: Simplifying Your Approach*. Guidance. UK National Cyber Security Centre. Accessed 29 October 2017. <https://www.ncsc.gov.uk/guidance/password-guidance-simplifying-your-approach>.
- . 2016b. 'Security Breaches as Communication: What Are Your Users Telling You?' Accessed 29 March 2018. <https://www.ncsc.gov.uk/blog-post/security-breaches-communication-what-are-your-users-telling-you>.
- O'Leary, Kathleen, Jacob O. Wobbrock and Eve A. Riskin. 2013. 'Q-Methodology As a Research and Design Tool for HCI'. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 1941–1950. CHI '13. New York, NY, USA: ACM. ISBN: 978-1-4503-1899-0. doi:[10.1145/2470654.2466256](https://doi.org/10.1145/2470654.2466256).
- Open Science Collaboration. 2015. 'Estimating the Reproducibility of Psychological Science'. *Science* 349 (6251): aac4716. ISSN: 0036-8075, 1095-9203. doi:[10.1126/science.aac4716](https://doi.org/10.1126/science.aac4716). pmid: [26315443](https://pubmed.ncbi.nlm.nih.gov/26315443/).
- Pahlila, S., M. Siponen and A. Mahmood. 2007. 'Employees' Behavior towards IS Security Policy Compliance'. In *System Sciences, 2007. HICSS 2007. 40th Annual Hawaii International Conference On*, 156b–156b. doi:[10.1109/HICSS.2007.206](https://doi.org/10.1109/HICSS.2007.206).
- Pallas, Frank. 2009. 'Information Security Inside Organizations: A Positive Model and Some Normative Arguments Based on New Institutional Economics'. PhD Thesis, Technische Universität Berlin. [http://opus4.kobv.de/opus4-tuberlin/files/2209/pallas\\_frank.pdf](http://opus4.kobv.de/opus4-tuberlin/files/2209/pallas_frank.pdf).
- Parkin, Simon, Ingolf Becker, Albesë Demjaha, Julienne Park, Nissy Sombatruang and M. Angela Sasse. 2017. *Quantifying the Impact of Password Policy Change*. Preliminary report, GCHQ Small Grants Scheme. London, UK: UCL.
- Parkin, Simon, Samy Driss, Kat Krol and M. Angela Sasse. 2015. 'Assessing the User Experience of Password Reset Policies in a University'. In *Technology and Practice of Passwords*, 21–38. Lecture Notes in Computer Science. Springer, Cham. ISBN: 978-3-319-29937-2 978-3-319-29938-9. doi:[10.1007/978-3-319-29938-9\\_2](https://doi.org/10.1007/978-3-319-29938-9_2).

- Parkin, Simon, Kat Krol, Ingolf Becker and M. Angela Sasse. 2016. ‘*Applying Cognitive Control Modes to Identify Security Fatigue Hotspots*’. In *Twelfth Symposium on Usable Privacy and Security (SOUPS 2016), Workshop on Security Fatigue*. Denver, CO: USENIX Association. <https://www.usenix.org/conference/soups2016/workshop-program/wsf/presentation/parkin>.
- Parkin, Simon, Aad van Moorsel, Philip Inglesant and M. Angela Sasse. 2010. ‘*A Stealth Approach to Usable Security: Helping IT Security Managers to Identify Workable Security Solutions*’. In *Proceedings of the 2010 Workshop on New Security Paradigms*, 33–50. NSPW ’10. New York, NY, USA: ACM. ISBN: 978-1-4503-0415-3. doi:[10.1145/1900546.1900553](https://doi.org/10.1145/1900546.1900553).
- Parsons, Kathryn Marie, Elise Young, Marcus Antanas Butavicius, Agata McCormac, Malcolm Robert Pattinson and Cate Jerram. 2015. ‘*The Influence of Organizational Information Security Culture on Information Security Decision Making*’. *Journal of Cognitive Engineering and Decision Making* 9 (2): 117–129. doi:[10.1177/1555343415575152](https://doi.org/10.1177/1555343415575152).
- Parsons, Kathryn, Dragana Calic, Malcolm Pattinson, Marcus Butavicius, Agata McCormac and Tara Zwaans. 2017. ‘*The Human Aspects of Information Security Questionnaire (HAIS-Q): Two Further Validation Studies*’. *Computers & Security* 66:40–51. doi:[10.1016/j.cose.2017.01.004](https://doi.org/10.1016/j.cose.2017.01.004).
- Parsons, Kathryn, Agata McCormac, Marcus Butavicius, Malcolm Pattinson and Cate Jerram. 2014. ‘*Determining Employee Awareness Using the Human Aspects of Information Security Questionnaire (HAIS-Q)*’. *Computers & Security* 42:165–176. ISSN: 01674048. doi:[10.1016/j.cose.2013.12.003](https://doi.org/10.1016/j.cose.2013.12.003).
- Patrick, Kathleen. 1996. ‘*Reg E: To Pay Or Not To Pay*’. *Bankers’ Hotline* 6 (9). Accessed 5 June 2017. <https://www.bankersonline.com/articles/103455>.
- Paulk, Mark. 2002. ‘*Capability Maturity Model for Software*’. In *Encyclopedia of Software Engineering*. John Wiley & Sons, Inc. ISBN: 978-0-471-02895-6. doi:[10.1002/0471028959.sof589](https://doi.org/10.1002/0471028959.sof589).
- Pfleeger, Shari Lawrence, M. Angela Sasse and Adrian Furnham. 2014. ‘*From Weakest Link to Security Hero: Transforming Staff Security Behavior*’. *Journal of Homeland Security and Emergency Management* 11 (4): 489–510. ISSN: 1547-7355. doi:[10.1515/jhsem-2014-0035](https://doi.org/10.1515/jhsem-2014-0035).
- Pink, Daniel H. 2011. *Drive: The Surprising Truth about What Motivates Us*. Penguin.
- Pinkas, Benny, and Tomas Sander. 2002. ‘*Securing Passwords Against Dictionary Attacks*’. In *Proceedings of the 9th ACM Conference on Computer and Communications Security*, 161–170. CCS ’02. New York, NY, USA: ACM. ISBN: 978-1-58113-612-8. doi:[10.1145/586110.586133](https://doi.org/10.1145/586110.586133).

- Posey, Clay, Tom L Roberts, Paul Benjamin Lowry and Ross T Hightower. 2014. 'Bridging the Divide: A Qualitative Comparison of Information Security Thought Patterns between Information Security Professionals and Ordinary Organizational Insiders'. *Information & Management* 51 (5): 551–567. doi:10.1016/j.im.2014.03.009.
- Prichard, Janet J., and Michael B. Hayden. 2008. 'Assessing the Readability of Free-ware End-User Licensing Agreements'. *Issues in Information Systems* 9 (2): 452–459. Accessed 27 April 2017. <https://pdfs.semanticscholar.org/d735/37e9f628730f00ca80cc3b9fe2ded8005eef.pdf>.
- Prolific. 2018. 'Participant Pool Demographics'. Accessed 30 December 2019. <https://prolific.ac/demographics/>.
- Rajivan, Prashanth, Pablo Moriano, Timothy Kelley and L Jean Camp. 2017. 'Factors in an End User Security Expertise Instrument'. *Information & Computer Security* 25 (2): 190–205. doi:10.1108/ICS-04-2017-0020.
- Reason, James T. 2008. *The Human Contribution: Unsafe Acts, Accidents and Heroic Recoveries*. Ashgate Publishing, Ltd.
- Renaud, Karen, and Wendy Goucher. 2014. 'The Curious Incidence of Security Breaches by Knowledgeable Employees and the Pivotal Role of Security Culture'. In *Human Aspects of Information Security, Privacy, and Trust*, 8533:361–372. LNCS. Springer. doi:10.1007/978-3-319-07620-1\_32.
- Renaud, Karen, and Judith Ramsay. 2014. 'How Helpful Is Colour-Cueing of PIN Entry?' Accessed 18 July 2018. arXiv: 1407.8007 [cs]. <http://arxiv.org/abs/1407.8007>.
- Rhee, Hyeun-Suk, Cheongtag Kim and Young U Ryu. 2009. 'Self-Efficacy in Information Security: Its Influence on End Users' Information Security Practice Behavior'. *Computers & Security* 28 (8): 816–826. ISSN: 0167-4048. doi:10.1016/j.cose.2009.05.008.
- Röder, Nina, Manuel Wiesche, Michael Schermann and Helmut Krcmar. 2014. 'Why Managers Tolerate Workarounds—the Role of Information Systems'. In *Twentieth Americas Conference on Information Systems*. Savannah. <http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.666.3949>.
- Rogers, T.B. 1995. *The Psychological Testing Enterprise: An Introduction*. Psychology Series. Brooks/Cole Publishing Company. ISBN: 978-0-534-21648-1. <https://books.google.co.uk/books?id=iy4QAQAIAAJ>.
- Saltzer, Jerome H., and Michael D. Schroeder. 1975. 'The Protection of Information in Computer Systems'. *Proceedings of the IEEE* 63 (9): 1278–1308. ISSN: 0018-9219. doi:10.1109/PROC.1975.9939.
- SANS Securing The Human. 2016. *Awareness Is Hard: A Tale of Two Challenges*, Security Awareness Report. SANS. Accessed 23 November 2017. <https://securingthehuman.sans.org/media/resources/STH-SecurityAwarenessReport-2016.pdf>.



- . 2017. *It's Time to Communicate*, Security Awareness Report. SANS. Accessed 23 November 2017. <https://securingthehuman.sans.org/media/resources/STH-SecurityAwarenessReport-2017.pdf>.
- Sasse, M. Angela, Matthew Smith, Cormac Herley, Heather Lipford and Kami Vaniea. 2016. 'Debunking Security-Usability Tradeoff Myths'. *IEEE Security & Privacy* 14 (5): 33–39. doi:10.1109/MSP.2016.110.
- Sasse, M. Angela, Michelle Steves, Kat Krol and Dana Chisnell. 2014. 'The Great Authentication Fatigue—and How to Overcome It'. In *International Conference on Cross-Cultural Design*, 8528:228–239. LNCS. Springer, Springer. doi:10.1007/978-3-319-07308-8\_23.
- Sasse, Martina Angela, Sacha Brostoff and Dirk Weirich. 2001. 'Transforming the 'Weakest Link' - a Human/Computer Interaction Approach to Usable and Effective Security'. *BT Technology Journal* 19 (3): 122–131. ISSN: 1358-3948, 1573-1995. doi:10.1023/A:1011902718709.
- Schein, Edgar H. 2010. *Organizational Culture and Leadership*. Vol. 2. John Wiley & Sons.
- Segreti, Sean M., William Melicher, Saranga Komanduri, Darya Melicher, Richard Shay, Blase Ur, Lujo Bauer, Nicolas Christin, Lorrie Faith Cranor and Michelle L. Mazurek. 2017. 'Diversify to Survive: Making Passwords Stronger with Adaptive Policies'. In *Thirteenth Symposium on Usable Privacy and Security (SOUPS 2017)*, 1–12. Santa Clara, CA: USENIX Association. ISBN: 978-1-931971-39-3. <https://www.usenix.org/conference/soups2017/technical-sessions/presentation/segreti>.
- Shay, Richard, Lujo Bauer, Nicolas Christin, Lorrie Faith Cranor, Alain Forget, Saranga Komanduri, Michelle L. Mazurek, William Melicher, Sean M. Segreti and Blase Ur. 2015. 'A Spoonful of Sugar?: The Impact of Guidance and Feedback on Password-Creation Behavior'. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, 2903–2912. CHI '15. New York, NY, USA: ACM. ISBN: 978-1-4503-3145-6. doi:10.1145/2702123.2702586.
- Shay, Richard, Saranga Komanduri, Adam L. Durity, Phillip (Seyoung) Huh, Michelle L. Mazurek, Sean M. Segreti, Blase Ur, Lujo Bauer, Nicolas Christin and Lorrie Faith Cranor. 2014. 'Can Long Passwords Be Secure and Usable?' In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 2927–2936. CHI '14. New York, NY, USA: ACM. ISBN: 978-1-4503-2473-1. doi:10.1145/2556288.2557377.
- Siponen, Mikko, Seppo Pahlila and Adam Mahmood. 2007. 'Employees' Adherence to Information Security Policies: An Empirical Study'. In *New Approaches for Security, Privacy and Trust in Complex Environments*, 133–144. IFIP International Federation for Information Processing. Springer, Boston, MA. ISBN: 978-0-387-72366-2. doi:10.1007/978-0-387-72367-9\_12.

- Siponen, Mikko, and Anthony Vance. 2010. 'Neutralization: New Insights into the Problem of Employee Information Systems Security Policy Violations'. *MIS quarterly* 34 (3): 487. doi:10.2307/25750688.
- . 2014. 'Guidelines for Improving the Contextual Relevance of Field Surveys: The Case of Information Security Policy Violations'. *European Journal of Information Systems* 23 (3): 289–305. ISSN: 0960-085X, 1476-9344. doi:10.1057/ejis.2012.59.
- Smith, H. Jeff. 1993. 'Privacy Policies and Practices: Inside the Organizational Maze'. *Commun. ACM* 36 (12): 104–122. ISSN: 0001-0782. doi:10.1145/163298.163349.
- Sohrabi Safa, Nader, Rossouw Von Solms and Steven Furnell. 2016. 'Information Security Policy Compliance Model in Organizations'. *Computers & Security* 56:70–82. ISSN: 0167-4048. doi:10.1016/j.cose.2015.10.006.
- Sommestad, Teodor, Jonas Hallberg, Kristoffer Lundholm and Johan Bengtsson. 2014. 'Variables Influencing Information Security Policy Compliance: A Systematic Review of Quantitative Studies'. *Information Management & Computer Security* 22 (1): 42–75. doi:10.1108/IMCS-08-2012-0045.
- Squire, Larry R. 1989. 'On the Course of Forgetting in Very Long-Term Memory'. *Journal of Experimental Psychology: Learning, Memory, and Cognition*. 15 (2): 241–245. ISSN: 0278-7393. doi:10.1037/0278-7393.15.2.241.
- Stanley, Anne. 2003. *Voting With Your Feet: Consumers' Problems With Credit Cards and Exit Behaviors*. Discussion Paper, Payment Cards Center. Philadelphia, PA: Federal Reserve Bank of Philadelphia. Accessed 27 September 2018. <https://pdfs.semanticscholar.org/ce23/2f2c0e75a7f722adfc8569f7e2f16d2e06d9.pdf>.
- Stanton, Jeffrey M, Kathryn R Stam, Paul Mastrangelo and Jeffrey Jolton. 2005. 'Analysis of End User Security Behaviors'. *Computers & Security* 24 (2): 124–133. doi:10.1016/j.cose.2004.07.001.
- Stephenson, W. 1935. 'Correlating Persons Instead of Tests'. *Journal of Personality* 4 (1): 17–24. ISSN: 1467-6494. doi:10.1111/j.1467-6494.1935.tb02022.x.
- Steves, Michelle, Dana Chisnell, Martina Angela Sasse, Kat Krol, Mary Theofanos and Hannah Wald. 2014. *Report: Authentication Diary Study* NIST IR 7983. National Institute of Standards and Technology. Accessed 15 July 2014. <http://nvlpubs.nist.gov/nistpubs/ir/2014/NIST.IR.7983.pdf>.
- Stewart, Rachelle S. 2007. 'Protective Measures for Private Health Information'. *Perspectives in Health Information Management* 4:5. ISSN: 1559-4122. pmid: 18066355.
- Straub, Detmar W. 1989. 'Validating Instruments in MIS Research'. *MIS Quarterly* 13 (2): 147–169. ISSN: 0276-7783. doi:10.2307/248922. JSTOR: 248922.

- Sullivan, Richard J. 2014. *Controlling Security Risk and Fraud in Payment Systems*. Economic Review Third Quarter 2014. Federal Reserve Bank of Kansas City.
- Tamjidyamcholo, Alireza, Mohd Sapiyan Bin Baba, Nor Liyana Mohd Shuib and Vala Ali Rohani. 2014. 'Evaluation Model for Knowledge Sharing in Information Security Professional Virtual Community'. *Computers & Security* 43:19–34. ISSN: 0167-4048. doi:10.1016/j.cose.2014.02.010.
- Tamjidyamcholo, Alireza, Mohd Sapiyan Bin Baba, Hamed Tamjid and Rahmatollah Gholipour. 2013. 'Information Security – Professional Perceptions of Knowledge-Sharing Intention under Self-Efficacy, Trust, Reciprocity, and Shared-Language'. *Computers & Education* 68:223–232. ISSN: 0360-1315. doi:10.1016/j.compedu.2013.05.010.
- Tariq, Muhammad Adnan, Joel Brynielsson and Henrik Artman. 2014. 'The Security Awareness Paradox: A Case Study'. In *Advances in Social Networks Analysis and Mining (ASONAM)*, 704–711. IEEE. <https://dl.acm.org/citation.cfm?id=3191974>.
- Thayer-Hart, Nancy, Jennifer Dykema, Kelly Elver, Nora Cate Schaeffer and John Stevenson. 2010. *Survey Fundamentals: A Guide to Designing and Implementing Surveys*. University of Wisconsin.
- Thomson, Kerry-Lynn, and Rossouw von Solms. 2006. 'Towards an Information Security Competence Maturity Model'. *Computer Fraud & Security* 2006 (5): 11–15. doi:10.1016/S1361-3723(06)70356-6.
- Thurstone, L L. 1931. 'The Measurement of Social Attitudes.' *The Journal of abnormal and social psychology* 26 (3): 249–269. ISSN: 0096-851X.
- Topa, Ioanna, and Maria Karyda. 2015. 'Identifying Factors That Influence Employees' Security Behavior for Enhancing ISP Compliance'. In *International Conference on Trust and Privacy in Digital Business*, 9264:169–179. LNCS. Springer. doi:10.1007/978-3-319-22906-5\_13.
- Tsohou, Aggeliki, Maria Karyda and Spyros Kokolakis. 2015. 'Analyzing the Role of Cognitive and Cultural Biases in the Internalization of Information Security Policies: Recommendations for Information Security Awareness Programs'. *Computers & security* 52:128–141. doi:10.1016/j.cose.2015.04.006.
- UCISA. 2015. *Chapter 8: Roles and Competencies*, UCISA Information Security Management Toolkit. Accessed 1 November 2017. <https://www.ucisa.ac.uk/representation/activities/ismt/chapt8>.
- Ur, Blase, Sean M. Segreti, Lujo Bauer, Nicolas Christin, Lorrie Faith Cranor, Saranga Komanduri, Darya Kurilova, Michelle L. Mazurek, William Melicher and Richard Shay. 2015. 'Measuring Real-World Accuracies and Biases in Modeling Password Guessability'. In *24th USENIX Security Symposium (USENIX Security 15)*, 463–481. Washington, D.C.: USENIX Association. ISBN: 978-1-931971-23-2. <http://blogs.usenix.org/conference/usenixsecurity15/technical-sessions/presentation/ur>.

- Van Oorschot, Paul C., and Stuart Stubblebine. 2006. 'On Countering Online Dictionary Attacks with Login Histories and Humans-in-the-Loop'. *ACM Trans. Inf. Syst. Secur.* 9 (3): 235–258. ISSN: 1094-9224. doi:10.1145/1178618.1178619.
- Vaniaea, Kami E., Emilee Rader and Rick Wash. 2014. 'Betrayed by Updates: How Negative Experiences Affect Future Security'. In *Proceedings of the 32Nd Annual ACM Conference on Human Factors in Computing Systems*, 2671–2674. CHI '14. New York, NY, USA: ACM. ISBN: 978-1-4503-2473-1. doi:10.1145/2556288.2557275.
- Waddell, T. Franklin, Joshua R. Auriemma and S. Shyam Sundar. 2016. 'Make It Simple, or Force Users to Read?: Paraphrased Design Improves Comprehension of End User License Agreements'. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, 5252–5256. CHI '16. New York, NY, USA: ACM. ISBN: 978-1-4503-3362-7. doi:10.1145/2858036.2858149.
- Wash, Rick, Emilee Rader and Chris Fennell. 2017. 'Can People Self-Report Security Accurately?: Agreement Between Self-Report and Behavioral Measures'. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, 2228–2232. ACM. doi:10.1145/3025453.3025911.
- Weir, Charles Alexander Forbes, Lynne Blair, James Noble, Ingolf Becker and M. Angela Sasse. 2018. 'Light-Touch Interventions to Improve Software Development Security'. In *IEEE SecDev*, 1:85–93. Cambridge, MA, USA: IEEE Computer Society. doi:10.1109/SecDev.2018.00019.
- Weir, Matt, Sudhir Aggarwal, Michael Collins and Henry Stern. 2010. 'Testing Metrics for Password Creation Policies by Attacking Large Sets of Revealed Passwords'. In *Proceedings of the 17th ACM Conference on Computer and Communications Security*, 162–175. CCS '10. New York, NY, USA: ACM. ISBN: 978-1-4503-0245-6. doi:10.1145/1866307.1866327.
- Wheeler, Daniel Lowe. 2016. 'Zxcvbn: Low-Budget Password Strength Estimation'. In *25th USENIX Security Symposium (USENIX Security 16)*, 157–173. Austin, TX: USENIX Association. ISBN: 978-1-931971-32-4. <https://www.usenix.org/conference/usenixsecurity16/technical-sessions/presentation/wheeler>.
- Whitten, Alma. 2004. 'Making Security Usable'. Unpublished Ph. D. thesis, CS, CMU.
- Whitten, Alma, and J. Doug Tygar. 1999. 'Why Johnny Can't Encrypt: A Usability Evaluation of PGP 5.0'. In *Proceedings of the 8th USENIX Security Symposium*, vol. 99. Washington, D.C.: USENIX Association.
- Witherspoon, Candace L., Jason Bergner, Cam Cockrell and Dan N. Stone. 2013. 'Antecedents of Organizational Knowledge Sharing: A Meta-Analysis and Critique'. *Journal of Knowledge Management* 17 (2): 250–277. ISSN: 1367-3270. doi:10.1108/13673271311315204.

- Wogalter, Michael S., and Matthew Ryan Hayes. 2014. 'Online and Software Licensing Agreements: User Beliefs and Expectations of Risks'. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting* 58 (1): 1391–1394. ISSN: 1541-9312. doi:[10.1177/1541931214581290](https://doi.org/10.1177/1541931214581290).
- Wogalter, Michael S., Julie E. Howe, Alla H. Sifuentes and James Luginbuhl. 1999. 'On the Adequacy of Legal Documents: Factors That Influence Informed Consent'. *Ergonomics* 42 (4): 593–613. ISSN: 0014-0139. doi:[10.1080/001401399185504](https://doi.org/10.1080/001401399185504).
- Woon, Irene M. Y., and Atreyi Kankanhalli. 2007. 'Investigation of IS Professionals' Intention to Practise Secure Development of Applications'. *International Journal of Human-Computer Studies*, Information Security in the Knowledge Economy, 65 (1): 29–41. ISSN: 1071-5819. doi:[10.1016/j.ijhcs.2006.08.003](https://doi.org/10.1016/j.ijhcs.2006.08.003).
- Workman, Michael, William H. Bommer and Detmar Straub. 2008. 'Security Lapses and the Omission of Information Security Measures: A Threat Control Model and Empirical Test'. *Computers in Human Behavior*, Including the Special Issue: Electronic Games and Personalized eLearning Processes, 24 (6): 2799–2816. ISSN: 0747-5632. doi:[10.1016/j.chb.2008.04.005](https://doi.org/10.1016/j.chb.2008.04.005).
- Yan, Jeff, Alan Blackwell, Ross J. Anderson and Alasdair Grant. 2004. 'Password Memorability and Security: Empirical Results'. *IEEE Security Privacy* 2 (5): 25–31. ISSN: 1540-7993. doi:[10.1109/MSP.2004.81](https://doi.org/10.1109/MSP.2004.81).
- Young, Richard M. 1983. 'Surrogates and Mappings: Two Kinds of Conceptual Models for Interactive Devices'. *Mental models* 37:35–52.
- Zeuschwitz, Emanuel von, Alexander De Luca and Heinrich Hussmann. 2013. 'Survival of the Shortest: A Retrospective Analysis of Influencing Factors on Password Composition'. In *Human-Computer Interaction – INTERACT 2013*, 460–467. Lecture Notes in Computer Science. Springer, Berlin, Heidelberg. ISBN: 978-3-642-40476-4 978-3-642-40477-1. doi:[10.1007/978-3-642-40477-1\\_28](https://doi.org/10.1007/978-3-642-40477-1_28).
- Zhang-Kennedy, L., S. Chiasson and P. van Oorschot. 2016. 'Revisiting Password Rules: Facilitating Human Management of Passwords'. In *2016 APWG Symposium on Electronic Crime Research (eCrime)*, 1–10. doi:[10.1109/ECRIME.2016.7487945](https://doi.org/10.1109/ECRIME.2016.7487945).
- Zhang, Yinqian, Fabian Monrose and Michael K. Reiter. 2010. 'The Security of Modern Password Expiration: An Algorithmic Framework and Empirical Analysis'. In *Proceedings of the 17th ACM Conference on Computer and Communications Security*, 176–186. CCS '10. New York, NY, USA: ACM. ISBN: 978-1-4503-0245-6. doi:[10.1145/1866307.1866328](https://doi.org/10.1145/1866307.1866328).





## Productive Security

This Appendix contains a number of additional statistical tables that may be of interest to the reader. The content is discussed in [Section 6.5](#). [Appendices A.3](#) and [A.4](#) contain the scenarios that have been used in Companies A and B respectively. Finally, [Appendix A.5](#) contains the Maturity Model that underpins the Maturity Levels described in [Section 6.4.5.2](#). Neither the formulation of any of the scenarios or the maturity model are my own work, but are not publicly available otherwise.

## A.1 Company A additional tables

Business location	Scenario Accept <sup>**</sup>	Behaviour 1		Behaviour 2		Behaviour 3		Behaviour 4		$\tau$
		Rank <sup>*</sup>	Sev <sup>**</sup>	Rank <sup>*</sup>	Sev	Rank	Sev	Rank <sup>**</sup>	Sev <sup>**</sup>	
Other	-0.14 <sup>**</sup>	-0.04	0.11 <sup>**</sup>	0.07 <sup>**</sup>	0.01	-0.01 <sup>**</sup>	0.01	-0.02 <sup>**</sup>	0.21 <sup>**</sup>	-0.28 <sup>**</sup>
UK	0.12 <sup>**</sup>	-0.05	0.01 <sup>**</sup>	-0.05 <sup>**</sup>	0.02 <sup>*</sup>	0.01 <sup>**</sup>	0.04	0.08 <sup>**</sup>	-0.04	-0.30 <sup>**</sup>
US	-0.04	0.08 <sup>**</sup>	-0.09 <sup>**</sup>	0.00 <sup>**</sup>	-0.03	-0.01 <sup>**</sup>	-0.06 <sup>*</sup>	-0.08 <sup>**</sup>	-0.10 <sup>**</sup>	-0.26 <sup>**</sup>
mean	2.43	1.93	4.35 <sup>**</sup>	2.62 <sup>**</sup>	3.51	2.34 <sup>**</sup>	3.52	3.11 <sup>**</sup>	3.46	-0.28 <sup>**</sup>

(a) Behaviour types rankings and behaviour severity score split by business locations.

Business location	Attitude 1		Attitude 2		Attitude 3		Attitude 4		$\tau$
	Rank	Approp	Rank	Approp <sup>*</sup>	Rank <sup>**</sup>	Approp	Rank <sup>*</sup>	Approp	
Other	0.03	-0.04	0.02 <sup>**</sup>	0.01 <sup>**</sup>	0.02 <sup>**</sup>	0.04 <sup>**</sup>	-0.07 <sup>**</sup>	-0.03 <sup>**</sup>	0.74 <sup>**</sup>
UK	-0.02	-0.01	-0.03 <sup>**</sup>	-0.04 <sup>**</sup>	0.03 <sup>**</sup>	0.03 <sup>**</sup>	0.02 <sup>**</sup>	-0.00 <sup>**</sup>	0.73 <sup>**</sup>
US	0.01	0.04	0.01 <sup>**</sup>	0.04 <sup>*</sup>	-0.05 <sup>**</sup>	-0.06 <sup>**</sup>	0.03 <sup>**</sup>	0.03 <sup>**</sup>	0.70 <sup>**</sup>
mean	1.27	1.28	3.58 <sup>**</sup>	4.58 <sup>**</sup>	2.92 <sup>**</sup>	4.13 <sup>**</sup>	2.23 <sup>**</sup>	2.56 <sup>**</sup>	0.72 <sup>**</sup>

(b) Attitude type rankings and appropriateness score split by business locations.

**Table A.1: Additional Factor Analysis in Company A** The values in each cell of the tables above describe the variation from the mean in their column, with the mean being shown at the bottom (the mean is the value for the organisation as a whole). Based on the scoring in Section 6.4.6, higher ranks imply more popular choices. Similarly, the higher the Approp/Sev score, the more appropriate/severe the participants take the option to be. In the second row, the <sup>\*\*/\*</sup> after Rank/Approp/Sev show statistical significant variations from the median rank or acceptability or severity score respectively based on the Kruskal-Wallis H-test for independent samples at  $p < 0.01/p < 0.05$  confidence respectively. If this Kruskal-Wallis test shows statistical significance, for each subgroup a two-sided Mann-Whitney rank test between this subgroup and the union of all other subgroups is carried out; the results of these tests are shown by further <sup>\*\*/\*</sup> at each number, showing statistical significance at  $p < 0.01/p < 0.05$  confidence respectively.

Further, the colours show the order of mean Rank/Approp/Sev for each of the groups (i.e., ranking them horizontally). The largest mean is given the darkest colour, and the colour changes to a lighter shade if there is a statistically significant difference between the distribution of ranks/scores of the current mean and the next largest mean, based on a one-sided paired Wilcoxon rank test. This statistical test is further shown by <sup>\*\*/\*</sup> at the value of the higher cell, showing  $p < 0.01/p < 0.05$  confidence respectively. If more than one cell contains the same colour, there is no statistical significant variation between the rankings/scores for these options.

Lastly, the rightmost column  $\tau$  lists Kendall's  $\tau$  correlation coefficients between the rank and the acceptability/severity score respectively for each of the groups. Kendall's  $\tau$  ranges from  $-1$  (perfect anti-correlation) to  $1$  (perfect correlation). <sup>\*\*/\*</sup> signifies rejecting the null hypothesis of independence (i.e.  $\tau = 0$ ) with statistical significance at  $p < 0.01/p < 0.05$  confidence respectively.



Age	Attitude 1		Attitude 2		Attitude 3		Attitude 4		$\tau$
	Rank	Approp	Rank	Approp	Rank	Approp*	Rank	Approp	
Less than 25	0.04	0.10	0.01**	-0.03**	-0.02**	-0.20**	-0.03**	-0.09**	0.73**
25 – 29	0.03	0.10	0.04**	-0.01**	-0.06**	-0.17**	-0.01**	-0.01**	0.70**
30 – 34	0.04	-0.03	0.04**	0.05**	-0.06**	0.00**	-0.01**	0.07**	0.74**
35 – 39	0.02	-0.02	-0.02**	0.04**	0.04**	0.07**	-0.04**	-0.07**	0.75**
40 – 44	0.01	-0.04	-0.02**	-0.01**	0.01**	0.03**	0.00**	-0.06**	0.73**
45 – 49	-0.02	-0.01	-0.01**	-0.04**	0.03**	0.03**	0.01**	0.04**	0.71**
50 – 54	-0.03	-0.07	-0.02**	0.06**	0.01**	0.09**	0.04**	-0.02**	0.74**
55 or Over	-0.07	0.07	0.00**	-0.02**	0.05**	0.03**	0.02**	0.09**	0.69**
mean	1.27	1.28	3.58**	4.58**	2.92**	4.13**	2.23**	2.56**	0.72**

(c) Attitude types rankings and appropriateness score split by age groups.

Age	Scenario Accept	Behaviour 1		Behaviour 2		Behaviour 3		Behaviour 4		$\tau$
		Rank	Sev	Rank	Sev*	Rank	Sev*	Rank	Sev	
Less than 25	0.01	0.09	-0.21**	-0.20**	-0.10	0.09	-0.25*	0.03**	0.08	-0.23**
25 – 29	-0.18	-0.06	-0.05**	-0.04*	-0.11	0.07**	-0.20**	0.03**	-0.03	-0.27**
30 – 34	-0.02	-0.02	-0.05**	-0.02**	-0.13*	-0.01**	-0.04	0.05**	-0.07	-0.28**
35 – 39	-0.09	0.00	0.04**	-0.02**	-0.01	0.03**	0.06	-0.01**	0.17	-0.27**
40 – 44	-0.06	0.01	0.06**	0.09**	-0.01	-0.09**	0.07	-0.01**	-0.02	-0.29**
45 – 49	0.06	0.01	0.03**	0.01**	0.08*	0.01**	0.07	-0.03**	-0.02	-0.30**
50 – 54	0.26	-0.04	0.08**	0.08**	0.06**	-0.08**	0.06	0.04**	-0.12	-0.30**
55 or Over	0.02	0.04	-0.04**	-0.03*	0.15*	0.06**	0.04	-0.06**	0.08	-0.25**
mean	2.43	1.93	4.35**	2.62**	3.51	2.34**	3.52	3.11**	3.46	-0.28**

(d) Behaviour types rankings and behaviour severity score split by age groups.

Employee category	Attitude 1		Attitude 2		Attitude 3		Attitude 4		$\tau$
	Rank**	Approp**	Rank**	Approp**	Rank**	Approp**	Rank**	Approp**	
External, Offsite	0.05	-0.03	-0.18**	-0.42**	0.28**	0.20	-0.14**	-0.21**	0.75**
External, Onsite	0.12**	0.04	-0.28**	-0.04**	-0.14**	-0.03**	0.30**	0.42**	0.63**
Internal, Offsite	-0.02	0.07**	-0.09*	-0.66**	0.46**	0.36**	-0.34**	-0.76**	0.66**
Internal, Onsite	0.00*	-0.03**	0.11**	0.04**	0.02**	-0.02**	-0.13**	-0.12**	0.75**
Other	-0.06**	0.03	-0.14**	0.21**	-0.23**	-0.09**	0.43**	0.52**	0.68**
mean	1.27	1.28	3.58**	4.58**	2.92**	4.13**	2.23**	2.56**	0.72**

(e) Attitude type rankings and appropriateness score split by employee category.

Employee category	Scenario Accept	Behaviour 1		Behaviour 2		Behaviour 3		Behaviour 4		$\tau$
		Rank**	Sev**	Rank**	Sev*	Rank**	Sev**	Rank**	Sev**	
External, Offsite	0.44*	-0.10	-0.16**	0.30**	-0.12**	-0.57**	0.74**	0.38**	-0.89**	-0.55**
External, Onsite	0.65**	-0.05	0.16**	0.09**	-0.20**	-0.33**	0.04**	0.29**	-0.31**	-0.27**
Internal, Offsite	-0.12	-0.29**	0.08**	0.08**	0.13**	0.40**	0.13	-0.20**	-0.11	-0.46**
Internal, Onsite	-0.33**	0.10**	0.05**	0.02**	0.00**	0.07**	-0.14**	-0.19**	0.41**	-0.19**
Other	0.71**	-0.16**	-0.24**	-0.16**	0.05**	-0.17**	0.27**	0.49**	-0.96**	-0.43**
mean	2.43	1.93	4.35**	2.62**	3.51	2.34**	3.52	3.11**	3.46	-0.28**

(f) Behaviour types rankings and behaviour severity score split by employee category.

Table A.1: Concluded

## A.2 Company B additional tables

Age	Level 2		Level 3		Level 4		Level 5		$\tau$
	Rank	Accept	Rank*	Accept**	Rank**	Accept	Rank**	Accept**	
Less than 25	-0.10	-0.10	0.12**	0.31**	-0.09**	-0.00**	0.07**	0.16**	0.63**
25 – 29	-0.18	-0.08	-0.01**	0.24**	0.10**	0.15**	0.10**	-0.12**	0.69**
30 – 34	0.04	0.06	0.11**	0.24**	-0.10**	0.03**	-0.06**	-0.02**	0.64**
35 – 39	0.10	0.09	0.02**	-0.05**	0.13**	0.06**	-0.24**	-0.18**	0.58**
40 – 44	0.06	0.00	0.03**	-0.13**	-0.12**	-0.08**	0.03**	0.10**	0.60**
45 – 49	0.06	-0.03	-0.18**	-0.41**	-0.04**	-0.01**	0.16**	0.20**	0.59**
50 – 54	-0.08	0.01	-0.18**	-0.45**	0.10**	-0.33**	0.16	0.04**	0.63**
55 or over	0.04	-0.07	-0.20*	-0.32**	0.17**	-0.23**	-0.01	0.02**	0.56**
mean	1.48	1.50	2.13**	2.19**	3.08**	3.98**	3.30**	4.51**	0.62**

(a) Maturity level rankings and acceptability score split by age groups.

**Table A.2: Additional Factor Analysis in Company B** The values in each cell of the tables above describe the variation from the mean in their column, with the mean being shown at the bottom (the mean is the value for the organisation as a whole). Based on the scoring in Section 6.4.6, higher ranks imply more popular choices. Similarly, the higher the Accept/Sev score, the more acceptable/severe the participants take the option to be. In the second row, the \*\*/\* after Rank/Accept/Sev show statistical significant variations from the median rank or acceptability or severity score respectively based on the Kruskal-Wallis H-test for independent samples at  $p < 0.01/p < 0.05$  confidence respectively. If this Kruskal-Wallis test shows statistical significance, for each subgroup a two-sided Mann-Whitney rank test between this subgroup and the union of all other subgroups is carried out; the results of these tests are shown by further \*\*/\* at each number, showing statistical significance at  $p < 0.01/p < 0.05$  confidence respectively.

Further, the colours show the order of mean Rank/Accept/Sev for each of the groups (i.e., ranking them horizontally). The largest mean is given the darkest colour, and the colour changes to a lighter shade if there is a statistically significant difference between the distribution of ranks/scores of the current mean and the next largest mean, based on a one-sided paired Wilcoxon rank test. This statistical test is further shown by \*\*/\* at the value of the higher cell, showing  $p < 0.01/p < 0.05$  confidence respectively. If more than one cell contains the same colour, there is no statistical significant variation between the rankings/scores for these options.

Lastly, the rightmost column  $\tau$  lists Kendall's  $\tau$  correlation coefficients between the rank and the acceptability/severity score respectively for each of the groups. Kendall's  $\tau$  ranges from  $-1$  (perfect anti-correlation) to  $1$  (perfect correlation). \*\*/\* signifies rejecting the null hypothesis of independence (i.e.  $\tau = 0$ ) with statistical significance at  $p < 0.01/p < 0.05$  confidence respectively.

Age	Scenario Sev	Individualist		Egalitarian		Hierarchist		Fatalist		$\tau$
		Rank	Sev	Rank	Sev	Rank**	Sev	Rank**	Sev	
Less than 25	-0.02	0.05	-0.01*	0.11	0.17	-0.12	-0.19	-0.04	0.34*	-0.17**
25 – 29	-0.30	-0.19	0.24	-0.04	-0.12*	0.38**	0.05	-0.14*	0.21	-0.14**
30 – 34	0.04	0.02**	-0.02	-0.06	0.00*	0.25**	-0.19	-0.20*	-0.08**	-0.17**
35 – 39	0.09	0.15	-0.13	0.04	-0.14	-0.18	0.07	-0.01**	-0.05	-0.22**
40 – 44	-0.01	0.04	-0.07	0.01	-0.21	-0.18**	0.27	0.13	0.09	-0.21**
45 – 49	-0.02	0.01	-0.06	0.15	0.07**	-0.17	0.09	0.00	-0.03	-0.30**
50 – 54	0.36	-0.07	0.04*	-0.22	0.53**	-0.26**	-0.13	0.55**	-0.54	-0.41**
55 or over	0.28	-0.18*	0.10*	-0.04	0.58*	-0.28	-0.25	0.50*	-0.53	-0.24**
mean	2.24	2.68**	3.49	2.02	3.76**	2.80*	3.20	2.50**	3.44**	-0.20**

(b) Behaviour types rankings and behaviour severity score split by age groups.

Business location	Level 2		Level 3		Level 4		Level 5		$\tau$
	Rank**	Accept**	Rank	Accept**	Rank	Accept	Rank**	Accept*	
1: HQ	0.28**	0.28**	0.23**	-0.07**	-0.01**	0.01**	-0.51**	-0.15**	0.44**
2	-0.23*	-0.19	-0.00**	0.23**	0.05**	0.07**	0.19*	0.03**	0.74**
3	-0.14	-0.20**	-0.12**	0.20**	-0.05**	-0.07**	0.30**	0.09**	0.72**
4	-0.30**	-0.21*	-0.16**	-0.05**	0.24**	-0.10**	0.22	-0.10**	0.72**
5	0.13	0.08	0.03**	-0.13**	-0.03**	0.10**	-0.13	-0.04**	0.59**
Homeworker	0.06	0.09	-0.04**	-0.26**	-0.08**	0.06**	0.06**	0.11**	0.60**
Minor offices	0.28	-0.05	-0.04	-0.71**	-0.08**	-0.22**	-0.16	0.40**	0.52**
Other	-0.19**	-0.12	-0.10**	0.22**	0.00**	-0.05**	0.29**	0.04**	0.73**
mean	1.48	1.50	2.13**	2.19**	3.08**	3.98**	3.30**	4.51**	0.62**

(c) Maturity level rankings and acceptability score split by business locations.

Business location	Scenario Sev**	Individualist		Egalitarian		Hierarchist		Fatalist		$\tau$
		Rank**	Sev**	Rank	Sev**	Rank**	Sev*	Rank**	Sev**	
1: HQ	0.12*	0.21**	-0.30**	0.01	-0.27**	-0.36**	-0.05	0.15	-0.32**	-0.31**
2	0.32	-0.05*	0.01	0.13	0.08*	0.42*	-0.13	-0.50**	-0.06	-0.07
3	0.07	-0.10	0.78**	0.01	0.73**	0.33*	0.66**	-0.24	0.98**	-0.23**
4	-0.31*	-0.02*	-0.00	0.18	-0.09	0.28*	-0.00	-0.44**	0.15	-0.24**
5	-0.03	0.09	-0.03*	-0.02	0.03**	-0.28**	-0.05	0.21*	-0.15	-0.24**
Homeworker	0.51**	0.19*	-0.03	0.24	-0.10	-0.37*	0.01	-0.06	-0.09	-0.31**
Minor offices	0.42*	-0.35*	-0.37**	-0.16	0.29**	-0.32	-0.08	0.83**	-0.99**	-0.67**
Other	-0.29**	-0.20**	0.05**	-0.12	-0.07	0.37**	-0.10	-0.05**	0.17*	-0.11**
mean	2.24	2.68**	3.49	2.02	3.76**	2.80*	3.20	2.50**	3.44**	-0.20**

(d) Behaviour types rankings and behaviour severity score split by business locations.

Table A.2: Concluded

### A.3 Company A Behaviour and Attitude scenarios

Below, we provide the texts of the behaviour and attitude scenarios used in the study. Labels are included to indicate which option related to which behaviour and attitude type, but these were not displayed to participants in the study and the order of the options was randomised as well.

#### A.3.1 Scenario A (Behaviour): File Sharing

Jessica, a Business Analyst at CompA, needs to share a large volume of files with colleagues in her department as part of a high priority task she is undertaking. These files contain “Confidential” company information for “internal use only”.

Jessica has made the files available through Microsoft SharePoint, restricting access to certain team members. Some team members tell her they cannot access the files due to incorrect permissions, so Jessica has submitted a request for changes to be made to her colleagues’ permissions and escalated this due to the urgency.

However, she knows from past experience that it may take up to 1 week for the changes to be approved and applied. If these files are not made available within the next 2 working days, this will severely impact delivery and quality. As not all of her colleagues require access to all the files, to manually distribute them would involve her identifying the subset of files for each person—this will be a very time-consuming task, so Jessica creates an archive of all the relevant documents and considers how best to deliver it to the group.

If you were Jessica, how acceptable do you think it would be NOT to take any further action at this stage? [5-point Likert from Not Acceptable at All to Very Acceptable]

Assuming that Jessica chooses to take some action at this stage, if you were her, what would you do in these circumstances?

**Behaviour Type 1:** Move the files to an unrestricted folder on the internal network to allow the work group to have continued access to it.

**Behaviour Type 2:** Burn a copy of the files onto a CD/DVD and distribute to the work group.

**Behaviour Type 3:** Request that those with access share their (main log-in) account details and passwords with those without to allow them access to the information.

**Behaviour Type 4:** Email the document archive directly to the general work group mailing list using your company email address.

### A.3.2 Scenario B (Behaviour): Managing Permissions

John is a System Administrator at CompB responsible for deciding who has access to confidential information.

John normally reviews each request and then applies the most appropriate permissions, or the request is denied according to established procedures and guidelines. He undertakes this task every 24 hours to ensure there is no risk of maintenance schedules slipping due to a lack of access to records.

John is called away from the office on short notice by a family emergency and he is concerned about how this task will be managed during his absence. The system used to set the permissions does not easily allow him to deputise the task to another account, so he must find another way to ensure this activity is completed while he is away. He is also concerned that as the guidelines are not always clear and require some degree of discretion when granting access, deputising the task may mean there is a higher risk of incorrect permissions being granted. This makes the choice of who to trust this task to in his absence an important one.

How acceptable do you think it would be for John to leave without delegating the task in this situation? [5-point Likert from Not Acceptable at All to Very Acceptable]

Assuming that John decides not to leave without putting some kind of solution in place to cover his absence, if you were John, what would you do in these circumstances?

**Behaviour Type 1:** Grant blanket access rights to the whole department (clone of the permissions of an individual with the most access rights) for the duration of your absence to forestall many of the access requests you are usually asked to deal with.

**Behaviour Type 2:** Give out login details of a range of access permissions (used by temporary workers) with instructions that they be used where existing permissions do not allow access.

**Behaviour Type 3:** Leave your password with a trusted member of the department and ask them to handle “all decision making” while you are away.

**Behaviour Type 4:** Leave your password with your secretary, who although temporary, is a trusted employee, with instructions to use your account to resolve “emergency situations”.

### A.3.3 Scenario C (Behaviour): USB Stick Usage

Jason works for CompA as a Commercial Analyst and is currently involved in an important project that requires him to present progress updates to clients, often in offsite locations.

Jason would normally use his laptop to take presentations to clients but his laptop is currently in for maintenance. Instead he decides to use an encrypted USB memory stick to transfer the required files to the client site. Unfortunately, shortly before he is due to leave for the meeting, Jason realises he lent out his encrypted USB stick and cannot recall who to. He knows he will not get a replacement at such short notice. In the meantime he still needs some way to transfer information. The presentation includes embedded media and is therefore too large to email and he knows that offsite locations cannot access the internal network.

How acceptable do you think it would be to risk not delivering the presentation in this situation? [5-point Likert from Not Acceptable at All to Very Acceptable]

If you were Jason, what would you do in these circumstances?

**Behaviour Type 1:** Upload the files to a public online data storage service and recover them at the client's site.

**Behaviour Type 2:** Take the required data on an unencrypted USB stick—you have one to hand.

**Behaviour Type 3:** An employee of the client has been visiting CompA and is due to travel back with you. Use the available unencrypted stick to put a copy of the data onto their laptop and ask them to take it to the client's site.

**Behaviour Type 4:** Borrow an encrypted stick from a colleague. You would have to also make a note of their password so you can access the data at the client's site. The colleague had asked that you do not share / erase the confidential data already on the stick.

### A.3.4 Scenario D (Attitude): Tailgating

Agnes works for CompA as a Customer Account Manager and often has meetings on site with external visitors.

Agnes is aware that visitors need to be supervised at all times and security / reception are made aware of all visitors. She therefore personally receives visitors and allows them entry/exit through the barrier door which requires an ID pass and further supervises them whilst they are on site.

Whilst collecting visitors she often sees unsupervised people without a visible visitor's badge waiting near the barrier door and occasionally 'tailgate' (follow closely behind another person) to get into the main building. Although Agnes appreciates that this is a security risk, she is also aware that this is a common occurrence which is normally overlooked.

If you were Agnes, what would you do in these circumstances?

**Attitude Type 1:** Assume the people have access and have been checked by the reception staff and continue with your work so as not to disrupt their work or yours.

**Attitude Type 2:** Notify security that you have observed visitors tailgating past the barrier.

**Attitude Type 3:** Confront the people and then report their names to either your manager or security.

**Attitude Type 4:** Confront the people you see tailgating directly and ask them to show you some ID (if they are not known to you) and supervise them back to reception.

#### A.3.5 Scenario E (Attitude): Document Control

Anne works for CompA as a Compliance Officer and is responsible for managing and handling documents containing sensitive Critical National Infrastructure (CNI) information. Only certain people can access and use this information, provided they have gone through the appropriate vetting.

Anne recently received an angry phone call from Bob (Senior Manager) who is not happy that his staff no longer have access to important information they need. Anne explains the vetting procedure to him, but he is still not happy and asks to speak to her boss Cyril, who supports Anne and tells Bob that his staff require clearance to access the documents

A couple of weeks later Anne reviews the access logs to the documents, and notices that Dave (who works for Bob) has been accessing a large number of CNI documents. Anne looks into this further and finds that Dave is widely sharing this information to others in his team, none of whom have been through the vetting and managing of privileged information training.

If you were Anne, what would you do in these circumstances?

**Attitude Type 1:** Do nothing—If something goes wrong, Bob has to deal with it as he is the Senior Manager in charge of the department that is sharing the information.

**Attitude Type 2:** Report your observations to Cyril, and urge him to tell Bob formally that this is not acceptable.

**Attitude Type 3:** Initiate an audit of Bob's Department to attempt to track the use and distribution of the CNI documents.

**Attitude Type 4:** Send Dave an informal email, reminding him that sharing CNI documents with non-cleared employees is not allowed.

### A.3.6 Scenario F (Attitude): Information Disposal

James works for CompA as a Senior Contracts Manager and regularly reviews confidential contracts, which contain sensitive commercial information and customer data. He prefers to review documentation in paper form, so he often prints out confidential documents to read and make notes on whilst travelling on the train to/from home.

When he is finished with these documents, as an environmentally conscious person, he places them in the recycling bin.

At home there is no secure disposal available so he uses the standard recycling service. The risk of 'dumpster diving' (people stealing documents from rubbish bins) has been communicated in a recent Internal Awareness Campaign. It specifically recommends disposing of confidential information in the confidential bin but James feels that this is overly cautious and does not always use the Confidential but the normal recycle bin as he thinks the paper will be destroyed when it is recycled anyway, so there is no need to be concerned.

What is your attitude to James' working practises?

**Attitude Type 1:** James is right to work in the way that suits him best—without access to the company systems even if someone did get hold of a few bits of information they couldn't damage CompA anyway.

**Attitude Type 2:** James is being totally reckless with customer's information—the major threat caused by hard copies is to the customers via identity theft and he should stop printing out work unless it is absolutely necessary.

**Attitude Type 3:** James should ensure any paper copies he makes are disposed of specifically in a confidential recycle bin to ensure they are securely shredded once he has finished with them—hard copies are a major source of information leaks.

**Attitude Type 4:** James' working practises are acceptable; recycling the paper is good for the environment and destroys any sensitive information at the same time.



### A.3.7 Scenario G (Behaviour): Backing Up Information

Emilia works for CompA as a Finance Analyst and is a very conscientious individual who occasionally works from home in the evening to catch up on things she could not complete during the day.

Emilia normally uses the train to get home. She chooses to leave her laptop as she has recently had her laptop stolen when travelling home from work.

Emilia keeps a backup of all her work files on her personal computer so she can access files without having to connect to the CompA system as her home network connection is not always reliable. She knows this is against company policy, but she lives in a safe neighbourhood and does not consider this to be a great security risk. In order to transfer files to her home computer she uses a variety of methods.

How acceptable do you think it would be to risk not completing her work tasks due to not having the right information at home in this situation? [5-point Likert from Not Acceptable at All to Very Acceptable]

Assuming you had some vital work to complete at home, if you were Emilia, how would you create a backup of your work for home use?

**Behaviour Type 1:** Email the files to your personal email account from your work account and download them at home.

**Behaviour Type 2:** Use one of your USB sticks to carry your current work with you on the train.

**Behaviour Type 3:** Use an online storage service such as Dropbox as an interim transfer location, deleting the files once you have made a local copy.

**Behaviour Type 4:** Log in to the company VPN and make local copies via that connection.

### A.3.8 Scenario H (Attitude): External Threats

Andrew works at a CompA site, and he walks there each morning from the local station. One morning he notices a blue van parked outside the entrance gates. He thinks he has seen the van parked in the same spot several times before over the last couple of weeks.

Andrew becomes suspicious so he notes down the van details so he can check again if it is the same van.

A few days later, Andrew notices the same van parked in the same location. As he passes the van he observes two individuals, one of whom appears to be taking

pictures of the building/people around the building. As soon as the individuals see Andrew, the van pulls away in a hurry.

If you were Andrew, what would you do in these circumstances?

**Attitude Type 1:** Put it out of your mind; you have seen the van several times and nothing has happened at the site so it probably isn't a threat.

**Attitude Type 2:** Report your suspicions directly to security so they can take the appropriate action.

**Attitude Type 3:** Report the incident to your line manager, it is better to report such incidents even if nothing has happened.

**Attitude Type 4:** Do nothing now but keep an eye out for the van in the future to confirm his suspicions. If it shows up again then report the incident.

### A.3.9 Scenario I (Behaviour): Information Requests

Mohammed is a Contract Support Assistant at CompA who manages 3rd-party contracts. One afternoon, he receives a phone call from Alison who used to work with him at CompA but now works for one of CompA's trusted 3rd-party companies. She asks Mohammed for some commercially sensitive information that is not publicly available through the company's web site.

While the company she works for is allowed access to the information, Mohammed is aware that there is a procedure 3rd-parties need to go through to obtain that information. Mohammed politely refuses the request and reminds Alison of the procedure she should follow. Alison now becomes very persistent and reminds Mohammed that they used to be colleagues as well as mentioning the names of several senior people in both companies, saying they will be extremely unhappy if she does not get this information that day. She further says she is still in contact with his line manager and will explain everything to him later, so Mohammed should be ok with providing this information today.

How acceptable do you think it would be for Mohammed to not share the information with Alison at this stage? [5-point Likert from Not Acceptable at All to Very Acceptable]

Assuming Mohammed decides to share information with Alison at this stage, if you were Mohammed what would you do in these circumstances?

**Behaviour Type 1:** Accede to the request for information to ensure that the senior personnel are satisfied and Alison's productivity isn't hampered.

**Behaviour Type 2:** Ask Alison specifically which pieces of information she needs and send through a redacted or edited version of the documents.

**Behaviour Type 3:** Send Alison the information she requested but immediately inform your line manager of the call and that information has been provided.

**Behaviour Type 4:** Send the information through but password protect the file and wait until you have spoken to your line manager before releasing the password to Alison.

### A.3.10 Scenario J (Attitude): Working Practises

Sanjeeta has worked with Kevin at CompA for a number of years. Kevin has always been an effective member of the team, but is known for ‘having his own way of getting things done’. A few months ago Kevin left CompA to work for one of CompA’s Service Partners. They still maintain a close working relationship and are located at the same site. Recently Sanjeeta noticed that several confidential documents/records were missing and there was no audit trail of who had used them last.

Sanjeeta then recalls that Kevin had accessed the documents to resolve a query associated with a project he had recently been working on, so she decides to ask Kevin about the missing documents next time she saw him.

When asked about the missing documents, Kevin becomes very defensive and objects to being challenged, telling Sanjeeta that she should “stick to her own work and stay out of mine”. Sanjeeta was very taken aback by this response.

If you were Sanjeeta, what would you do in these circumstances?

**Attitude Type 1:** Do nothing, Kevin’s working practises have always been eccentric and this seems to be no more than a product of his usual attitude.

**Attitude Type 2:** Discuss Kevin’s behaviour with the department manager—it isn’t acceptable for an individual in the department to have their own methods that conflict with the company best practice and policy.

**Attitude Type 3:** Call the Business Conduct helpline and make a report about Kevin’s behaviour—it is suspicious that there appears to be no proper audit of his work.

**Attitude Type 4:** Accommodate Kevin’s work practises by adjusting your own, it will be easier and more productive for you both.

## A.4 Company B Behaviour and Attitude scenarios

Below, we provide the texts of the behaviour and attitude scenarios used in the study. Labels are included to indicate which option related to which behaviour type and maturity level, but these were not displayed to participants in the study and the order of the options was randomised as well.

### A.4.1 Scenario A (Attitude): ID Badges

Jemima is a member of the Operations team working in Location 1. While sat working at her desk, she notices someone she doesn't recognise walk past without a visible ID badge. This prompts her to do one of the following:

- Level 2:** Nothing, the security badges are only used for accessing the building and once you are in serve no other real purpose.
- Level 3:** Nothing, although security badges are meant to be visible at all times it is a formality and it is the job of the security guards to check not hers.
- Level 4:** Make sure that her own ID badge is visible, seeing someone without theirs reminds her that she should have hers on display.
- Level 5:** Go and talk to the person and ask if they have a badge. If they have, remind them to have it on display, if not then politely escort them to security.

### A.4.2 Scenario B (Attitude): Clear Desk Policy

When leaving his desk to go for lunch with some colleagues Darren, a member of the HR team, notices that one of them has left his screen unlocked. The rest of the people he is with don't seem to have noticed, or seem to be OK with leaving it as it is. Darren got into the habit of locking his screen some years ago while working in a different company. As his colleagues start to walk away he decides to:

- Level 2:** Do nothing, there is no risk here as no-one could get into the office without passing through security. The screen locks are there just as a formality.
- Level 3:** Do nothing, the screen will automatically lock after a few minutes and this will keep things secure.
- Level 4:** Lock the screen himself.
- Level 5:** Quickly find out whose desk it is from the group and ask them to lock it before they leave for lunch.

### A.4.3 Scenario C (Behaviour): Password Manager

Hina, a member of the Operations division, has recently been required as part of her job to use a new piece of software about once a week. This requires her to log in to the service using a new username and password combination. Unfortunately the password manager does not work correctly with this new software and fails to store or enter her password. Because of the lack of support Hina is worried about being able to use the service as she struggles to remember infrequently used passwords.

Assuming that Hina decides to continue using the service without the support of the password manager, if you were Hina, what would you do in these circumstances?

**Individualist:** Store the password using a method of your own devising—you can be trusted to keep it safe.

**Egalitarian:** Share your password with a trusted member of your working group so that if you forget it they can remind you.

**Hierarchist:** Stop trying to remember the password and just use the password reset feature to generate a new password each time you need to use the service.

**Fatalist:** Re-use a password from another service that you have committed to memory.

### A.4.4 Scenario D (Behaviour): VPN

Robert, an analyst in the Operations team, has a set of logs from secure company hardware that he needs to upload to the manufacturer's website for analysis. He is working from home and unfortunately while connected to the VPN, he is unable to utilise the upload function on the manufacturer's site. It is necessary that the logs are analysed each day so he cannot wait until he is next in the office if he is to successfully complete this task.

Assuming that Robert decides to upload the logs via a different method, if you were Robert, what would you do under these circumstances?

**Individualist:** Make a local copy of the logs, disconnect from the VPN and upload the logs over your home connection.

**Egalitarian:** Give the password to the server to a trusted colleague not working from home and ask them to download the logs from the server before uploading to them to the manufacturer.

**Hierarchist:** Email the logs directly to the manufacturer's customer support email, and ask them to conduct the analysis and send the file back.

**Fatalist:** Email the logs to a colleague not working from home and see if they can upload the logs via a direct LAN connection.

#### A.4.5 Scenario E (Attitude): Tailgating

Jessica is heading toward an access controlled entry door and notices a man she does not recognise gain entry by following close behind someone else who had tagged in at the door. The two men are walking close together although they do not appear to obviously be in conversation. The second man is holding a cup of coffee in one hand and his laptop in the other. His ID badge is not immediately visible. Jessica decides to:

- Level 2:** Return to her desk, she sees this sort of thing quite regularly and it is probably because his hands were full that he did not swipe through himself.
- Level 3:** Do nothing, if he is up to some mischief the security guards will catch him later on.
- Level 4:** Find a security guard at one of the manned turnstiles and tell them what happened.
- Level 5:** Follow the man and ask to see his ID badge.

#### A.4.6 Scenario F (Behaviour): File Storage

Concerned about the safety of his current work, Shamal decides to back up his data, some of which is confidential. As he uses his own laptop under the ‘bring your own device’ scheme, he usually stores all his work on his drive on the central server but he wants to have a second copy just in case something happens or he loses connectivity to the company network. He thought about using one of the common drives but none of the ones he regularly uses have sufficient space.

- Individualist:** Create a local copy on the hard drive of your BYOD laptop, it is the only machine you work on so you know it will be safe and this ensures you will always have access to it if needed.
- Egalitarian:** Use a common drive that you used for an old project and still have access to, as your credentials were never revoked. It has enough space although you do not know who manages it now.
- Hierarchist:** Use an online service, such as Dropbox, to store the data as it is more under your control.
- Fatalist:** Back your work up onto a USB stick—you have ordered an encrypted one but while you wait for it to arrive you use a personal stick you have to hand.

#### A.4.7 Scenario G (Attitude): Secure Disposal

John works as a Sales Advisor in a company store in London. During a busy period of the day he notices that a customer, served by one of his colleagues, has left their paperwork behind. John's colleague grabs the paperwork and throws it into a wastepaper bin under the desk. Seeing this John decides to:

- Level 2:** Carry on serving customers in the store, all the rubbish will be thrown out at the end of the day anyway so it is no big deal, and using the shredder in the back area, locked by a keypad, is inconvenient when the bin is right there.
- Level 3:** Make a note to check with his manager what the appropriate action would be, as it has been some time since he took the Data Protection training module and he cannot clearly remember the details.
- Level 4:** Go and grab the paperwork out of the bin when he has a spare moment and take it to the shredder in the back of the store.
- Level 5:** Go over immediately and ask his colleague to take the paperwork out of the bin and put it in the shredder, having documents lying around exposes both the store and the customer to the risk of identity theft.

#### A.4.8 Scenario H (Behaviour): Credit Check

Karina works as a Sales Assistant in a company store. Her manager has asked her to increase her sales, in order to meet the store's monthly target. In her experience, customers can be put off by the need for credit and ID checks, and sometimes fail them altogether. She knows of a few unofficial ways of making the checks seem less of a problem, or to increase the chance of customers passing them.

- Individualist:** Attempt multiple credit checks in quick succession in order to try to figure out which details are causing the problem and amend them.
- Egalitarian:** Give information about the credit check to a few of your personal contacts so that they can prime potential customers on what they need to do to beat the system before referring them to the store.
- Hierarchist:** Use your employee discount to offer the customer a more attractive deal.
- Fatalist:** Give the benefit of the doubt when encountering IDs with indicators of possible fraud, such as dates of birth that do not seem to align with the apparent age of the customer, or addresses in different cities.

## A.5 Maturity model

This model expresses the maturity of the security culture within an organisation in terms of how aligned with the policy employee behaviour is, and also how integrated the policy is with the primary business process of the organisation. Most critically the model does not represent a checklist of required behaviours for employees, but aims to reinforce the synergy and co-operation required between employer and employee to deliver effective security. As such it is not possible to reach the highest levels of the model in an environment with an inefficient or poorly implemented policy that is in conflict with the primary process of the organisation. Thus the model is capable of guiding change both for the organisation and the individuals that work for it.

This model is based on the Carnegie Mellon Capability Maturity Model (Paulk 2002). This model expresses the degree of formality associated with various processes. What we need from our security behaviour model is a characterisation of what represents effective employee security behaviour, as observed by the organisation. This will then act as a scale against which progress can be measured, as well as a tool for identifying the current state of security behaviour. The CMM consists of five levels, moving from unplanned/unmanaged through a managed state to one of optimisation through incremental innovation. These levels are listed below with definitions for reference.

### **Level 1: Initial (Chaotic)**

It is characteristic of processes at this level that they are (typically) undocumented and in a state of dynamic change, tending to be driven in an ad hoc, uncontrolled and reactive manner by users or events. This provides a chaotic or unstable environment for the processes.

### **Level 2: Repeatable**

It is characteristic of processes at this level that some processes are repeatable, possibly with consistent results. Process discipline is unlikely to be rigorous, but where it exists it may help to ensure that existing processes are maintained during times of stress.

### **Level 3: Defined**

It is characteristic of processes at this level that there are sets of defined and documented standard processes established and subject to some degree of improvement



over time. These standard processes are in place (i.e., they are the AS-IS processes) and used to establish consistency of process performance across the organisation.

#### **Level 4: Managed**

It is characteristic of processes at this level that, using process metrics, management can effectively control the AS-IS process (e.g., for software development). In particular, management can identify ways to adjust and adapt the process to particular projects without measurable losses of quality or deviations from specifications. Process Capability is established from this level.

#### **Level 5: Optimizing**

It is a characteristic of processes at this level that the focus is on continually improving process performance through both incremental and innovative technological changes/improvements.

When considering a Security Behaviour version of this model we must consider how to convert these organisational indicators to indicators of personal behaviour. One approach is to consider how the individual is managing or motivating their own behaviour: what factors they are considering when planning their security actions. At the highest level, they will be actively working toward an improved and improving security culture. At the lower levels employees will be following the policy by rote (possibly reluctantly, ineffectively or incompletely) or simply taking actions as they see fit, based on their own internal security model with no input from the organisation. The following levels represent this range of behaviours.

#### **Level 1: Uninfluenced**

At this level, user behaviour is mediated only by their own knowledge, instincts, goals and tasks. Their actions will reflect only the needs of their primary task and will only deviate from that where their internal security schema conflicts with those actions. While some members of the organisation may be sufficiently knowledgeable to act securely it is expected that employees at this level will introduce a range of vulnerabilities into the system. In practice this level can only exist where employees are working on non-organisational systems, as even the act of logging in to a managed network means that organisational security is exerting an influence.

### **Level 2: Technically Controlled**

Employees at this level act as in level 1 except where technical controls exist that enforce policy on a case-by-case basis. Technically controlled employees will follow their own security rules except where they must use organisational systems in the execution of their primary task, and those systems enforce policy at the software or hardware level. Realistically, this is the lowest practical level that employees working in an office environment could function at.

### **Level 3: Ad-hoc Knowledge and Application**

Employees at level 3 follow policy without necessarily a deep knowledge of what it contains. Their security knowledge comes from the ‘best practise’ or habits associated with their work environment, rather than from being aware of, and understanding, organisational policy.

### **Level 4: Policy Compliant**

Level 4 behaviour demonstrates knowledge and understanding of the policy, and compliance with it, even in situations where the local work environment may include the use of workarounds and frequently made excuses. At Level 4, employees can be considered to be useful role models and guides for security culture within the organisation.

### **Level 5: Active Approach to Security**

At Level 5, employees take an active role in the promotion and advancement of security culture within the organisation. They serve not just the letter of the policy but the intent as well and will challenge breaches at their level appropriately. They see security as a valuable part of the function of the organisation and have internalised this motivation. Level 5 employees are not security zealots, but rather understand the need to balance the security and business processes and champion the cause of security intelligently and effectively.

## SANS analysis

### B.1 Analysis, 2016

I have focused on four specific questions of the survey. These are

Q14: How would you classify the maturity of your organisation's security awareness program?

Q15: What are you measuring in your metrics program?

Q16: What is your organisation's approach when people fail to exhibit the correct secure behaviors?

Q17: What is the single, biggest challenge you are facing with your security awareness program?

Questions 14 and 16 are single choice questions, question 15 is multiple choice with an *other* option and question 17 is open entirely.

Response	Frequency
Non-existent	23
Promoting awareness and behaviour change	161
Compliance focused	78
Robust metrics framework	7
Long-term sustainment and culture change	39

**Table B.1: Frequency of coded responses to 2016's Q14** *'How would you classify the maturity of your organisation's security awareness program?'*

Table B.1 lists the pure quantitative frequencies of the responses to the Q14. The answers are in strictly increasing order of awareness maturity.

Response	Frequency
Phishing assessments	159
Number of security violations	121
Number of infected devices	118
We do not have any metrics	95
Number of times data has been accidentally exposed	89
Number of lost devices	80
Training attendance	20
How much money our awareness program is saving us	19
Number of incidents reported	17
Intranet site traffic	12
Tests	10
Number of compromised accounts	4
Amount of outreach work	3
Number of security advocates	2
Incident responds time	1
Non-compliance	1
Customer complains	1

**Table B.2: Frequency of coded responses to 2016’s Q15 ‘What are you measuring in your metrics program?’**

Table B.2 lists the frequency of responses to question 15. There was a substantial amount of free-text responses which have been manually grouped together into the groups

- Amount of outreach work,
- Customer complains,
- Incident responds time,
- Intranet site traffic,
- Non-compliance,
- Number of compromised accounts,
- Number of incidents reported,
- Number of security advocates,
- Tests,
- Training attendance.

It is interesting to note that a number of these groups have significant frequencies, and are certainly under reported due to the nature of the ”other” answer group in questionnaires.

Table B.3 plots the responses to Question 14 (on the horizontal axis) against the response to question 15 (along the vertical axis). The number in each cell represents the percentage of responses of all participants which have chosen the combination of answers from Questions 14 and 15 for the corresponding cell. For example, 31.5% of participants both report to conduct *Phishing assessments* as well as their organisation is *Promoting awareness and behaviour change*. As question 15

	Non-existent	Compliance focused	Promoting awareness and behaviour change	Long-term sustainment and culture change	Robust metrics framework
We do not have any metrics	5.0%	11.7%	10.1%	1.3%	0.3%
How much money our awareness program is saving us	0.0%	1.6%	0.9%	2.5%	0.6%
Phishing assessments	1.3%	6.9%	31.5%	7.9%	1.9%
Number of infected devices	0.6%	7.9%	20.5%	6.0%	1.6%
Number of lost devices	0.3%	4.1%	15.8%	4.1%	0.9%
Number of security violations	0.6%	7.9%	21.8%	6.0%	1.3%
Number of times data has been accidentally exposed	0.9%	5.0%	14.5%	6.0%	1.3%
Other	1.3%	3.8%	10.4%	2.2%	0.6%
Amount of outreach work	0.0%	0.0%	0.3%	0.6%	0.0%
Customer complains	0.0%	0.0%	0.0%	0.3%	0.0%
Incident responds time	0.0%	0.0%	0.0%	0.3%	0.0%
Intranet site traffic	0.0%	0.0%	3.2%	0.6%	0.0%
Non-compliance	0.0%	0.3%	0.0%	0.0%	0.0%
Number of compromised accounts	0.0%	0.9%	0.3%	0.0%	0.0%
Number of incidents reported	0.0%	0.9%	3.5%	0.9%	0.0%
Number of security advocates	0.0%	0.0%	0.3%	0.3%	0.0%
Tests	0.3%	0.6%	1.9%	0.3%	0.0%
Training attendance	0.3%	0.6%	4.1%	0.9%	0.3%

**Table B.3: Frequency analysis of 2016's Q14 vs Q15** Q14: 'How would you classify the maturity of your organisation's security awareness program?' vs. Q15: 'What are you measuring in your metrics program?'

is a multiple choice question the percentages do not add up in a meaning full manor. The background colours of the cells rank the values globally. Note that the *other* row represents the sum of all textual responses.

The data is most easily interpreted by taking each column in isolation and establishing the most common responses to question 15. This is straight forward due to the colour ranking.

The data gives rise to a number of trends: responses with awareness metrics of *non-existent* or *compliance focused* do mostly undertake no metrics to establish the impact of their awareness program. As we progress right this shifts: three times as many respondents reported to conduct *phishing assignments* alone than not having any metrics at all.

*Phishing assessments* are by far the most popular method to measure awareness success. As the awareness programs become more mature, automatic measures such as *number of infected devices* and *number of security violation* gain shares. Lastly and suprisingly, *money saved* due to the awareness program is rarely measured.

Response	Frequency
Punitive	14
Progressive	157
Flexible	146

**Table B.4: Frequency of coded responses to 2016’s Q16** ‘*What is your organisation’s approach when people fail to exhibit the correct secure behaviors?*’

Question 16 splits the population on their responses to incorrect behaviour. The potential answers are here:

**Punitive** Security policies and secure behaviors are the ”law of the land” and people who violate them should be punished.

**Progressive** Security policies and secure behaviors are enforced on a sliding scale. First time offenders are not punished, but people are punished the more times they fail.

**Flexible** Security policies are more like guidelines and should not get in the way of effectiveness.

While the arguably ideal awareness program should not punish incorrect behaviour, the majority of responses claim a *progressive* response (Table B.4).

Table B.5 supports this hypothesis: as more metrics become available, the focus shifts away from a *flexible* approach to incorrect behaviour. Without evaluation metrics a punitive approach is not possible (and indeed only 1.6% of respondents claimed to have no metrics but a punitive response to incorrect behaviour). Interestingly none of the manually grouped evaluation metrics coincide with a punitive response. This may well be a selection bias on the survey’s participants.

Question 17 was entirely open. I coded the 258 responses in three iterations into 12 groups:

**Resources:** A shortage of technical resources as well as money,

**Adoption:** Users fail to change their behaviour through training. One quote from the participants describes this group very well: ‘*Users still clicking on links & opening malicious attachments despite attending training sessions, sometimes only hours before*’.

**Support from Management:** Superiors do not see the necessity for the awareness campaigns and/or fail to collaborate and facilitate.

**End-user support:** The people supposed to take part in the awareness campaign do not feel it is necessary and are generally unwilling to take part.

**Finding time to take part:** End-users struggle to find time to take part in awareness campaigns, they are always a low priority.

	Punitive	Progressive	Flexible
We do not have any metrics	0.6%	9.5%	19.9%
How much money our awareness program is saving us	0.6%	3.5%	1.9%
Phishing assessments	2.8%	29.7%	17.7%
Number of infected devices	1.9%	22.1%	13.2%
Number of lost devices	1.3%	15.1%	8.8%
Number of security violations	2.8%	24.0%	11.4%
Number of times data has been accidentally exposed	1.9%	18.0%	8.2%
Other	0.0%	10.1%	8.2%
Amount of outreach work	0.0%	0.6%	0.3%
Customer complains	0.0%	0.3%	0.0%
Incident responds time	0.0%	0.0%	0.3%
Intranet site traffic	0.0%	2.2%	1.6%
Non-compliance	0.0%	0.3%	0.0%
Number of compromised accounts	0.0%	0.6%	0.6%
Number of incidents reported	0.0%	2.8%	2.5%
Number of security advocates	0.0%	0.3%	0.3%
Tests	0.0%	2.2%	0.9%
Training attendance	0.0%	3.8%	2.5%

**Table B.5: Frequency analysis of 2016’s Q16 vs Q15** Q16: ‘What is your organisation’s approach when people fail to exhibit the correct secure behaviors?’ vs. Q15: ‘What are you measuring in your metrics program?’

Response	Frequency
Resources	50
Adoption	48
Support from Management	47
End-user support	27
Finding time to take part	24
Content	23
Not enough awareness staff	22
Non-mandatory	9
Legal department	4
Non-punishable	2
Translations	1
Metrics	1

**Table B.6: Frequency of answers to 2016’s Q17** ‘What is the single, biggest challenge you are facing with your security awareness program?’

**Content:** Participants struggle to create appropriate content.

**Not enough awareness staff:** Self-explanatory.

**Non-mandatory:** Taking part in awareness campaigns is not-mandatory. This blocks progress as take-up is low.

**Legal department:** The legal department is blocking parts of the awareness campaign.

**Non-punishable:** End-users cannot be punished for failing awareness evaluations.

**Translations:** The awareness campaign needs to be translated into many languages.

**Metrics:** Correct metrics are difficult to establish.

Table B.6 lists the frequency of each of these categories. It should be noted that *Resources* and *Not enough awareness staff* are linked, as well as *Adoption* and *End-user support*: the first is post training and the latter prior training. Contextually there are dependencies between the lack of *support from management*, *end-user support* and *finding time to take part*: if one’s boss is unwilling to provide time for an awareness campaign and does not support it, end-users are unlikely to think highly of the awareness campaign either.

	Non-existent	Compliance focused	Promoting awareness and behaviour change	Long-term sustainment and culture change	Robust metrics framework
Adoption	1.6%	3.2%	8.5%	1.9%	0.0%
Content	0.0%	3.5%	2.5%	1.3%	0.0%
End-user support	0.6%	1.3%	4.4%	1.6%	0.0%
Finding time to take part	0.9%	1.6%	3.8%	0.6%	0.6%
Legal department	0.0%	0.0%	0.9%	0.3%	0.0%
Metrics	0.0%	0.0%	0.0%	0.3%	0.0%
Non-mandatory	0.6%	0.6%	0.9%	0.3%	0.3%
Non-punishable	0.0%	0.0%	0.3%	0.3%	0.0%
Not enough awareness staff	0.3%	1.3%	4.4%	0.9%	0.0%
Resources	0.6%	2.8%	9.8%	1.6%	0.3%
Support from Management	0.9%	6.3%	6.0%	1.3%	0.0%
Translations	0.0%	0.0%	0.0%	0.3%	0.0%

**Table B.7: Frequency analysis of 2016’s Q14 vs Q17** Q14: ‘How would you classify the maturity of your organisation’s security awareness program?’ vs. Q17: ‘What is the single, biggest challenge you are facing with your security awareness program?’

Next, Table B.7 contrasts the single biggest challenge to the maturity of the awareness program. Again there are some note worthy trends: as the awareness program matures, *support from management* cedes to be an issue. *Resources* as well as *not enough awareness staff* display similar trends. Interesting are also the behaviour of *adoption* and *end-user support*: while significant issues for maturity levels 2,3 and 4, for the case of the *robust metrics framework* not a single response was recorded.



	Adoption	Content	End-user support	Finding time to take part	Not enough awareness staff	Resources	Support from Management
We do not have any metrics	3.8%	1.6%	1.9%	1.6%	2.2%	4.4%	6.3%
Money saved	1.6%	0.9%	0.3%	0.0%	0.3%	0.0%	0.6%
Phishing assessments	8.5%	3.5%	4.4%	4.7%	3.2%	9.8%	4.1%
Number of infected devices	7.3%	2.8%	2.2%	3.2%	3.2%	4.1%	4.4%
Number of lost devices	3.8%	1.9%	1.6%	1.9%	2.5%	5.0%	2.2%
Number of security violations	6.9%	2.8%	3.5%	2.5%	1.3%	7.3%	4.4%
Number of times data accidentally exposed	5.7%	1.9%	2.8%	1.6%	1.3%	4.7%	2.8%
Other	1.6%	1.9%	2.5%	1.6%	0.9%	3.5%	3.2%
Intranet site traffic	0.3%	0.6%	0.3%	0.0%	0.3%	1.6%	0.3%
Number of incidents reported	0.6%	0.3%	1.3%	0.0%	0.3%	1.6%	0.6%
Training attendance	0.3%	0.9%	0.6%	0.0%	0.3%	2.2%	1.3%

**Table B.8: Frequency analysis of 2016’s Q17 vs Q15** Q17: ‘What is the single, biggest challenge you are facing with your security awareness program?’ vs. Q15: ‘What are you measuring in your metrics program?’ for answers with > 10 responses.

Finally, [Table B.8](#) contrasts the awareness metrics with the single biggest challenges. For readability, all categories with less than 11 responses have been omitted. There are a number of interesting correlations: *We do not have any metrics* coincides with a lack of *support from management*, in fact a lack of metrics in general is the strongest contributing factor to a lack of support from management. *Phishing assessments* are the strongest contributing factor to all other single biggest challenges. Other metrics have significantly less strong contributions.

### B.1.1 Conclusions

There are a number of conclusions that can be drawn from the analysis.

#### Awareness professionals use metrics to punish

Rather than using additional metrics to improve and tailor education and create proactive defenses, the awareness professionals are inclined to use more detailed results to punish users who are unwilling or unable to change. This is surprisingly linked to the maturity of the awareness model.

#### Metrics are widespread, but do not measure monetary value

The vast majority of respondents noted that they were using metrics to evaluate their awareness campaign. Phishing campaigns are most widely used, but

virtually no awareness program attempts to measure its financial impact.

### **Little support from management**

The organisations management often does not support the awareness campaigns. This has the consequence of little end-user support and low uptake of awareness training. Metrics are extremely important to the management, but there does not seem to be a clear preference of which metrics are preferred.

### **End-users keep on making the same mistakes**

The adoption of awareness training is low; the second most frequently mentioned single biggest challenge is low adoption.

### **Too few resources available**

Especially less mature awareness campaigns struggle with too few monetary as well as human resources. This also extends to technical support from different parts of the organisation.

Awareness campaign in their current format suffer from a number of self reinforcing problems. Without strong metrics support from management is low, which influences participation in awareness training negatively which in turn weakens metrics. Low adoption of awareness make awareness campaigns appear to be a waste of time, which limits resources available which causes poor awareness campaigns.

## B.2 Analysis, 2017

There were a number of questions that appeared in both this years and last years survey. The questions of interests are:

Q14: How would you classify the maturity of your organisation’s security awareness program?

Q16: What is your organisation’s approach when people fail to exhibit the correct secure behaviors?

Q17: What is the single, biggest challenge you are facing with your security awareness program?

Response	2017
Promoting awareness and behavior change	445
Compliance focused	219
Long-term sustainment and culture change	80
Non-existent	60
Robust metrics framework	7

**Table B.9: Frequency of answers to 2017’s Q14 ‘How would you classify the maturity of your current security awareness program, based on the Security Awareness Maturity Model?’**

Response	2017
Flexible: security policies are more like guidelines and should not get in the way of effectiveness.	394
Progressive: security policies and secure behaviors are enforced on a sliding scale. First time offenders are not punished, but people are punished the more times they fail.	382
Punitive: security policies and secure behaviors are the ‘flaw of the land’ and people who violate them should be punished.	35

**Table B.10: Frequency of answers to 2017’s Q16 ‘What is your organisation’s approach when people fail to exhibit the correct secure behaviors?’**

Tables B.9 to B.11 are the 2017 response frequencies corresponding to the 2016 versions in Tables B.1, B.4 and B.6.

As in the previous year, Q17 remained an open question. Again, in a multi-round coding, I reduced the 811 responses down to the following categories:

**enforcement** punishing non-compliant users

**security fatigue** security fatigue

**skills** not the correct skills in order to carry out awareness

Response	2017
enforcement	15
high level support	127
organisational structure	48
outcome	45
relevance	147
resources	286
security fatigue	9
skills	14
users	53

**Table B.11: Frequency of answers to 2017’s Q17** *‘In as few words as possible, what is the biggest challenge you are facing with your security awareness program?’*

**content** content

**culture** existing culture

**organisational structure** the structure of the organisation is a barrier

**relevance** people don’t care, the

**outcome** achieving something / wanting to achieve too much

**high level support** programs do not get approved

**resources** too few money, staff, material

**users** the stupid user is the barrier

### B.2.1 Memorable quotes

The respondents gave a number of memorable quotes:

*‘the perception that security gets in the way of productivity’*

*‘Getting the organisations attention and not using fear’*

Table B.12 is the 2017 version of Table 5.4 of this years data set. The difference between these two figures is staggering. In 2016, the most commonly selected intersection was a *Progressive* approach combined with a maturity at the level of *Promoting awareness and behaviour change*. This was closely followed by a *Flexible* approach with the same maturity level. In 2017 however this These two options have dropped to 1.7% and 0.1% (from previously 27.4% and 21.1%). Instead the focus on the enforcement side has shifted towards a combination of *Punitive* and *Progressive*. The most common maturity is now *Long-term sustainment and culture change*. Further there has been a sharp increase (from 0% to 15.5% in the combination of a *Non-existent* maturity level with a *Punitive* approach, while the *Flexible* approach for organisations without any maturity of their awareness program as decreased.

Table B.13 is the 2017 version of Table B.7 of this years data set. Comparing these two tables there is also some change to be noticed, albeit less pronounced.

	Non-existent	Compliance focused	Promoting awareness and behaviour change	Long-term sustainment and culture change	Robust metrics framework
Punitive	15.5%	3.1%	5.5%	24.4%	0.0%
Progressive	10.4%	5.7%	1.7%	28.6%	0.7%
Flexible	1.1%	1.1%	0.1%	1.8%	0.1%

**Table B.12: Frequency analysis of 2017’s Q14 vs Q16** Q14: ‘How would you classify the maturity of your current security awareness program, based on the Security Awareness Maturity Model?’ vs. Q16: ‘What is your organisation’s approach when people fail to exhibit the correct secure behaviors?’

	Non-existent	Compliance focused	Promoting awareness and behaviour change	Long-term sustainment and culture change	Robust metrics framework
enforcement	0.5%	0.1%	0.0%	1.2%	0.0%
high level support	5.7%	0.6%	2.2%	7.2%	0.0%
organisational structure	1.6%	0.5%	0.4%	3.5%	0.0%
outcome	1.2%	1.0%	0.1%	3.2%	0.0%
relevance	3.0%	2.1%	0.6%	12.1%	0.4%
resources	9.0%	3.3%	2.6%	20.1%	0.2%
security fatigue	0.0%	0.5%	0.0%	0.6%	0.0%
skills	0.5%	0.2%	0.1%	0.9%	0.0%
users	2.1%	0.7%	0.2%	3.5%	0.0%

**Table B.13: Frequency analysis of 2017’s Q14 vs Q17** Q14: ‘How would you classify the maturity of your current security awareness program, based on the Security Awareness Maturity Model?’ vs. Q17: ‘In as few words as possible, what is the biggest challenge you are facing with your security awareness program?’

While in 2016 *resources* was the most common option for organisations attempting to *Promoting awareness and behaviour change*, *resources* are now the most common biggest challenge across all awareness maturity levels. *Relevance* and *Adoption* as well as *Management support* are still prominent features in both years.

### B.2.2 Conclusions

There has been a marked shift of organisations awareness programs towards a punitive approach when handling people who fail to exhibit the correct secure behaviour. This has not been matched by a shift in awareness maturity, where the dominant approach remains *Promoting awareness and behaviour change*. The number of organisations that support their approach with a *robust metrics framework* has remained constant, and as this years survey response rate has increased dramatically this is a surprising feature.

Many organisation are now punishing their users insecure behaviour without having an metrics framework in place in order to objectively evaluate their performance.