

Persuasive Password Security

Dirk Weirich

A dissertation submitted in partial fulfilment
of the requirements for the degree of

**Doctor of Philosophy
of the
University of London**

Department of Computer Science
University College London

October 2005

ABSTRACT

An organization that gives users access to computing resources via a password mechanism needs to ensure that they perform certain secure behaviours if it wants those resources to be protected adequately. The research problem this thesis seeks to address is the question of how the likelihood of users performing these behaviours can be increased when some of those behaviours can neither be enforced nor monitored adequately.

The primary substantive contribution of the thesis is a grounded theory model of the process users go through when choosing password-related behaviours in the absence of any organizational efforts to influence this choice. The model is subsequently extended to incorporate the effect on user behaviour of password regulations and their associated punishment regimes.

The thesis then presents a discourse-analytic investigation of the interpretative repertoires users draw on to describe aspects of password security, and of the effect of those repertoires on users' password practices. This investigation also shows that users might at times structure their discourse about password security issues in a manner that makes it possible for them to justify malpractice. The use of discourse analysis to investigate these issues is a methodological contribution to the field of human-computer interaction.

The opportunistic use of quantitative data that had been collected prior to a re-conceptualisation of the research approach is used to examine the extent to which users violate password regulations. An analysis of all the qualitative data collected allows a first insight into the specific insecure behaviours that users choose in particular situations.

Persuasive password security, an integration of all these findings into an applicable approach to improving user behaviour, is presented, and specific recommendations on how to improve users' password practices in organizations are made.

ACKNOWLEDGEMENTS

I would like to thank my supervisor, Professor Martina Angela Sasse, for her support and advice throughout the time this research and the subsequent write-up have taken. Thanks are also due to my second supervisor, Professor Mel Slater, and to my industrial supervisor, Dr. Marek Rejman-Greene.

Numerous UCL colleagues have helped during the research process, none more so than Dr. Sacha Brostoff, with whom I collaborated on some of the studies presented in this thesis.

I would like to acknowledge that part of this research was funded by a British Telecom CASE studentship.

Finally, I would like to thank all the participants in the studies for their invaluable contribution.

TABLE OF CONTENTS

ABSTRACT	2
ACKNOWLEDGEMENTS	3
TABLE OF CONTENTS	4
LIST OF FIGURES.....	8
LIST OF TABLES.....	9
1 INTRODUCTION	10
1.1 BACKGROUND TO THE RESEARCH PROBLEM	10
1.2 RESEARCH QUESTIONS.....	11
1.3 RESEARCH APPROACH	13
1.4 RESEARCH SCOPE.....	13
1.5 THESIS CONTRIBUTIONS.....	14
1.6 THESIS STRUCTURE.....	15
1.7 TERMINOLOGY, TYPOGRAPHY AND GENDER-SPECIFIC PRONOUNS	18
2 BACKGROUND	19
2.1 OVERVIEW	19
2.2 HUMAN-COMPUTER INTERACTION (HCI).....	19
2.3 PASSWORD SECURITY	23
2.3.1 <i>Computer security and the importance of access control.....</i>	<i>23</i>
2.3.2 <i>The dominance of password security – and why it will continue.....</i>	<i>24</i>
2.3.3 <i>How password security is implemented.....</i>	<i>25</i>
2.3.4 <i>How password security is breached</i>	<i>26</i>
2.3.4.1 Online guessing.....	27
2.3.4.2 Offline dictionary attacks.....	27
2.3.4.3 Obtaining the password.....	29
2.3.4.3.1 Persuading the user to disclose the password.....	29
2.3.4.3.2 Finding the password	30
2.3.4.3.3 Exploiting the login process.....	30
2.3.5 <i>Interim summary.....</i>	<i>30</i>
2.3.5.1 Users’ tasks.....	30
2.3.5.2 The research problem.....	32
2.4 PREVIOUS RESEARCH ON THE HUMAN ISSUES AFFECTING COMPUTER AND PASSWORD SECURITY	33
2.4.1 <i>Threats in the civil world: (re)discovering the human link</i>	<i>33</i>
2.4.2 <i>A taxonomy of previous research on the human issues involved in computer and password security</i>	<i>34</i>
2.4.3 <i>Studying user behaviour</i>	<i>37</i>
2.4.4 <i>Usability of computer security.....</i>	<i>38</i>
2.4.4.1 Usability of computer security in general.....	38
2.4.4.2 Usability of password security	42
2.4.4.2.1 Password memorability	42
2.4.4.2.2 Identifying factors that affect user behaviour.....	45
2.4.5 <i>User knowledge and skills</i>	<i>49</i>
2.4.6 <i>Motivation.....</i>	<i>50</i>
2.4.7 <i>Investigating the wider organisational context.....</i>	<i>52</i>
2.4.8 <i>Interim summary.....</i>	<i>54</i>
2.4.8.1 Previous research	54
2.4.8.2 The research questions	55
2.5 CHAPTER SUMMARY: THE RESEARCH AGENDA	56

3	METHODOLOGY	58
3.1	OVERVIEW	58
3.2	DATA COLLECTION TECHNIQUES.....	58
3.2.1	<i>Questionnaires</i>	58
3.2.1.1	Justification.....	58
3.2.1.2	Application.....	58
3.2.2	<i>Semi-structured in-depth interviews</i>	59
3.2.2.1	Justification.....	59
3.2.2.2	Application.....	59
3.2.3	<i>Focus groups</i>	60
3.2.3.1	Justification.....	60
3.2.3.2	Application.....	60
3.3	GROUNDING THEORY	61
3.3.1	<i>Justification</i>	61
3.3.2	<i>Application</i>	62
3.4	DISCOURSE ANALYSIS	64
3.4.1	<i>Justification</i>	64
3.4.2	<i>Application</i>	67
3.5	PROTECTION MOTIVATION THEORY	68
3.5.1	<i>Justification</i>	68
3.5.2	<i>Application</i>	70
3.6	CHAPTER SUMMARY: THE RESEARCH APPROACH.....	71
4	RESEARCH CHRONOLOGY	72
4.1	OVERVIEW	72
4.2	STUDY 1: INTERVIEWS	73
4.3	STUDY 2: FIELD TRIAL 1	73
4.4	STUDY 3: FEAR APPEAL SCENARIO.....	76
4.5	STUDY 4: INVESTIGATING A DIFFERENT USER GROUP	77
4.6	STUDY 5: FIELD TRIAL 2	78
4.7	STUDY 6: FIELD TRIAL FOCUS GROUPS.....	79
4.8	STUDY 7: SINGLE-SIGN-ON SCENARIOS	79
4.9	CHAPTER SUMMARY	80
5	PASSWORD PRACTICES AND ATTITUDES TOWARDS PASSWORD SECURITY.....	81
5.1	OVERVIEW	81
5.2	PASSWORD PRACTICES	82
5.2.1	<i>Password choice and password changing</i>	82
5.2.1.1	Choosing an individual password.....	82
5.2.1.2	Choosing multiple passwords that have to be changed regularly	84
5.2.2	<i>Memorizing the password and writing it down</i>	86
5.2.3	<i>Sharing the password</i>	88
5.3	ATTITUDES	90
5.3.1	<i>The risk of malpractice</i>	90
5.3.2	<i>The likelihood of an attack</i>	91
5.3.3	<i>The threat of punishment</i>	92
5.4	CHAPTER SUMMARY	94
6	HOW USERS CHOOSE PASSWORD-RELATED BEHAVIOURS: THE CORE MODEL	95
6.1	OVERVIEW	95
6.2	CHOICE: THE CENTRAL CATEGORY AND THE GROUNDING THEORY MODEL BASED ON IT	97
6.2.1	<i>A grounded theory model of choice</i>	97
6.2.2	<i>The individual elements of the model</i>	101
6.2.2.1	Required security level	101
6.2.2.1.1	Attack likelihood.....	101
6.2.2.1.2	Attack consequences.....	103
6.2.2.2	Provided security level.....	105
6.2.2.3	User cost.....	107
6.2.2.4	User benefit.....	109

6.2.2.5	Personality.....	109
6.3	HABITUAL CHOICES AND THE RECURSIVE APPLICATION OF THE MODEL.....	110
6.4	TESTING THE WATERS.....	111
6.5	CHAPTER SUMMARY.....	112
7	EXTENSIONS TO THE MODEL.....	114
7.1	OVERVIEW.....	114
7.2	EXTENDING THE MODEL TO INCORPORATE THE EFFECT OF PASSWORD REGULATIONS AND THEIR ASSOCIATED PUNISHMENT REGIMES.....	115
7.2.1	<i>Extending the model</i>	115
7.2.2	<i>The additional model elements</i>	119
7.2.2.1	Punishment risk.....	119
7.2.2.1.1	Detection likelihood.....	119
7.2.2.1.2	Severity of expected punishment.....	120
7.2.2.2	Authority personality.....	122
7.3	REMARKS ON SECURITY AWARENESS AND EDUCATION CAMPAIGNS.....	122
7.4	REACHING OUT TO USERS.....	123
7.5	CHAPTER SUMMARY.....	124
8	A DISCOURSE-ANALYTIC STUDY OF USERS' ACCOUNTS OF PASSWORD SECURITY ISSUES.....	125
8.1	OVERVIEW.....	125
8.2	IMPOVERISHED REPERTOIRES USED TO DESCRIBE POTENTIAL ATTACKERS.....	126
8.2.1	<i>Kids</i>	126
8.2.2	<i>Vandals</i>	127
8.2.3	<i>Criminals</i>	128
8.2.4	<i>Vengeful people</i>	128
8.2.5	<i>Disgruntled employees</i>	128
8.2.6	<i>Industrial spies</i>	128
8.2.7	<i>Terrorists</i>	129
8.2.8	<i>Jokers</i>	129
8.3	RESOURCE AND PROBLEM REPERTOIRES.....	129
8.4	QUESTIONING USERS' ACCOUNTS.....	131
8.4.1	<i>Using impoverished repertoires</i>	132
8.4.2	<i>Using problem repertoires</i>	133
8.4.3	<i>Resistance to punishment</i>	134
8.4.3.1	Users' proposals to improve password practices.....	135
8.4.3.2	Punishment.....	136
8.5	CHAPTER SUMMARY.....	137
9	PERSUASIVE PASSWORD SECURITY.....	138
9.1	OVERVIEW.....	138
9.2	RECOMMENDATIONS BASED ON THE CORE MODEL.....	140
9.2.1	<i>Required security level</i>	140
9.2.2	<i>Provided security level</i>	143
9.2.3	<i>User cost</i>	145
9.2.3.1	Primary cost.....	145
9.2.3.2	Secondary cost.....	146
9.2.3.3	Social cost and image cost.....	147
9.2.4	<i>User benefit</i>	148
9.2.4.1	Creating primary benefits.....	148
9.2.4.2	Social benefit and image benefit.....	148
9.2.5	<i>Personality</i>	148
9.2.6	<i>Choosing out of habit</i>	149
9.3	RECOMMENDATIONS BASED ON THE EXTENDED MODEL.....	150
9.3.1	<i>Regulations and their associated punishment regimes</i>	150
9.3.1.1	Punishment risk.....	150
9.3.1.2	Authority personality.....	151
9.3.2	<i>Other possible extensions</i>	152
9.4	GETTING USERS' ATTENTION.....	152
9.5	RECOMMENDATIONS BASED ON THE DISCOURSE-ANALYTIC STUDY.....	153

9.6	PERSUASIVE PASSWORD SECURITY: RECOMMENDATIONS.....	153
9.7	CHAPTER SUMMARY	156
10	CONCLUSIONS.....	158
10.1	THESIS CONTRIBUTIONS.....	158
10.1.1	<i>Overview</i>	158
10.1.2	<i>Substantive contributions</i>	159
10.1.2.1	User behaviour and attitudes.....	159
10.1.2.2	The core model.....	160
10.1.2.3	Extensions to the model	161
10.1.2.4	Interpretative repertoires	161
10.1.2.5	Discursive practices	161
10.1.2.6	Persuasive password security.....	162
10.1.3	<i>Methodological contributions</i>	162
10.1.3.1	Extension of HCI scope	162
10.1.3.2	Fear appeals	162
10.1.3.3	Discourse analysis.....	162
10.2	CRITICAL REVIEW OF THE THESIS.....	163
10.2.1	<i>Research scope</i>	163
10.2.2	<i>Thesis complexity</i>	164
10.2.3	<i>Discourse analysis invalidating grounded theory findings?</i>	164
10.3	FURTHER RESEARCH	165
10.3.1	<i>Testing the model's predictions</i>	165
10.3.2	<i>Expanding the scope of the results</i>	165
10.3.3	<i>Adding further extensions to the model</i>	166
	REFERENCES.....	167
	APPENDIX A: GROUNDED THEORY CATEGORIES AND PROPERTIES IN ALPHABETICAL ORDER WITH EXAMPLE QUOTES.....	173
	APPENDIX B: STATISTICAL TESTS	189
	APPENDIX C: FLYER 1.....	192
	APPENDIX D: WEBSITE.....	193
	APPENDIX E: FLYER 2.....	194
	APPENDIX F: QUESTIONNAIRE.....	195
	APPENDIX G: FLYER 3	203

LIST OF FIGURES

FIGURE 1: POOLS OF BEHAVIOURS USERS CAN CHOOSE FROM IN SPECIFIC SITUATIONS	95
FIGURE 2: THE THREE PROPERTIES OF AVAILABLE BEHAVIOURS.....	98
FIGURE 3: THE STORYLINE OF CHOICE	99
FIGURE 4: THE FACTORS INFLUENCING THE CHOICE PROCESS.....	100
FIGURE 5: HOW THE DIFFERENT ATTACK TYPES ARE CONSTRUCTED ON THE BASIS OF USERS' PERCEPTIONS OF ATTACKERS	101
FIGURE 6: HOW USERS ESTIMATE THE ATTACK LIKELIHOOD.....	103
FIGURE 7: HOW USERS ESTIMATE ATTACK CONSEQUENCES.....	105
FIGURE 8: HOW USERS ESTIMATE THE PROVIDED SECURITY LEVEL.....	107
FIGURE 9 : THE FOUR PROPERTIES OF AVAILABLE BEHAVIOURS.....	116
FIGURE 10: STORYLINE OF THE EXTENDED MODEL.....	117
FIGURE 11: THE FACTORS INFLUENCING THE CHOICE PROCESS IN THE EXTENDED MODEL.....	118
FIGURE 12: HOW USERS ESTIMATE THE DETECTION LIKELIHOOD	120
FIGURE 13: HOW USERS ESTIMATE THE SEVERITY OF EXPECTED PUNISHMENT	122

LIST OF TABLES

TABLE 1: AXIAL CODING EXAMPLE	64
TABLE 2 : HOW RESPONDENTS CHOSE THEIR PASSWORD (N=113).....	84
TABLE 3 : THE RISKS ASSOCIATED WITH BREAKING PASSWORD RULES (ON A SCALE OF 1 (STRONGLY DISAGREE) TO 7 (STRONGLY AGREE)) (N=165)	91
TABLE 4: THE PERCEIVED LIKELIHOOD OF CERTAIN TYPES OF ATTACKS (ON A SCALE OF 1 (VERY UNLIKELY) TO 7 (VERY LIKELY)) (N=165)	91
TABLE 5 : PUNISHMENT EXPECTANCY FOR BREAKING PASSWORD REGULATIONS (N=165).....	92
TABLE 6 : SCENARIO USED IN THE QUESTIONNAIRE.....	92
TABLE 7: FAIRNESS (ON A SCALE OF 1 (VERY UNFAIR) TO 7 (VERY FAIR)) AND LIKELIHOOD (ON A SCALE OF 1 (VERY UNLIKELY) TO 7 (VERY LIKELY)) OF THE FOUR ARGUMENTS (N=165)	93

1 INTRODUCTION

1.1 Background to the research problem

In the past, research on computer security in general, and password mechanisms in particular, has focussed almost entirely on technical issues, such as encryption algorithms or firewalls. This approach has been criticised as early as 1975, when Saltzer & Schroeder (1975) included usability in the stated goals of the secure system. However, it is only in recent years that the security community at large has realised that a significant number of security breaches that are being reported have been enabled and facilitated by user behaviour. The fact that user behaviour can be an important contributor to security breaches is not surprising, since the way in which password mechanisms are set up requires users to perform a number of tasks in order to ensure that the resources in question are protected adequately. The exact tasks users need to perform depend on the specific implementation of the password mechanism that is in place, and on the particular threats that an organization faces, as usually captured in a threat model. Section 2.3.5.1. lists the tasks that were deemed necessary in the organizations that participated in the studies conducted as part of the research presented in this thesis, and introduces the guidelines given to users throughout those studies:

1. Users need to choose cryptographically strong passwords.
2. They need to choose a unique password for each password-protected resource they access.
3. They need to memorise their passwords (as opposed to keeping a physical copy).
4. They must not share their passwords with third parties.
5. They need to change their passwords at regular intervals.

There are numerous reports of users regularly performing behaviours that undermine security (e.g. Schneier (2000), Winkler (1997)): they choose weak passwords, use them for several systems, write them down, and share them easily. This is exploited by attackers of computing resources. Kevin Mitnick, arguably the world's most famous hacker, testified to the US Senate committee that he had obtained 9 out of 10 passwords by tricking users, rather than through cracking. In his new role as security evangelist, he emphasises that

“The human side of computer security is easily exploited and constantly overlooked. Companies spend millions of dollars on firewalls, encryption and secure access devices, and it's money wasted, because none of these measures address the weakest link in the security chain.” (Poulsen (2000)).

It is now accepted by the security community that

“... security is only as good as its weakest link, and people are the weakest link in the chain.” (Schneier (2000)).

An organisation that gives its members access to its resources via a password mechanism needs to ensure that they perform the necessary behaviours if it wants those resources to be protected adequately. However, some of these behaviours can neither be enforced nor monitored. For example, a user might disclose his password to somebody else while he is away from the organisational premises. This would certainly make it difficult, and most likely impossible to monitor him. Organisations then find themselves in a situation where they rely on users to co-operate and perform certain behaviours, when those users can choose other behaviours which would undermine security, but which cannot necessarily be detected (see section 2.3.5.2). Failure to behave in a security-conscious fashion may lead to security breaches, but it will almost always be impossible to prove beyond reasonable doubt that a security breach has occurred as the direct result of a user choosing an insecure behaviour. The research question addressed in this thesis is how organisations can increase the likelihood of users performing the expected password-related behaviours when some of these behaviours can not be enforced or monitored adequately.

1.2 Research questions

There is still only a limited amount of research on the human issues involved in computer security in general, and even less on the human issues that affect password security in particular. Moreover, there is no agreed research agenda, and different researchers focussing on human issues in computer security are sometimes unaware of each other's efforts. The majority of research has assumed that one factor has a strong influence on user behaviour, and has then tried to find ways in which this factor can be changed in order to improve security practices. The factor that has been researched the most is the usability of various security mechanisms (e.g. Zurko & Simon (1996), Whitten & Tygar (1999)). With respect to password security, this has focussed on the issue of password memorability, i.e. on the question of which type of password content is easier to memorize (e.g. Zviran & Haga (1993)). Other password usability issues that have been investigated are the effect of enforced password changes (Sasse et al. (2001)) and the compatibility between password procedures and work practices (Adams & Sasse (1999)). Apart from usability, two other factors have been proposed, but have only been researched to a very limited extent: users'

knowledge and skills (e.g. Thomson & Solms (1998), Adams & Sasse (1999)) and users' motivation (e.g. Parker (1998), Whitten & Tygar (1999)).

The approach of assuming one factor to have a strong influence on user behaviour, followed by an investigation of how this factor can be altered in order to improve security practices, has two immediate drawbacks. Firstly, it makes it less likely that any other factors that influence user behaviour will be identified. Secondly, it does not make it possible to determine the way in which the interaction of different factors causes specific behaviours. However, one strand of research has taken a different approach. Adams & Sasse (1999) focussed on the password mechanism, which is the subject of the research presented in this thesis, and did not start out by assuming certain factors to be responsible for users' password behaviour, but instead aimed to *identify* these factors. They used grounded theory (Strauss & Corbin (1990)) to analyse qualitative interview data and were able to identify a number of factors, some of which, such as the extent to which users perceive password procedures and work practices to be compatible, had been not been addressed by research up to that point. While the work presented in Adams & Sasse (1999) has to be considered seminal in both its substantive and its methodological contribution to the field, it does only identify a small number of factors which is unlikely to be comprehensive, and it does not determine the interplay of these factors to a sufficient extent. The research presented in this thesis is strongly influenced by that of Adams & Sasse (1999) and aims to answer the following three main research questions:

1. What are the factors that affect the password behaviour of users in organisations?
2. How do these factors interact in order to cause specific behaviours?
3. How can knowledge about these factors and their interplay be used to improve user behaviour?

One particular factor that has been identified – the *interpretative repertoires* users draw on to describe aspects of password security – has been investigated further by using discourse analysis (Potter & Wetherell (1987)) on the qualitative data collected. Discourse analysis has also been used to find at least a partial answer to the question of whether users might structure their discourse about password security issues in a manner that makes it possible for them to justify their malpractice. The opportunistic use of data that had been collected prior to a re-conceptualisation of the research approach (see chapter 4) has made it possible to examine the extent to which users violate password regulations. Finally, an analysis of all the qualitative data that has

been gathered has allowed for a first insight into the specific insecure behaviours users choose in certain situations.

1.3 Research approach

The research presented in this thesis is situated in the research discipline of human-computer interaction (HCI), to the substantial and methodological knowledge base of which it contributes. However, in order to address the research problem and the research question, it has been necessary to extend the traditional scope of HCI (see section 2.2). Grounded theory has been used as the primary research methodology, because it is ideally suited to the systematic and empirically-based creation of a theory of complex high-level phenomena about which little is known (see section 3.3). It also had already been successfully applied in the problem domain by Adams & Sasse (1999). Discourse analysis has been used as a supplementary research methodology (see section 3.4), because it makes it possible to investigate how the linguistic resources users draw on to describe aspects of password security can influence their password behaviour. In addition, the question of whether users may structure their discourse in a manner that justifies their malpractice has been examined by analysing (1) the linguistic resources they draw on, and (2) the discursive practices they engage in. Data collection has taken the form of interviews and focus groups, which has made it possible to gain a detailed understanding of the participants' points of view with respect to password issues. This was necessary to deal with a phenomena as complex as the one addressed in this thesis (see section 3.2). Questionnaires and protection motivation theory have been used in the parts of the research that took place prior to the re-conceptualisation of the research approach that is described in chapter 4.

1.4 Research scope

The main objective of this research was to identify the factors that influence the password behaviour of users in organisations. The grounded theory model of these factors and their interplay is based on qualitative data gathered from very specific user groups. However, the use of grounded theory ensures that the model can be transferred to situations that are similar to the one it has been developed in. This makes it necessary to clarify the scope of this research so that readers wishing to apply its findings can identify contexts in which this is appropriate. There are three relevant aspects with respect to which the scope of the thesis research is limited:

1. User population

Users' knowledge about security issues is likely to play a key part in the way in which they approach security issues. An effort to use participants with a varying degree of knowledge about security issues was made in the studies carried out as part of the research. It would have been desirable also to have participants who varied in their general computing expertise, but this was not possible due to access issues. As a result, no novice users of computers were present in the studies.

2. Organisation population

The majority of the participants in the studies came from two organisations, namely the computer science department of a university (University College London) and the research park of a large technology company (British Telecom). These organisations shared certain characteristics (e.g. a focus on knowledge work) and differed with respect to others (e.g. one was a teaching and research institution and the other the research arm of a commercial enterprise). While one study (see section 4.5) gained access to a limited number of participants from an operational arm of British Telecom and confirmed the validity of some of the findings in this context, it is unlikely that the thesis findings will be applicable to organisations that exhibit fundamentally different characteristics (e.g. the military).

3. Factors considered

This research aimed to identify those factors that directly affect user behaviour, and does not take into consideration factors that are further removed and might affect user behaviour indirectly through their effect on other factors. As an example, the research has confirmed a finding of previous research in the area, namely that the usability of the password mechanism has an effect on users' password practices. However, no attempt has been made to identify those factors that are responsible for the poor usability of existing mechanisms and which are likely to be part of the wider organisational context (see section 2.4.7 for a more thorough discussion of this point).

1.5 Thesis contributions

There is only a limited amount of research on the human issues involved in computer security in general, and password security in particular. The majority of this research has focussed on individual factors that are assumed to influence user behaviour, but the research presented here has taken a different approach. It has identified the factors that influence users' password behaviour, and created a model that shows how these factors interact when users choose specific behaviours in particular situations. It has also indicated that the linguistic resources that are used to describe password-related issues can influence user behaviour, and that users might in fact use specific resources or discursive practices to justify their malpractice. The research has also supported the claim that users regularly behave in an insecure manner and has identified ways in which they do so. Finally, an integration of all these findings into an applicable approach to improving user behaviour has been performed. The research findings therefore provide:

1. An insight into the extent to which users violate password regulations and the specific ways in which they do so.
2. A fuller understanding of the way in which users choose their password-related behaviours and of the factors that interact to determine the outcome of this choice process in specific situations.
3. An increased appreciation of how the use of specific linguistic resources to describe password-related issues can influence user behaviour.
4. A heightened awareness of the fact that users might construct their discourse about password-related issues in a manner that makes it possible for them to justify their malpractice.
5. A integrated approach to improving user behaviour, which puts together the findings of the research in a manner that is easily applicable by people who have been entrusted by their organisation with the task of ensuring the effectiveness of password security.

1.6 Thesis structure

Chapter 2 (Background) begins by placing the thesis research within the research discipline of HCI and explains how the work is related to this discipline. It then provides background information on password security and the way in which password mechanisms are typically implemented in an organisational setting. It will be shown that users are required to perform a number of tasks for the password mechanism to function effectively, some of which can be neither enforced nor monitored by the organisation. The research problem is identified as the question of how the likelihood of users performing the required secure behaviours can be increased. The chapter then goes on to present and discuss previous research that has tried to tackle this problem. The substantive and methodological contributions of this research are discussed, and its weaknesses are pointed out. The research questions that have directed the research documented in this thesis are presented, and it is shown which of the weaknesses of previous research they aim to address. The chapter concludes with a summary of the research agenda.

Chapter 3 (Methodology) introduces the methodologies that were used to address the research problem and the research questions developed in chapter 2. The specific methodologies that have been used in the thesis are introduced, specific background information for each of them is given, and their use to approach specific parts of the research is explained and justified. The chapter finishes with a summary of the way in which the different methodologies have been used in conjunction to create the specific research approach used in this thesis.

Chapter 4 (Research chronology) puts together into one place a chronological description of all the studies that have been carried out as part of this research. This is

necessary since the research problem, the research questions and the research approach had not been defined in their current form at the beginning of the research process that led to the findings that will be presented in this thesis. Instead, they were re-conceptualised after initial studies had already been carried out, and data that had been collected up to this point was now re-used and re-analysed in a manner that had not been envisaged at the outset. This means that the connection between individual studies and specific findings that are presented in this thesis is not as clear-cut as it might have been had the final research problem and research approach been used from the beginning. As a consequence of all this, it is felt that a thesis structure in which individual results chapters were preceded by the studies that have informed them could not only be misleading, but would also make the thesis cumbersome reading material. The chronological description of the studies in this chapter also makes it possible to outline the development of the thoughts that have led to the formulation of the research problem, the research questions and the research approach in their final form.

Chapter 5 (Password practices and attitudes towards password security) provides the results of an opportunistic use of the data that has been collected in the studies conducted for this thesis. Quantitative data collected in two field trials at University College London (see sections 4.3 and 4.6) has made it possible to assess the extent to which the password regulations were broken by the students. Questionnaire data collected as part of the same studies has been used to detail students' attitudes towards password security. The qualitative data gathered in all of the interviews and focus groups of this thesis has been used to determine the specific ways in which participants violated password regulations, and the contextual circumstances in which they did so.

Chapter 6 (How users choose password-related behaviours: the core model) presents a grounded theory model of the way in which users choose their password-related behaviours in the absence of any organisational efforts to influence this choice. This model identifies four factors that influence users' choice of password-related behaviours: the level of security that the resource that is being protected requires, the security levels provided by the various password-related behaviours that are available to the user, the user cost incurred by each behaviour and the user benefit provided by it. The model shows how these factors interact to lead to the choice of specific behaviours in particular situations.

Chapter 7 (Extensions to the model) extends the core model presented in chapter 6 to incorporate the effect of password regulations and their associated punishment regimes on users' choice of password-related behaviours. It is shown that this extension can be achieved by adding two additional factors to the model, without needing to change it structurally: users' estimate of the risk of punishment they create for themselves by violating password regulations, and the aspect of a user's personality that determines how he responds to orders. The chapter also offers a summary of the limited insights into the effect of security awareness and education campaigns that can be gained on the basis of the qualitative data collected in the studies. Finally, it is pointed out that users employ the same strategy that they have used to choose their (possibly insecure) password behaviours to determine their response to any measures by the organisation that are aimed at making them aware of security issues.

Chapter 8 (A discourse-analytic study of users' accounts of password security issues) presents the results of analysing the qualitative data gathered as part of this research using discourse analysis. This presentation focuses on the linguistic resources and discursive techniques that were found to have the strongest effect on user behaviour and which could be conglomerated to address larger issues. It is shown that users draw on a limited set of linguistic resources to describe the potential attackers of computing resources, and that this can directly lead to insecure password practices. Linguistic resources that improve user behaviour are juxtaposed to those that worsen it and examples of both are given. Finally, it is shown that there are indications in the data of users at times structuring their discourse in a manner that justifies their malpractice and reduces the likelihood of punishment regimes being introduced.

Chapter 9 (Persuasive password security) shows how the research findings presented in chapters 5 to 8 can be used by organisations to improve user behaviour. Measures are listed that can be undertaken to configure those factors that according to the grounded theory model influence users' choices of password-related behaviours and which organisations have some degree of control over in a manner that will improve these choices. It is shown how organisations can use the discourse-analytic findings of the thesis to construct the discourse about password security in their communications with users so that user behaviour will be affected positively. The difficulty of getting users' attention in security awareness and education campaigns is discussed, and measures that can be undertaken to address this issue are proposed. Finally, an

integrated approach to improving user behaviour is presented, which is based on all the findings in the thesis and presented in a manner that makes it easily applicable by people who have been entrusted by their organisation with the task of ensuring the effectiveness of password security. This approach has been named *persuasive password security*.

Chapter 10 (Conclusions) summarises both the substantive and the methodological contributions made by the research presented in this thesis. It also contains a critical review of the thesis research and points out areas for further research.

1.7 Terminology, typography and gender-specific pronouns

Throughout the thesis, the term *user* refers to the user of a password mechanism within an organizational context. *Computer system* and *system* will be used interchangeably to denote the actual computer system he uses within the organization (e.g. a PC, or the larger network behind it). *Computing resource* and *resource* both refer to the actual resources he gets access to by authenticating himself successfully to the password mechanism (e.g. his email, databases or word-processing facilities). The term *beneficiary* will be used to refer to the envisaged user of the research results presented in this thesis (e.g. other researchers continuing this research, or security managers applying it within an organization).

This thesis presents the results of both a grounded theory and a discourse-analytic analysis of the qualitative data that has been collected. A specific typography is used for grounded theory categories and properties once they have been introduced in the text. In the same way, another typography is used for discourse-analytic interpretative repertoires.

Throughout this thesis, the male pronouns he/his/him will be used in a non gender-specific way, i.e. they can be read as (s)he, his/her and him/her. Cases where pronouns are used in a gender-specific manner that is not obvious from the surrounding text will be pointed out explicitly.

2 BACKGROUND

2.1 Overview

This chapter serves three functions within the wider context of the thesis:

1. It presents the background information that is necessary to understand the research presented in this thesis.
2. It develops the research agenda that has driven this research, consisting of
 - a fundamental research problem and
 - a set of questions, the answers to which will contribute to the solution of this problem.
3. It provides the justification for the research agenda and for the research presented in this thesis.

The chapter is divided into three parts:

1. Section 2.2 describes human-computer interaction (HCI) as the research discipline within which the work documented in this thesis is situated, and explains how the work is related to this discipline. In addition, a number of extensions to the traditional scope of HCI will be suggested in order to make it possible to tackle the problem addressed here.
2. Section 2.3 provides background information on password security, and on how password mechanisms are typically implemented in an organisational setting. It will be shown that users are required to perform a number of tasks in order for the password mechanism to function effectively. This means that organisations rely on users to perform certain behaviours for their resources to be protected. Some of these behaviours cannot be enforced, or monitored adequately. The question of how the likelihood of users performing secure behaviours can be increased is the research problem that has driven the research documented in this thesis.
3. Section 2.4 presents and discusses previous research that has tried to tackle this problem. The substantive and methodological contributions of this research will be discussed, and its weaknesses will be pointed out. The research questions that have directed the research documented in this thesis will be presented, and it will be shown which of the weaknesses of previous research they aim to address.

The chapter summary (section 2.5) will recapitulate the research agenda and the following chapter (chapter 3) will describe the methodological approach that was taken to implement this agenda.

2.2 Human-Computer Interaction (HCI)

Human-computer interaction (HCI) originally was a joining of software engineering and human factors engineering, and consequently lies at the intersection between information technology on the one hand, and the social and behavioural sciences on the other (Carroll & Campbell (1989)). As a discipline, it is concerned with how people make use of devices that incorporate or embed computation, and with how both the utility and the usability of such devices can be increased. It is generally

agreed upon that the aim of HCI research is to generate a knowledge base that can be utilized by designers to create usable computer systems (Preece et al. (1994)).

Human-computer interaction initially had a vision of itself as an applied science that was to bring cognitive science methods and theories to bear on software development (Carroll & Campbell (1989)). However, in the late 1980s, a large number of new scientific ideas entered the HCI mainstream. This has led to a scientific foundation that is far more expansive and diverse than was initially envisaged. As a result, HCI encompassed nearly all of the social and behavioural sciences by the mid-1990s. The commitment to multidisciplinary science has been an important element in the development of HCI up to this point, and will remain a key to its future (Carroll & Campbell (1989)).

Nowadays, HCI does not only analyse and design user interfaces and new user interface technologies, but also studies and aims to improve the work and organisational processes of technology development in general (Carroll & Campbell (1989)). A large variety of methods for understanding the tasks and work practices of people and their organisations have been developed, and these can be used to frame new possibilities for the application of technology to support human activities. The methods employed in HCI range from the development of checklists to cognitive walkthroughs, from field studies of workplaces to laboratory experiments, and from interviews and surveys to various kinds of analytic models. In addition, HCI has become a primary test bed for participatory design and ethnographically driven design, with the former directly involving users in design work and the latter studying work practices in order to ensure that new technology supports work as it is practiced in reality, and not as it is normatively described in procedures.

The influx of new ideas in HCI has widened the scope of the discipline in ways that are relevant to the research documented in this thesis:

1. Acknowledgement of the wider context of the interaction between user and computer

It has been recognised that the interaction between a user and an interactive computer system does not occur within a vacuum, but takes place in a specific context. This wider physical, social and organisational context can influence the behaviour and the motivation of the user. Dix et al. (1998) point out that this wider context may contain factors over which the designer of interactive systems may not have any control, but that it is important to be aware of such factors in order to understand the user and the work domain fully.

2. Recognition of further stakeholders

Socio-technical approaches have recognised that technology is not developed in isolation but as part of a wider organisational environment. As a result, some of these approaches consider all stakeholders, rather than just the end-user (Dix et al. (1998)). In this context, the term stakeholder refers to anyone who is affected by the success or failure of the interactive computer system. As an example, CUSTOMS (Kirby (1991)), a socio-technical approach developed for use in smaller organisations, distinguishes 4 categories of stakeholders: *primary stakeholders* are the users of the system; *secondary stakeholders* do not use the system directly, but receive output from it or provide input to it; *tertiary stakeholders* are neither primary nor secondary stakeholders, but are affected by the success or failure of the system (e.g. the owner of the resource that is protected by a password mechanism); and *facilitating stakeholders* are involved in the design, development and maintenance of the system (e.g. system administrators).

3. Beyond usability

Research on pleasure-based approaches in HCI (e.g. Jordan (2000)) aims to create interactive computer systems that are not only usable, but which also provide the user with some degree of pleasure.

The work documented in this thesis primarily contributes to the substantive and methodological knowledge base of the field of HCI, to which it is most closely related. However, it also goes beyond the scope of what traditionally constitutes HCI research in a number of respects, some of which are covered by the wider scope of the new approaches presented above:

1. Inclusion of behaviours that are not part of the human-computer interaction

The research problem the work presented in this thesis addresses is the question of how the likelihood of users in an organisation performing the expected password-related behaviours can be increased in a situation where some of these behaviours can be neither enforced nor monitored adequately (see section 2.3.5.2). The research questions it tries to answer focus on the individual factors that affect users' password behaviour in organisations. Some of these behaviours will be expressed while the user is interacting with a computer, e.g. by entering a cryptographically strong password. However, other behaviours that are of importance actually take place while the user is not interacting with a computer. One example of this would be a user who is away from the office and who is being asked for his password by a stranger, with the expected behaviour being a refusal to disclose the password. The model aims to capture both the factors that are expressed while a user interacts with a computer and those that are performed while he does not interact with a computer.

2. Beyond usability

Section 2.4.6 will point out that authentication to a system is in most cases an *enabling task*, which means it creates an overhead for the user, who is using the system as a tool to achieve a *primary, real-world task*. It is predictable that most users will cut corners to reduce that extra load given a chance, unless they are *motivated* to make the effort to behave in a security-conscious fashion. Users in an organisation will very often not put themselves, but the organisation at risk when neglecting their security-related duties. This might lead to them lacking the motivation to make the extra effort that is needed to behave in a security-

conscious fashion. One way of dealing with this issue is to increase the usability of the password mechanism to a point where even users with little motivation are willing to make the small effort that the new, highly usable mechanism requires them to put in. However, this might not always be possible and even if it were, some users may still lack the required minimum level of motivation. Organisations often approach this problem by introducing measures such as punishment regimes for users who are found to behave in an insecure manner. Factors such as these are typically addressed by HCI research in a passing fashion only, since they fall outside of the control of the designers of the system in question (see point 3 below). However, they form a fundamental part of the model presented in this thesis. In addition, it will be suggested in chapter 9 that the designers of password mechanisms should also look beyond pure usability issues by taking into consideration the effect of their design decisions on users' attitudes and behaviour from the outset. This is a move not dissimilar to the one taken by researchers into pleasure-based approaches within HCI (e.g. Jordan (2000)).

3. **Envisaged beneficiaries of the model**

The model presented in this thesis will contain a number of factors that fall under the control of the designer of the password mechanism and which would most certainly be considered by traditional HCI research. The most notable example of such a factor is the usability of the password mechanism itself, which is also the factor that previous HCI-inspired research in this field has focussed on to an almost exclusive extent (see section 2.4.4). However, other factors will not be under the control of the designer of the password mechanism and would most probably only be investigated by HCI research in order to understand the user and work domain more fully. An example of one such factor are the security guidelines that are put into place by the security department in many organisations and which typically threaten users with some form of punishment should they perform certain insecure behaviours. As a result, the model presented in this thesis will include factors that are under the control of the designer of the password mechanism as well as factors that are under the control of other people, such as the creators of security guidelines. Unlike an approach that focussed solely on the HCI elements of the password mechanism, the latter factors would also be considered to be under the control of the beneficiary of the model, who would employ it to improve overall security. In other words, the envisaged beneficiary of the model within an organization will be anyone who has been entrusted by the owner (a tertiary stakeholder) of the resource that is protected by the password mechanism with ensuring the success of this mechanism.

These three extensions to the traditional scope of HCI make it possible to address the research problem and the research questions that have driven the research documented in this thesis. It is also suggested that these extensions should be considered for incorporation into the canon of what constitutes HCI research in order to equip the field with the ability to deal with problems of a similar nature to the one discussed here. As such, they form a methodological contribution of the thesis.

2.3 Password security

2.3.1 Computer security and the importance of access control

In its most basic sense, computer security deals with the protection of computer-related assets such as the computers themselves, the data and software stored and running on them and the networks they are connected to (Gollmann (1999)). The process of security engineering usually starts out with the development of a threat model, which identifies the specific threats the resources in question need to be protected against (Anderson (2001)). This drives the development of the security policies, which specify clearly and concisely the security objectives that are to be attained. The policies, in turn, inform the design of the security mechanisms, and are in effect implemented through these mechanisms. Since policies are written at a broad level, organisations will often develop standards, guidelines and procedures that offer a clearer approach to implementing policy and meeting organisational goals (National Institute of Standards and Technology (1995)). *Standards* specify the uniform use of specific technologies and procedures to secure systems, whereas *guidelines* assist users in developing system-specific standard procedures. Procedures themselves are simply detailed steps to be followed by all involved in securing the resources in question. Security standards, guidelines and procedures are often disseminated throughout an organisation via handbooks, regulations, or manuals.

There are three objectives which any secure system would be expected to attain (Schneier (2000)):

1. *Confidentiality* ensures that any data stored on a system or transmitted between systems is disclosed only to authorised individuals.
2. *Integrity* safeguards that no data stored on a system or transmitted between systems has been modified, deleted or created by unauthorized individuals.
3. *Availability* makes certain that the data and services provided by a system are available when users require them.

Ultimately, confidentiality, integrity and availability are provided through access control mechanisms (Schneier (2000)): we want to make sure that authorized users are able to do whatever they are authorized to do, whereas everyone else is not. Access control mechanisms usually are the first line of defence any malevolent would-be intruder to a computer system has to breach, and as such can be said to form the foundation on which much of computer security is built Anderson (2001).

2.3.2 The dominance of password security – and why it will continue

On most computer systems, access control is carried out by a process of user *identification* and *authentication* (Garfinkel & Spafford (1996)). Identification establishes who the user in question claims to be. Authentication then verifies that the user is who he says he is. There are basically three ways to perform authentication. They are based on something the user knows, has or is:

1. *Knowledge-based authentication* uses a secret piece of information, such as a password or passphrase, known only to the user and the computer system.
2. *Token-based authentication* uses a physical token such as a smart card, which is difficult to forge and obtain and has to be in the user's possession when he uses the access control mechanism.
3. *Biometric authentication* uses some unique characteristic(s) in a person's physical appearance or behaviour, such as their fingerprint or keystroke dynamics.

Today, knowledge-based authentication, in particular in the form of the password mechanism, is the most widely used authentication mechanism. It is unlikely that this situation will change significantly in the foreseeable future. Password mechanisms have a number of vulnerabilities and usability problems, as will be pointed out throughout this thesis, but so do their possible replacements. Alternative knowledge-based mechanisms, such as passfaces, have not been tested sufficiently in real-world scenarios to ascertain their usability, and they also raise a number of implementation issues (Brostoff & Sasse (2000)). A physical token might be stolen, copied or left in its slot by the user (Schneier (2000)). Token construction and distribution is also far from trivial, and has led to documented financial loss (Anderson (1994)). Biometric solutions can be unpopular with users, who are afraid of a 'Big Brother' scenario, in which their every activity is being monitored (Deane et al. (1995)). Intruders might also obtain digital representations of biometrics by stealing them from a database, by intercepting them while they are being sent over a network during authentication, or by replicating them from analogue copies left behind by users (such as fingerprints on an empty beer glass). They can then be used to impersonate the corresponding user with impunity (Kim (1995)).

Apart from the specific vulnerabilities and usability problems all authentication mechanisms possess, a deciding factor for the choice of a specific solution will always be its cost and the ease with which it can be implemented. Password mechanisms are often chosen for their superiority in both respects (Anderson (2001)). In addition, physical tokens and biometrics are currently almost always used in conjunction with a password, as in smart cards that require the user to provide a PIN. All this means that

knowledge-based authentication in the form of the password mechanism will be with us for a long time to come.

2.3.3 How password security is implemented

The principle underlying password security is deceptively simple: the authorised user of a computing resource and the password mechanism that has been put into place to control access to this resource share exclusive knowledge of a secret password, which is disclosed by the user during login to authenticate himself. The manner in which this basic principle is implemented varies between different application software and operating systems. In a simplified manner, a typical implementation, as in existence in most organisations today, works as follows:

1. The administrator of the password mechanism issues every new user with a username (for identification) and a password generated randomly by the mechanism (for authentication), and sets up a new account with these details. Usually, the user will be handed this information in paper format and uses it to log in for the first time.
2. After the first login, the user will change the system-generated password to one of his own choice. However, most systems will only accept user-generated passwords if they comply with certain criteria believed to ensure cryptographic strength, which makes offline dictionary attacks more time-consuming (proactive password checking (Stallings (1995))). There is no agreed standard for the cryptographic strength of a password. It is usually assumed that it increases with the length of the password, the size of the character set it is drawn from, and the extent to which the sequence of characters in it is random (e.g. Anderson (2001), Schneier (2000)).
3. The system stores the usernames and associated hashed passwords in a specific file, which on some systems is public and on others is kept protected. Even if the hashed password file gets stolen, the attacker cannot directly recover the unencrypted passwords from the hashed ones, since a hash function is not reversible (Schneier (2000)). In addition, random bits may be added to the plain text password before hashing it in order to increase the search space for dictionary attacks (see section 2.3.4.2) even further. As a result of these extra bits, which usually are called the *salt*, identical plain text passwords will be coded into different hashed passwords.
4. Users are supposed to memorise their password and recall it during login. They use the keyboard to enter their username and password into a dialogue screen. Feedback is given on screen, but the password itself is not being echoed. Users are expected to ensure that they are not being observed typing in their password.
5. The password is then hashed by the system, and, together with the username, compared to the stored entries in the password file. If a match is found, the user will be given access to the computer system. A lot of systems implement what is called a three-strike policy and suspend a computer account if a certain number of unsuccessful login attempts has been made (Viega & McGraw (2001)).

6. If a user forgets his password, he will not be able to log on to the computer. Instead, he will have to contact the administrator (usually in the form of a helpdesk) for the password to be reset.
7. Most systems implement password ageing, which forces users to change their passwords at regular intervals. If a change is not made in time, the username in question will be suspended and the user has to contact the administrator to have it activated again. This aims to reduce the likelihood of successful dictionary attacks and also tries to prevent a situation where a password gets compromised without the user's knowledge and is used by an attacker indefinitely.

2.3.4 How password security is breached

An attacker trying to access a computing resource that is password-protected essentially has two options:

1. He can attempt to access it by logging in through the password mechanism, for which he needs knowledge of both the username and the password.
2. He can try to circumvent the password mechanism altogether.

The latter can be achieved in two ways: he can persuade the user, either overtly or covertly, to log him in, or he can exploit software weaknesses to hack his way into the resource. Obviously, the overt way of getting a user to log him in is to ask him to do so, for whatever reason. Covert ways include using a computer that an authorized user is logged into while he has left his place momentarily, or sending emails with virus attachments, which, if opened, effectively give the attacker access to some or all of the resources. These types of attacks will not be considered in this thesis. Neither will the exploitation of software weaknesses be dealt with, but it is still important to keep this possibility in mind, since it forms the basis of some attacks that are initially dependent on the knowledge of a password. In some cases, the attacker of a large networked computer system, which offers many and varied resources to its users, may choose to obtain access to a subpart of that system, e.g. a weakly protected user's account. Once he has this access, he can exploit software weaknesses to hack into other parts of the system. This is the basis of many dictionary attacks (see section 2.3.4.2), and it means that systems that get targeted like this are only as secure as their worst-protected account (Schneier (2000)).

An attacker trying to access the resource by logging in through the password mechanism needs knowledge of both the username and the password. It is usually easy to obtain the username, so his real challenge is to get hold of the corresponding password. The methods he can employ to do so broadly fall into three categories, which will be discussed individually in the following sections.

2.3.4.1 Online guessing

An attacker trying to access a specific resource via the password mechanism without knowing the password can repeatedly try to log in, using his best guesses as the password. When the log-in attempt is carried out manually, it will be so slow compared to offline dictionary attacks (see section 2.3.4.2) that the chances of success entirely depend on the quality of the guesses. Typically, these will be based on knowledge about the user who chose the password (e.g. his spouse's name) or about the general preferences users have when choosing passwords (e.g. the names of popular football teams). In any case, such an attack is only likely to succeed if the password chosen is extremely weak.

The attacker can instead choose to carry out the online attack in an automated fashion. If he does not want to target a specific account in a multi-user environment, but just wants to get access to *any* account in the system, he will be able to launch parallel attacks on a large number of accounts, making it possible for him to try out a hundred or more guesses in a second (Pinkas & Sander (2002)). There are two common countermeasures against such online attacks. Firstly, the server can introduce a small time delay before responding to a login attempt, which will prevent an attacker from making a sufficiently large number of guesses in a reasonable amount of time. Secondly, accounts can be locked after a few unsuccessful login attempts, which reduces the number of guesses an attacker can try out. However, both these measures have weaknesses (Wang et al. (2005)). Global password attacks on large multi-user environments can circumvent the effect of delayed responses by increasing the number of login attempts that are made in parallel. Account locking not only lays the system open to denial-of-service attacks, but also can be circumvented by ensuring that the number of login attempts that are made for a specific username never reaches the threshold that triggers the account locking mechanism. As a result, online attacks in large multi-user environments can often use the same search strategy as offline dictionary attacks, which will be discussed in the next section.

2.3.4.2 Offline dictionary attacks

Offline dictionary attacks require the attacker to have gained access to the password file, which contains the hashed passwords for all the accounts of a particular system. A dictionary attack tool, which can be downloaded easily from the Internet, is then used to carry out a form of high-speed guessing. The tool contains a substitute

password mechanism, which hashes passwords in the same manner as the password mechanism under attack. Hashed guesses can then be compared to the entries in the password file. If a match is found, the attacker has identified the password to one of the accounts of the system. The tool also contains a dictionary of guesses, which are ordered based on the preferences of users when they choose passwords (e.g. names, dictionary words or famous football teams). The tool will first run a *dictionary attack* by using all the entries in the dictionary as its guesses (Yan (2001)). This is typically followed by a *hybrid attack*, which performs string substitutions on the dictionary entries by using methods that are often employed by users (e.g. appending letters or substituting letters for visually similar numbers), before using them as guesses. This is more time-consuming than a dictionary attack, since there are many possible alterations for each entry in the dictionary. Finally, the tool may carry out a *brute-force attack*, in which all possible password combinations of increasing length are used as guesses. This is the slowest attack form, but has the highest chances of success. Schneier (2000) reports that L0phtcrack, a password recovery hacker tool optimized for Windows NT passwords, can try every possible keyboard password for NT's weaker password function (where the passwords are case-insensitive, and cannot be much stronger than seven characters) in 480 hours on a 400 MHz Quad Pentium II. The effectiveness of offline dictionary attacks has obviously improved with the increase in computing power over the past decades, and it will continue to do so. As an example, Perrine & Kowatch (2003) reports that the Unix crypt() function, on which the Unix password protection system has depended for over 30 years, has to be considered obsolete in the face of attacks that use high-performance computing resources to pre-compute and store hashed passwords

The possibility of both online and offline dictionary attacks has two immediate consequences for users. Firstly, it requires them to choose passwords of higher cryptographic strength than would be necessary in the absence of these threats. Secondly, it makes it necessary to change passwords regularly, since these attacks will crack any password, given sufficient time. However, offline dictionary attacks are only possible if the attacker gains access to the password file. While Unix famously makes this file world-readable, more recent operating systems, such as NT, aim to ensure that it is only accessible to the systems administrator. Assuming that such operating systems do not have flaws, and are configured properly, an attacker can only get access to the password file by achieving administrator status. Once he has

managed to do this, it is difficult to see how cryptographically strong passwords will reduce the damage he can cause. He already has access to the whole system, and can even change users' passwords without needing to know their current one. Sadly, most operating systems currently in use do not protect the password file adequately (Anderson (2001)).

2.3.4.3 Obtaining the password

2.3.4.3.1 Persuading the user to disclose the password

The most straightforward way to obtaining a password is to tell someone a lie to extract it directly from someone who knows it (which could be someone other than the user in question, e.g. his systems administrator). This practice is called social engineering, and it is both widespread and highly successful (e.g. Winkler (1997), Mitnick & Simon (2002)). In a study at the University of Sydney, 138 out of 336 computer science students returned a valid password after being emailed with a false request to supply their password so that the password database could be validated after a break-in (Greening (1996)). This example might be harmless, but Ira Winkler provides a number of case studies in which he demonstrates the effectiveness of social engineering techniques as a means for corporate espionage, which according to the FBI costs U.S. companies anywhere from \$24 billion to \$100 billion annually (Winkler (1997)).

Another approach is password harvesting, where an attacker exploits the fact that many users use the same password for different applications. He sets up a website with some interesting content, makes it password-protected, adds a questionnaire asking new users which other systems they use, and tries out the password entered on these systems (Schneier (2000)). This makes it vitally important for users to use their password only on the system it has been created for.

One final approach worth mentioning is often dubbed 'spoofing'. An attacker might run a program on an unattended computer, displaying the usual logon screen. When a user enters his username and password, the system stores it, replies "Sorry, wrong password" and terminates itself, invoking the proper password program. One way of avoiding this attack is to have a trusted path (such as the CTRL-ALT-DELETE key combination in Windows NT, which is guaranteed to take the user to a genuine password prompt (Anderson (2001))), which can be considered a way of ensuring the computer system authenticates itself to the user trying to log in.

2.3.4.3.2 Finding the password

Users often keep material copies of the password, either to avoid having to memorise it, or to aid them in case they forget it. Each material copy of the password is a potential target for attackers, who can, among other things, search a user's workspace for pieces of paper with passwords on them, steal PDAs and try to find passwords, or even scour an account they have compromised for electronic copies of passwords for other systems.

2.3.4.3.3 Exploiting the login process

Attackers can try to identify which keys a user is pressing when he is typing in his password during login. This can be done by personal observation, a practice commonly dubbed *shoulder-surfing*, or by technical means, such as keyboard tapping.

2.3.5 Interim summary

2.3.5.1 Users' tasks

The primary goal of a user attempting to log into a computer system via a password mechanism is to access the resource(s) protected by that mechanism so he can perform tasks such as reading his email or writing a report. From the point of view of the organization providing these resources to the user a second goal, of equal importance to the first one, should be to protect the resource(s) as much as is possible from access by unauthorized third parties (section 2.4.6 will discuss the wider issue of users' motivation to make the effort to achieve this second goal). The tasks the user needs to perform in order to accomplish *both* these goals include:

1. Enrolment

- 1.1. The user obtains the username and the initial password (typically from the helpdesk).
- 1.2. He uses the username and initial password to log into the system for the first time.
- 1.3. He selects a new password in accordance with guidelines to ensure sufficient cryptographic strength (often enforced by the system through proactive password checking, see section 2.3.3), and different from passwords used to access other computing resources.
- 1.4. He changes the password to this newly-chosen one.
- 1.5. He memorises the password.

- 1.6. He destroys any trace of the initial password and the newly-chosen one (e.g. a paper copy provided by the helpdesk).
2. Normal use
 - 2.1. He retrieves the correct username.
 - 2.2. He retrieves the correct password.
 - 2.3. He uses the username and password to log into the system, making sure that this information cannot be observed by third parties as it is being entered.
3. Prevent password expiry
 - 3.1. At regular intervals, usually enforced by the password mechanism through password aging (see section 2.3.3), he selects a new password in accordance with guidelines and different from passwords used to access other computing resources.
 - 3.2. He memorises the password and destroys any trace of it.
 - 3.3. He changes the password to this newly chosen one.
4. Handle login failure
 - 4.1. He contacts helpdesk if it is not possible to retrieve the username and/or the password, or if the system refuses to accept a combination believed to be correct.
 - 4.2. He authenticates himself to the helpdesk.
 - 4.3. He obtains his username and a new password for initial login, and continues as off step 1.2.
5. Handle attempts by third parties to extract password from user
 - 5.1. Should the user be approached by any third party to disclose his password, he has to refuse to do so.

Throughout the studies carried out as part of the research conducted for this thesis, guidelines about the behaviours they were expected to perform had to be given to users. These were chosen on the basis of the list above, but had to be altered slightly to satisfy the wishes of the relevant authorities both at University College London and at British Telecom, where the studies were carried out. The resulting list of guidelines took the following form:

1. Users need to choose cryptographically strong passwords.
2. They need to choose a unique password for each password-protected resource they access.
3. They need to memorise their passwords (as opposed to keeping a physical copy).
4. They must not share their passwords with third parties.

5. They need to change their passwords at regular intervals.

2.3.5.2 The research problem

An organisation which gives its members access to its resources via a password mechanism needs to ensure that they perform the behaviours described in section 2.3.5.1 if it wants those resources to be protected adequately. There are two obvious ways in which the organisation can aim to achieve this:

1. The organisation can enforce a behaviour, as is done in the case of enforced password ageing. Here, users are forced to change their password at regular intervals and will not be able to log in unless they have done so. However, users might find ways of undermining this. As an example, they might just switch between two passwords on a monthly basis. If the mechanism keeps a history list of past passwords and makes it impossible to choose them again, they might just choose a password with an index that increases with every enforced change (e.g. 'peter1', 'peter2', 'peter3'...). The organisation might find ways of protecting itself against this, but it is clear that not all of the behaviours that users are required to perform are also enforceable.
2. The organisation can monitor users and punish misbehaviours. However, this might seriously clash with the organisational culture and result in users rebelling against a 'Big Brother' scenario. More importantly, some of the required behaviours cannot be monitored. As an example, a user might disclose his password to somebody else while he is away from the organisational premises. This would make it difficult, if not impossible to monitor him.

This means that organisations are faced with a fundamental problem: they require their users to perform a number of behaviours, some of which can be neither enforced nor monitored. In other words, they rely on the users to co-operate and perform certain behaviours when they could choose other behaviours which would undermine security but which cannot necessarily be monitored. Failure to behave in a security-conscious fashion may lead to security breaches, but it will almost always be impossible to prove beyond reasonable doubt that a security breach has occurred as the direct result of a user choosing an insecure behaviour. The research problem this thesis tries to address then is the question of how the likelihood of users in an organisation performing the expected password-related behaviours can be increased in a situation where some of these behaviours can be neither enforced nor monitored adequately.

2.4 Previous research on the human issues affecting computer and password security

2.4.1 Threats in the civil world: (re)discovering the human link

Research on security mechanisms to date has focussed almost entirely on technical issues, such as encryption algorithms or firewalls. Very little work has been carried out with regards to the human issues that are involved. This approach has been criticized as early as 1975, when Saltzer & Schroeder (1975) included usability in the stated goals of the secure system. Davis & Price (1987) have pointed out that human factors should be considered in the design of security mechanisms, since security is ultimately implemented and breached by humans. Hitchings (1995) has suggested that the narrow, technology-oriented perspective has produced security mechanisms which are much less effective than they are generally considered to be. However, it is only in recent years that the security community has realized that a large number of security breaches that are being reported have been enabled and facilitated by user behaviour. With respect to password security, there are numerous reports of users regularly not performing the required tasks identified in section 2.3.5.1 (e.g. Schneier (2000), Winkler (1997)): they choose weak passwords, use them for several systems, write them down, and share them easily. This gets exploited by attackers of computing resources. Kevin Mitnick, arguably the world's most famous hacker, testified to the US Senate committee that he had obtained 9 out of 10 passwords by tricking users, rather than through cracking. In his new role as security evangelist, he emphasises that:

“The human side of computer security is easily exploited and constantly overlooked. Companies spend millions of dollars on firewalls, encryption and secure access devices, and it's money wasted, because none of these measures address the weakest link in the security chain.” (Poulsen (2000)).

It is now commonly accepted by the security community that *“security is only as good as its weakest link, and people are the weakest link in the chain.”* (Schneier (2000)). As a result, first research into the human issues involved in security has been carried out, which will be presented in the following sections.

How is it possible that the security community has ignored human factors, which are clearly of utmost importance, for such a long time, despite the early and repeated warnings? Zurko & Simon (1996) point out that the majority of security research has strong roots in the military, where users can be expected to follow even the most

onerous rules and procedures precisely. This has reduced the pressure on the designers of security systems to make them usable. It has also led to an emphasis on mathematical modelling, which does not guarantee a system that emulates user intuitions and is easy to use. This point is echoed by Dhillon & Backhouse (2001), who have mapped out the current territory of information systems and security research, paying particular attention to the socio-philosophical concerns of the various approaches. They stress that most security research has followed a functionalist paradigm, and shares major assumptions with the military. In particular, it assumes a predominant culture of trust among the members, a system of clear roles and responsibilities, hierarchical organisational structures and largely centralized information processing. It can obtain appropriate results for a reality that is well-defined, such as the military. However, as organisational structures, particularly those of commercial enterprises, have become flatter and more organism-like in nature, a socio-organisational perspective is needed to address social groupings and the behaviour of people.

2.4.2 A taxonomy of previous research on the human issues involved in computer and password security

There is still only a limited amount of research that deals with the human issues involved in computer security. Researchers in this field come from a variety of backgrounds, and at times seem to be unaware of each other's work. This is particularly true for research carried out in the two main contributing fields of human-computer interaction and information systems management. At present, there is also no agreed research programme that can guide and integrate the efforts that are being undertaken. However, it is possible to create a first taxonomy of existing approaches, which can then be used to position individual research efforts and to identify gaps that need to be filled if there is to be an integrated research programme in the future.

Siponen (2000b) has proposed a first taxonomy of the different approaches taken by research into the human issues affecting computer security that has been carried out so far. He bases this taxonomy on the methods that are employed to influence user behaviour, and divides the existing approaches into two categories: those that try to affect users' behaviour by introducing external deterrents, such as punishment, and those that try to affect it without any such deterrents. The latter category is subdivided into approaches that try to affect the human component per se, predominantly by

means of security awareness, education and training, and those that aim to improve the security mechanisms themselves, mainly by improving their usability.

The taxonomy proposed here is in parts similar to that put forward by Siponen (2000b), but the reasoning underlying it is slightly different and it also encompasses research not considered by Siponen. The central tenet of all the different strands of research is that users of security mechanisms regularly behave in a manner that undermines the security of the resources in question. There are a number of factors that are likely to affect users' security behaviour, such as the usability of the security mechanism or the motivation of the user in question. It will be argued that the overwhelming majority of research into the human issues involved in computer security has so far not tried to identify explicitly those factors that impact on user behaviour. Instead, the various strands of research have focussed on one or several of the factors that are likely to affect it, implicitly assuming that these factors are responsible for users' insecure behaviours, and that altering them will result in improved security practices.

The argument presented above makes it possible to classify the large majority of research according to the factor(s) that are assumed to be responsible for users' choosing insecure behaviours. Three factors in particular have been the subject of research:

1. Usability of the security mechanism

The majority of work has focussed on the low usability of existing security mechanisms, and provides advice on how the user cost of performing secure behaviours can be reduced by altering the mechanisms in question (e.g. Zurko & Simon (1996), Whitten & Tygar (1999), Adams & Sasse (1999), Brostoff & Sasse (2001)).

2. User knowledge and skills

Users' limited knowledge of security issues, and ignorance of the skills that are required to behave in a secure manner have been proposed as a cause of insecure practices (e.g. McLean (1992), Spurling (1995), Adams & Sasse (1999)).

3. User Motivation

Users' limited motivation to make the effort that is required to behave in a security-conscious fashion has been identified as a factor that can undermine security (e.g. Parker (1998), Schneier (2000)).

In addition to this, a limited amount of research has been carried out with respect to two additional issues, and cannot be classified as focussing on a factor that is assumed to be the cause for improper user behaviour:

4. Studying user behaviour

The claim that users regularly behave in a manner that undermines security has largely been substantiated with anecdotal evidence. There is only a very small

amount of research that has tried to validate this claim quantitatively (e.g. Greening (1996), Yan et al. (2000)). In addition, there has been hardly any research on the exact way in which users undermine security, i.e. on the actual insecure behaviours they choose instead of the ones they ought to perform.

5. Investigating the wider organisational context

Once individual factors that lead to users behaving in an insecure fashion have been identified, it will often be possible to alter them in order to improve user behaviour. However, another question can be asked: how did the situation where these factors were allowed to affect user behaviour negatively arise in the first place? In other words, what are the factors that affect the factors that impact on user behaviour in a direct fashion? Answering this question would make it possible to identify interventions at different levels of the organisational system and to determine which of these interventions will have the strongest positive effect on user behaviour and overall security. There has only been a single research effort that has tried to address this specific issue (Brostoff & Sasse (2001)).

The following subsections will present the research that has been carried out in each of these categories. The fourth category, i.e. research that has tried to study actual user behaviour with respect to the use of password mechanisms, will be dealt with first. This will set the scene for a discussion of the research that has looked at individual factors that have been assumed to influence user behaviour (categories 1 to 3 above). Finally, research that has tried to investigate the wider organisational context and the way in which it affects password security (category 5) will be presented. It is important to point out that the taxonomy presented above is meant as an aid in positioning previous research efforts in order to make it possible to identify areas that require additional work. It is not meant as a taxonomy of all issues that might possibly affect human factors in password security.

One final point that needs to be made before we can look at the research carried out in the individual categories is that most of this research has focussed on mechanisms other than the password mechanism that is the subject of this thesis. In addition, a substantial part of it has looked at personal users, and not at users in an organisational context. Only those aspects of this research that shed light on issues that affect password security in an organisational context, or that make points that are relevant to the overall argument presented in this thesis will be reviewed. Research that touches on several of the categories listed above will be discussed in detail in its main category and will be referred to briefly in the other appropriate categories.

2.4.3 Studying user behaviour

The central tenet of the different strands of research on human factors in computer security is that users regularly behave in a manner that undermines the security of the resources in question. However, most of the evidence for this claim is of an anecdotal nature (e.g. Winkler (1997), Schneier (2000)). There is only a limited amount of research that has tried to quantify the extent of the problem, and there is even less research that has tried to identify which specific insecure behaviours users choose over secure ones. With respect to password security, a particular problem such research faces is the very issue that has given rise to the research problem driving the work presented in this thesis: a lot of the required behaviours cannot be monitored easily. This also means that it is often difficult to get hard, quantifiable data on the extent to which these behaviours are performed by users. A result of this is that the research that has provided such hard data has focussed on the one behaviour that can be verified easily, namely the strength of the password chosen by users. An example of such work is the study by Yan et al. (2000), in which graduates were asked to select strong passwords and the authors then attempted to crack the passwords in the resulting password file using a relatively short-running dictionary attack. The results were alarming: 32% of the passwords were cracked this way. Similar results were found in studies conducted in a business environment (e.g. Belgers (1993), Klein & Myers (1999)).

Studies on cryptographic strength of the passwords users choose have focussed on the one behaviour that can be tested easily, provided access to the password file is obtained. In cases where this access was not given, researchers had to resort to questionnaires as a means of studying user behaviour. An example would be a recent poll (Petrie (2002)) that found that 90% of 1200 users reported choosing passwords that would easily be cracked, such as dictionary words or names. 47% of the users even chose their own name, their nickname or the names of their partners, children or pets, which could be guessed online rather than cracked offline. Only 9% of this sample reported using cryptographically strong passwords. Adams et al. (1997) collected questionnaire data on two other behaviours users need to perform. Of their 139 respondents, 50% stated that they wrote down their passwords in one form or another. In addition, 50% of respondents answered that they had devised a method to construct related passwords, i.e. instead of choosing different and unrelated passwords for each system, they based their passwords on a common theme or domain. Since

almost half the respondents left this question blank, it can be assumed that the overall percentage of respondents using this method was even higher.

There are only 2 studies that have tried to quantify the extent to which users perform an insecure behaviour that cannot be monitored as easily as the strength of the password they choose without resorting to questionnaires. The first study (Greening (1996)) was carried out at the University of Sydney, where 338 computer science students were emailed with a false request to supply their password so that the password database could be validated after a break-in. 138 out of the 338 students returned a valid password. In the second study (Leyden (2003)), a researcher approached 152 office workers at Waterloo Station and asked them to participate in a 'security survey' in exchange for a cheap pen. 90% of the participants disclosed their password. However, the researcher was obviously not able to validate whether participants had been honest when giving a password, so the results of this study have to be taken with caution.

In summary, it can be said that the studies presented in this section, together with the anecdotal evidence, make it possible to state with a sufficient degree of confidence that users *do* regularly behave in a way that undermines security. However, it would be highly desirable for more studies to be carried out so that these claims can be substantiated in a quantitative manner. This holds particularly true for behaviours such as sharing passwords with third parties, which cannot easily be monitored. Apart from such quantitative studies, it would also be desirable to have more information on the specific insecure behaviours users choose. As an example, users writing down their password might do so on a post-it attached to their monitor, on a piece of paper locked away in a drawer or on a PDA that itself is password-protected. These behaviours differ in the extent to which they compromise security, and they might require differing actions on behalf of the organisation in order to stop them or reduce the security risk incurred by them.

2.4.4 Usability of computer security

2.4.4.1 Usability of computer security in general

The majority of security research that takes human issues into consideration has focussed on improving the usability of security mechanisms. The underlying assumption is that users primarily behave in a manner that undermines the security of the resources in question because the security mechanism is not usable (Schultz et al.

(2001)). Zurko & Simon (1996) propose three approaches that can be employed to achieve the goal of creating usable security:

1. Established procedures for enhancing usability can be applied to developing or existing security mechanisms (e.g. Karat (1989), Mosteller & Ballas (1989)).
2. Security models and mechanisms can be developed for and integrated into software that is already known to possess a high degree of usability (e.g. Foley & Jacob (1995), Shen & Dewan (1992)).
3. Usability can be considered as a primary design goal at the start of developing security mechanisms (e.g. Holmstroem (1999), Jendricke & Markotten (2000)).

Zurko & Simon (1996) consider the third approach, which they have named 'user-centered security', the most promising area for future work, and were themselves the first to design a security mechanism with both security and usability as peer goals (Zurko et al. (1999)). They state that Adage, an authorisation service for distributed applications, has been designed with a user-centered security approach from the outset, but are unclear about the exact techniques they have employed. They have tested the early results of the design process using contextual interviews, a think-aloud protocol and a questionnaire. This showed that novice users were able to perform a number of basic tasks in an acceptable timeframe, without the need for documentation.

The work on user-centered security that has been initiated by Zurko & Simon (1996) is characterized by the application of established usability design and evaluation techniques to security mechanisms. This approach has been questioned by Whitten and Tygar (Whitten & Tygar (1998), Whitten & Tygar (1999)), who argue that security mechanisms require a usability standard that is different from that applied to other types of consumer software. They reason that the problems their target group, home users, experience with security software will only be solved sufficiently by improved user interface design techniques. The alternatives they envision, namely legal remedies, increased automation and user training seminars are seen by them to provide only limited solutions. To show that standard user interface design techniques do not result in security software that satisfies their usability standard, they perform a usability evaluation of PGP 5.0 (Garfinkel (1995)), which they regard as a representative example of general user interface design techniques being applied to security software. They perform a cognitive walkthrough, together with a heuristic evaluation, and a user test, and conclude that PGP 5.0 is unusable for users who are not already knowledgeable in the area of security. In their view, this means that domain-specific user interface design principles and techniques will need to be

created for security software. However, this argument is at best questionable. The natural conclusion would be that PGP 5.0 has not been well-designed with respect to its usability, and needs to be improved further.

The authors also claim that the need for domain-specific design guidelines stems from the fact that security possesses a number of problematic properties:

1. “The unmotivated user”
Security is usually a secondary goal, which can get ignored when users focus on their primary goals. It is easy to put off learning about security, or to assume in an optimistic fashion that it is working, in order to avoid the effort necessary to perform security-related actions.
2. “Abstraction”
The management of computer security often involves security policies, which are systems of abstract rules. Ordinary users might find it difficult and unintuitive to manage such rules.
3. “Lack of feedback”
It is difficult to provide good feedback to users about the state of a security configuration, which is usually complex and often can’t be summarized adequately.
4. “Barn door”
It is impossible to be sure that a secret that has been left unprotected even for a very short time has not already been accessed by an attacker. This means that the user interface design has to make sure that users understand their security well enough to avoid potentially high-cost mistakes.
5. “The weakest link”
Users need to be guided to attend to all aspects of their security, since a weakness of even a single one of these can be exploited by an attacker.

The merit of the work by Whitten and Tygar lies in its pointing out that security possesses problematic properties which make it unlikely that the application of usability techniques alone will solve all the problems users experience with security software. However, a critical analysis of the properties they list shows that most of them can be addressed using a standard usability approach, whereas others do not necessarily precipitate the need for a new definition of usability, and new usability guidelines. Instead, they can be addressed by other means. The definition of usability for security Whitten and Tygar offer states that security software is usable if the people who are expected to use it

1. “are reliably made aware of the security tasks they need to perform;
2. are able to figure out how to successfully perform those tasks;
3. don’t make dangerous errors; and
4. are sufficiently comfortable with the interface to continue using it.”

The second and third point of this definition can most certainly be addressed with standard usability techniques. The first point can be satisfied in part by improved

usability, and in part by user training, e.g. in the form of online tutorials. The final point is important, and it goes beyond what is traditionally achieved by usability-based approaches. In effect, it states that for security software to be useful, it does not only have to be usable, but also needs to be used (to a certain extent, the first point raises the same issue). This is closely connected to the “unmotivated user” property Whitten and Tygar have identified. While the other four properties they mention can be addressed by traditional usability techniques, this one cannot. However, this does not necessarily mean that the definition of usability should be extended. Instead, this issue can be addressed by other approaches, which go beyond usability. These approaches have to take into account that security almost always is a secondary goal for users, which incurs additional overheads and might get into the way of a primary goal they want to achieve. They will have to find ways in which users can be persuaded to use the security functionality at their disposal, and to use it in the proper manner.

The “unmotivated user” property has not been identified explicitly by Dufft et al. (1999), but is being addressed implicitly by proposing that security software should be designed not only to be usable, but also to be likeable. The authors suggest to design a total user experience, which among other things increases the general positive appeal of the security mechanism and the social prestige associated with using it. This mirrors a recent trend in the field of human-computer interaction, where the focus is on moving beyond providing usability, and towards creating total, and pleasurable user experiences (e.g. Jordan (2000)).

The paper by Dufft et al. (1999) also points out another problematic property of security, which is not covered by Whitten and Tygar. Security is never positive in itself, i.e. its only positive aspect is a lack of negative consequences. This means that making the effort to use security mechanisms in the proper manner is rarely ever rewarded directly, whereas not using them can result in punishment after an unforeseeable amount of time. As a result, security behaviour is most often learned by negative reinforcement, which is known to be less effective than positive reinforcement. One way in which this situation could be improved is by the use of the before-mentioned pleasure-based approaches, which might make it possible to reward users immediately upon performing a behaviour that increases security.

Research on improving the usability of security mechanisms is still in its infancy. Underlying it is the assumption that users behave in a manner that undermines the

security of the resources in question because the security mechanism is not usable. This is a valid point, and it can only be hoped that it will be taken up by researchers both from the security and the human-computer interaction community. However, the first attempts at improving the usability of existing security mechanisms have also shown that security possesses a number of properties that raise additional issues for the developer of security mechanisms. Not all of these issues can be addressed just by integrating security and usability. Instead, usability is one building block towards an integrated solution, which incorporates further measures that are needed to ensure that users behave in the proper manner. The basis of such an integrated solution will be the identification of all the factors that determine user behaviour, and of the way in which these factors interact.

2.4.4.2 Usability of password security

2.4.4.2.1 Password memorability

The users of password mechanisms are required to memorise and recall their password without resorting to the help of a written copy. An inability to perform this task can lead to two problems. Firstly, users might not be able to recall their password when trying to log into a system. As a result, they will have to contact the systems administrator and have the password reset, which is likely to cost them a significant amount of time and the organisation a significant amount of money. Secondly, users may be tempted to avoid the cost potentially incurred by forgetting their password(s), either by choosing weak yet memorable passwords, or by writing their passwords down. This immediately undermines security, since weak passwords are easier to guess, whereas written copies of passwords can be found and exploited by attackers of computer systems. Research to date on the usability of password mechanisms has focussed almost exclusively on this problem of password memorability.

The majority of research in this area has focussed on which type of password content is easier to memorise and has been conducted in an experimental setting rather than an operational one. A good example of this is the work undertaken by Zviran & Haga (1993). 103 graduate students were given 3 passwords, each generated by one of the following methods:

1. System-generated (alphanumeric);
2. System-generated (pronounceable);
3. Self-generated.

In addition to these, they were given a pass-phrase and tens of challenge-response pairs. They were asked to recall the passwords after three months. The system-generated alphanumeric passwords, which represent the cryptographically strongest password content, were most difficult to remember: 13% could be recalled, but only by resorting to a written copy of them. The system-generated pronounceable passwords could be recalled by 37% (83% of which were reported from memory). Finally, the self-generated passwords could be recalled by 27.2% of participants (42.9% of which achieved this from memory only). The difference between alphanumeric and pronounceable passwords was statistically significant.

There are a number of other studies of this type which have been carried out (e.g. Bunnell et al. (1997), Spector & Ginzberg (1994)). The main characteristic all these studies have in common is that they investigate the ability of users to memorise different types of password content in an experimental setting, with no interference from other tasks they would have to perform in an operational setting. In particular, issues such as the fact that users typically have to memorise not just one, but a number of passwords and that they have to change these at regular intervals are not factored into these studies. This makes it difficult to determine to what extent their findings make it possible to draw conclusions about the memorability of passwords when used in an operational setting.

Yan et al. (2000) presents the only study that observed password memorability in an operational setting. Carstens et al. (2000) report a limited case study that mimics an operational setting but which is based on a false understanding of what constitutes secure password content). 288 first year undergraduate students were separated into three groups and given different advice on how to choose the password for the central computing facility, which they were likely to use on a regular basis:

1. The control group were given the same advice as previous years, i.e. to choose a password of at least seven characters containing at least one non-letter.
2. The random password group were shown how to choose a random password by using a sheet of paper which was provided to them and which contained the letters A-Z and the numbers 1-9 repeatedly on it.
3. The passphrase group were told to choose a password based on a mnemonic phrase.

The passwords chosen by the participants were submitted to a number of attacks by the researchers in order to determine their cryptographic strength. In addition, participants were sent an email questionnaire after four months, asking them how

difficult it had been for them to memorise their password and for how long they had needed to keep a written copy to refer to. The findings can be summarized as follows:

1. Passwords based on mnemonic phrases are no more difficult to remember than naively selected passwords, and both of these are easier to remember than randomly selected passwords.
2. Passwords based on mnemonic phrases are cryptographically as strong as randomly selected passwords, and both of these are stronger than naively selected passwords.
3. Even when given explicit instruction on how to choose their passwords, 10% of participants in the random password group and in the passphrase group failed to comply and chose passwords that were either too short or not constructed according to the instructions.

The study reported in Yan et al. (2000) is a clear improvement on previous studies in that it investigates the effect of password content on password memorability in an operational setting. However, it does not explicitly take into consideration other factors that might impact on password memorability. One such factor is the frequency with which a password is being used. Adams et al. (1997) used a web-based questionnaire on security and password-based authentication systems to collect responses from 139 respondents. They found that users reported that infrequently used passwords are the ones most often forgotten. This result was confirmed in a study by Sasse et al. (2001), who collected questionnaires from 144 British Telecom employees, asking them to describe the cause of the last password problem they encountered and the frequency with which they used the password. In the same paper, another study is reported which shows the importance of studying password memorability issues in an operational setting, in order to identify factors that might be overlooked in an experimental setting. System logs of passwords use were taken of 32 students who used a web-based system to practice and submit assessed coursework. This made it possible not only to study the frequency of logins and login failures, but also to determine the cause of those failures. In studies up to this point, login failure had always been assumed to be the result of users forgetting passwords. However, in this study it could be shown that users hardly ever draw a complete blank. Instead, the login usually fails because users only recall the password partly, or because they recall a different password than the required one, i.e. a previously used password for the same system or a password for a different system. This confusion of passwords happens more often for heavy- or medium-used passwords than it does for lightly used passwords. The authors also analyzed 6 months of password reset logs from the British Telecom password helpdesk and found that another factor that had been

overlooked in previous studies on password memorability was the effect of forced password changes: 13% of all reported password problems occurred just after changing a PIN. Finally, Sasse et al. (2001) have also pointed out how users' having to memorise multiple passwords can cause memorability issues and may lead to them choosing cryptographically weaker passwords or writing down their passwords. This will be discussed in greater detail in section 2.4.4.2.2.

Research on human issues involved in password security has focused almost exclusively on the issue of password memorability. The main issue investigated is the relationship between password content and password memorability, and this has been studied almost always in an experimental rather than an operational setting. However, the few studies performed in an operational setting have already shown that there are other factors apart from password content that impact on password memorability. These factors clearly include

1. the frequency of password use,
2. the number of passwords users have to memorise,
3. the extent to which they are forced to change their passwords at regular intervals.

More such studies need to be conducted in an operational setting to further our understanding of the relative importance and interplay of these factors. Such studies are likely to identify further factors that impact on password memorability in particular, and user behaviour with respect to password security in general.

2.4.4.2.2 Identifying factors that affect user behaviour

A common theme running through the previous 2 sections is that of the identification of new factors that influence user behaviour, even when the studies in question did not expect to discover these at the outset. In this section, the only study in this field that explicitly set out to identify such factors will be reviewed. Both in its substantive and in its methodological contribution to the field, this study has to be considered seminal and is the strongest individual influence on the work that has been carried out as part of this thesis. The work has been presented in two papers (Adams et al. (1997), Adams & Sasse (1999)), the first of which provides more technical detail of the study, whereas the second offers a more readable overview. The discussion of the study in this section draws on both papers, and will point out when an individual item can only be found in one of them. The discussion will also cover the study in its entirety in this

section, even though parts of it point beyond pure usability issues. Other sections will refer back to those parts as presented in this section.

The aim of the study was “*to identify human and organisational factors which impact on the security and usability of password systems*” (Adams et al. (1997)). The study consisted of two parts. In the first part, the answers of 139 respondents to a web-based questionnaire on the security and usability of password mechanisms were collected and analyzed. Half of the respondents came from British Telecom¹, whereas the other half were users from organisations throughout the world. The main quantitative results of this part of the study can be summarized as follows:

1. There was a significant inverse correlation between the frequency with which a password was used and the memory problems it caused – less frequently used passwords caused more memory problems.
2. 50% of the respondents reported writing down their password in one form or another.
3. All of the respondents who answered the question (50% of the total number of respondents) reported producing related passwords, i.e. all or most of their passwords were drawn from a pool that was based on a common theme or domain.

The questionnaire also contained a number of open-ended questions, and the responses to these questions suggested that there were other factors which had an impact on user behaviour and/or led to user problems. In the second part of the study, an initial set of 15 semi-structured in-depth interviews was carried out in British Telecom, covering issues of password generation and recall, as well as more general system and organisational factors. These interviews were then analyzed using grounded theory (Strauss & Corbin (1990)) to build a model of users’ password behaviour. This model was substantiated through the analysis of a further 15 in-depth interviews in Organisation B (a company in the construction sector). Four major factors that influence password usage were identified:

1. **Multiple passwords**

Many users do not have to use a single password, but multiple passwords for the different systems they use. Often, they also need to change their passwords regularly. As a result, memorability problems occur and users might resort to insecure work practices. In particular, they might write down their passwords or choose cryptographically weak passwords. A lot of users have developed their own method of creating related passwords, often by linking them (e.g. *tom1*, *tom2*, *tom3*).

¹ The original papers did not reveal British Telecom’s identity, but only referred to ‘organization A’.

2. Password content

Users' knowledge of secure password design was shown to be inadequate in the study, and led to users creating their own password design strategies, which were anything but secure.

3. Users' perceptions of organisational security and information sensitivity

Users' perception of the organisational importance of security and the threats to security was shown to be a key element in motivating their work-practices. Without feedback from their organisation, they would often create their own models of these issues, which were littered with misconceptions and led to insecure practices. Users' security behaviour was also shown to be dependent on their perception of the sensitivity of the information that was meant to be protected. Again, this perception regularly was incorrect. In particular, confidential information about individuals was regarded as sensitive, whereas commercially sensitive information often was regarded as not sensitive.

4. Perceived compatibility between password procedures and work practices

The study revealed the importance of compatibility between work practices and password procedures. British Telecom forced its users to have individually owned passwords for group working, where users would have considered shared passwords as more compatible with their work practices. This can result in users circumventing the security mechanism, e.g. by disclosing their password to other group members. In organisation B, the opposite problem occurred: users emphatically rejected the fact that they were forced to use group passwords for individual personal information, such as email.

At the highest level of the model, the study combines these 4 factors further to identify two main problems in password usage:

1. System factors

System factors comprise password restriction mechanisms, passwords for multiple applications and multiple changes over time. Users can conform with these system factors but often feel forced to circumvent them, e.g. by choosing linked passwords for different systems.

2. External factors

External factors comprise the information's perceived sensitivity or importance and enforced password practices so that passwords are perceived to be allocated to the individual or the group. These external factors can be either compatible or incompatible with users' perceptions.

Adams & Sasse (1999) conclude that users do not necessarily behave in an insecure manner because they lack the motivation to behave in a security-conscious fashion. Instead, system factors can force them to undermine security procedures. External factors that are incompatible with users' perceptions can have the same effect. Adams & Sasse (1999) recommend a number of measures to be taken by organisations to improve user behaviour, among them improved user training, changes to the password mechanism that increase its usability (e.g. the introduction of single sign-on) and a higher degree of compatibility between work practices and password procedures.

The work reported in Adams & Sasse (1999) and Adams et al. (1997) is seminal in that it is the first study that explicitly set out to identify the factors that have an impact on users' security-related behaviour. The methodology chosen to do so (grounded theory) is a valid addition to the field and seems to be particularly suited to the problem in that it makes it possible to uncover and descriptively relate a complex web of variables while reducing any bias possibly introduced by the researcher. The strength of the study lies in its systematic identification of factors that influence users' password behaviour, such as the perceived importance of the information protected. The combination of these factors into higher-level factors is overall sound but not always completely convincing. In particular, the membership criteria for the high-level 'system factors' and 'external factors' are not made explicit, and seem to be somewhat blurred (e.g. enforced password practice with respect to group practices is an external factor, whereas password restriction mechanisms are a system factor). In addition, individual factors are being identified, but their actual interplay is not determined to a sufficient extent. As a result, the story-line that combines the different factors could easily be split into a separate story-line for each individual factor, without any information being lost.

The overall "message" of the work - that users are not necessarily the enemy of the organisation (or security department) - is important, but does not seem to have been questioned and qualified seriously enough. It is based on the analysis of interview data which has been taken at face value. Adams et al. (1997) report how several interviewees made contradictory statements in the course of a single interview, and claim that the grounded theory analysis showed that this was due to a complex issue involving several factors being discussed. They give an example of such an apparent contradiction: Users perceiving the organisation's general security level as low will decrease their perception of how sensitive the information protected is. This, in turn, increases insecure password practices. On the other hand, users who perceive the organisation's general security level as high will decrease their overall perception of threats to the information. This will also increase insecure work practices. This analysis makes sense as long as the statements of the interviewees are taken at face value. However, it is possible that some of the interviewees have structured their discourse about password security in a manner which makes it possible for them to justify their improper password practices.

In summary, the study discussed in this section has added significantly to both the substantial and the methodological knowledge in this field. It points in the right direction for further research, which should aim to achieve the following:

1. The factors that have been identified as affecting users' password behaviour need to be validated further.
2. Additional factors are likely to exist and need to be identified.
3. The interplay between the various factors needs to be identified.
4. A higher-level model needs to be developed which incorporates all these factors into a comprehensive story-line and which makes explicit the interplay between the individual factors.
5. The possibility of users possibly structuring their discourse in order to justify their improper password practices needs to be addressed.

2.4.5 User knowledge and skills

The work by Adams & Sasse (1999) has substantiated a point that should be self-evident: the knowledge users have of security issues and the skills that they possess in this area will strongly affect the extent to which they behave in a secure fashion. More importantly, the study conducted by Adams & Sasse (1999) has also begun to identify specific gaps in the knowledge and skill base of users which need to be addressed. In the absence of any organisational efforts to fill these gaps, users will construct their own, often wildly inaccurate models of issues such as the importance of security or the best way to construct a password. This almost invariably leads to insecure security practices.

The way in which organisations typically try to fill the gaps in users' knowledge and skills is by means of security awareness and training campaigns. The need for such campaigns is generally accepted (e.g. Siponen (2001)), but there has only been a very limited amount of research in this area. Published research consists almost entirely of papers which do not present any empirical results, but only suggest ways in which security awareness campaigns should be designed. A typical example of this is Thomson & Solms (1997). The paper starts out by stressing the importance of security awareness campaigns, and goes on to describe a general approach towards developing a security awareness program, which is based on suggestions taken from a NIST handbook on computer security (National Institute of Standards and Technology (1995)). This contains a number of points of a general nature, such as the identification of the program scope, goals and objectives or the identification of the target audience. The paper then suggests specific content for the three main target groups that have been identified, namely the top management, the IS management and

the end-users. Again, the paper draws on other sources for this information, and keeps it mostly at a general and vague level. As an example, the paper advises to make top-management aware of information security terminology and standards. Where the paper does get more specific is with respect to the content end-users should be educated about, but even here only a short sample list of possible subjects is given. On the whole, the paper ends up as a collection of general and vague material taken from other sources and provides a high-level description of a methodology to design security awareness campaigns. It is difficult to determine the effectiveness of such a methodology without any empirical results about its use in an operational or experimental setting. Other papers suggest that methods from social psychology (e.g. Siponen (2000a)) or marketing (e.g. McLean (1992)) should be used to make security awareness programs more effective, but yet again no empirical validation of these claims is provided. The only paper which describes an actual security awareness campaign and its effect on users is of an anecdotal nature, and does not provide any hard data on the effect of the campaign (Spurling (1995)).

There is a clear need for research in this area, which should not only aim to develop a high-level methodology that can be used to design security awareness and education campaigns, but which should also try to provide domain-specific knowledge about specific security systems and user groups which can be used to design the content of such campaigns. As importantly, the claims made in such research need to be substantiated with quantitative data that has been gained in an experimental fashion. In addition, the possibility of means other than security awareness and education campaigns to instil knowledge and skills in users has to be investigated.

2.4.6 Motivation

The work by Adams & Sasse (1999) has shown how users' knowledge and skills can impact on their motivation to behave in a security-conscious fashion. As an example, a user who does not believe the information that is protected by the password mechanism to be at risk of being targeted by third parties will be less motivated to behave in a secure manner than one who does. This could make it seem as if users who are properly educated about security issues should be motivated to behave in the proper manner. However, there is a larger issue that needs to be considered. In most cases, authentication to a system is an *enabling task*, which means it creates an overhead for the user, who is using that system as a tool to achieve a *primary*, real-

world *task*. It is predictable that most users will cut corners to reduce that extra load given a chance, unless they are *motivated* to make the effort to behave in a security-conscious fashion (as discussed in section 2.4.4.1, Whitten & Tygar (1999) have put forward this argument for security mechanisms in general). Oversimplifying for the sake of argument, users of password mechanisms can be divided into two groups: those that face personal damage if they do not behave in a security-conscious fashion, and those that do not put themselves, but others, at risk by cutting corners. Self-employed and home users fall into the first category – users in this group can, if educated about the possible consequences of their behaviour, make an informed choice about their behaviour, based on an assessment of the risks and the effort required to reduce these risks. Users in an organisational context fall into the second category, and for them education alone might not be sufficient to increase their motivation to the level where they will make the extra effort needed to behave in a security-conscious manner. This is a point that has been stressed by Parker (1998), who suggests removing the conflicts between job performance and security constraints by making security a part of job performance. He puts forward the use of rewards and penalties in annual job performance reviews as a motivator for end-users in organisations and calls this “*the mother of all security controls*”.

The importance of motivating users to make the extra effort that is needed to behave in a security-conscious fashion has also been stressed by the advocates of security awareness and education campaigns. However, this has again only resulted in a number of recommendations on how to use techniques taken from social psychology (e.g. Kabay (1993)) or marketing (e.g. Spurling (1995)) in order to obtain users’ commitment to proper security behaviours. Once more, the resulting recommendations are general in nature and lack domain-specific knowledge about specific security systems and user groups. In addition, none of them have been tested experimentally.

It is clear that more research is needed in this area. Such research should first of all study the effect of user motivation on security-related behaviours further. It then needs to identify the factors that affect user motivation. The results gained in this way need to be substantiated with quantitative data that has been gained in an experimental fashion.

2.4.7 Investigating the wider organisational context

The majority of the research discussed in the previous three sections (2.4.4 to 2.4.6) has assumed an individual factor to be largely responsible for users' behaving in an insecure fashion and has then proposed ways in which this factor can be altered in order to improve user behaviour. The main exception to this was the work presented in Adams & Sasse (1999), which did not *presuppose* any factors but rather tried to *identify* them from scratch. One factor that this work identified was an often poor fit between password practices and work practices with respect to group work: users were forced to use shared passwords where individual passwords would have been more appropriate, and vice versa. Once such a factor has been identified, an immediate reaction is to make changes that lead to improved password practices, in this case by tying in password practices with work practices in a more appropriate fashion. However, there is another question that can be asked: how did the poor fit between the two come about in the first place? In more abstract terms, this boils down to the question of how certain factors that affect user behaviour have come to be allowed to affect it *negatively*. This can be rephrased again as the question of what the factors are that affect the factors that affect user behaviour in a direct fashion. This question has been addressed in an informal manner in some of the research efforts presented previously (e.g. the strong roots of security research in the military, as discussed in Zurko & Simon (1996)). In this section, the only research effort in the field that has tried to answer this question in a systematic manner will be presented.

Brostoff & Sasse (2001) starts out with the argument that almost all security systems involve human users as well as technology, and should therefore be designed as socio-technical work systems. It is then proposed that safety-critical systems design has goals and issues that are similar to those encountered in security design, and should therefore provide a good starting point for this endeavour. Reason's (Reason (1990)) *Generic Error Modelling System* is suggested as the most suitable candidate for such a socio-technical approach (Spruit (1998) has used parts of this model in a less convincing fashion in an earlier paper). The paper, which is explicitly flagged as a position paper, then briefly presents Reason's model, discusses its suitability for the task at hand, gives an example of it being applied to a problem in the security arena and finishes with a discussion of the advantages and disadvantages of using the model.

Brostoff & Sasse (2001) argue that Reason's model makes it possible to capture failures at the individual and at the organisational level in a systematic fashion. At the individual level (i.e. end-users in large corporations), the model posits three kinds of human error, which together fall under the category of *active failures*. *Slips* (attentional failures) and *lapses* (memory failures) are unintended actions that lead to a bad result. *Mistakes* (rule-based or knowledge-based mistakes) are intended actions that lead to an unintended result. Together with *violations*, which are actions intended by the user but not by other people, these form the class of *unsafe acts* (re-labeled by the authors as *insecure acts* in the domain of information security). In order for a security breach to occur, *insecure acts* have to combine with *latent failures* and/or unusual environmental conditions. *Latent failures* can be thought of as weaknesses that are built into the system and predispose it to security breaches. They act by promoting *insecure acts* and by weakening the system's defenses. Reason's model describes a system or an organisation as consisting of a number of levels: *decision-makers*, *line managers*, *preconditions* (e.g. reliable equipment or a motivated workforce), *productive activities* and *defences* (e.g. firewalls). The causes of a disaster can be traced back to failures at all levels in this model, with deficiencies at a higher level being transformed to deficiencies at a lower level. One central assertion of the model is that removing latent failures in higher parts of the system will provide greater benefits than removing errors at a lower level.

The paper goes on by giving an example of the application of the model to a specific example of a security breach, in this case the assumed theft of a laptop with confidential information on it that has not been encrypted even though the relevant software for doing so was available to the user. The authors consider each of the insecure acts and use them to label the user's behaviour, and then identify the latent failures in higher parts of the system that could have led to those insecure acts. The authors state that this application of the model has made it possible to identify several potential causes of or contributory factors to the security breach. The authors also point out that, in a situation similar to the one the safety-critical community found itself in 10 years earlier, the application of the model makes it clear that end-users usually are not the main instigators of a security breach, but rather the inheritors of system defects that are caused by failures higher up in the organisation.

The work presented in Brostoff & Sasse (2001) is vitally important for the whole field of computer security in that it is the first to propose a way in which the larger

organisational context can be taken into consideration when trying to reduce the number of security breaches. As importantly, it provides a systematic way in which insecure user behaviours can be classified and in which the larger organisational failures that led to those behaviours can be identified. This should make it easier to identify the interventions that will have the strongest positive effect on overall security. However, the paper is a position paper only, and as such its claims need to be verified through empirical studies before a final verdict on their validity can be made. In particular, it has to be seen whether a model that has been developed for safety-critical systems - where errors can lead to life-threatening situations, and operators can probably be assumed to be intrinsically motivated to behave in the proper fashion - can be applied to security systems, where end-user motivation might be less strong. Safety and security also differ in that the former does not need to consider adversaries that actively seek to attack, which might have an influence on users' attitudes and behaviour. In addition, the model is not easy to operationalise, as the authors willingly admit when stating that "*the quality of its application depends on the expertise of the people who apply it within a particular organisation*". This can be explained by the fact that the model is of a high-level, abstract nature and relies in its application on additional domain-specific knowledge about the security system and the organisation in question.

2.4.8 Interim summary

2.4.8.1 Previous research

The central research problem this thesis tries to address is the question of how the likelihood of users in an organisation performing the expected password-related behaviours can be increased in a situation where some of these behaviours can be neither enforced, nor monitored adequately (see section 2.3.5.2). Previous research that is related to this problem has been presented and discussed in the preceding sections (2.4.3 to 2.4.7). The amount of research on human factors in computer security to date is limited, and there is even less research on the human factors in password security. Moreover, the field does not yet have an agreed research agenda, and the various contributing researchers are at times unaware of each other's efforts. The discussion of the previous research has also shown up a number of issues, the main ones of which can be summarised as follows:

1. *Anecdotal nature of some of the claims being made*
A considerable number of the claims that are made in this field are of an anecdotal nature and have not yet been substantiated with quantitative data. This holds particularly true for claims about actual user behaviour.
2. *Assumption of certain factors affecting user behaviour*
A large part of the research that has been discussed has assumed one or several factors to have a strong influence on user behaviour and has then tried to find ways in which these factors can be changed in order to improve that behaviour. With the exception of the work presented in Adams et al. (1997) and Adams & Sasse (1999), there have been no efforts to identify the factors that affect user behaviour from scratch, without any prior assumptions.
3. *Proposals of solutions without any proof of effectiveness*
Some of the research has suggested ways in which user behaviour can be improved by borrowing ideas and concepts from other fields, such as social psychology, and suggesting they be applied to computer security. However, none of these suggestions have actually been tested in a manner that would make it possible to determine their effectiveness in this particular area in a quantifiable fashion. As importantly, the suggestions that have been made are often so vague that it would be impossible to devise a way of testing them without resorting to additional, domain-specific knowledge.
4. *Lack of domain-specific knowledge*
There has been very little research that has generated domain-specific knowledge about specific security mechanisms and specific user groups. The majority of research that has done so has focussed on only one factor affecting user behaviour, namely the usability of certain security mechanisms.
5. *Lack of awareness of the wider organisational context*
There has to date only been one research effort that has tried to identify and address the causes of insecure user behaviour that can be found in the wider organisational context (Brostoff & Sasse (2001)).

2.4.8.2 The research questions

The previous section has summarized five issues that have been identified with respect to previous research that has been carried out in this field. The issue that the research presented in this thesis primarily tries to resolve is the second, i.e. the lack of any research that has tried to identify the factors that affect user behaviour without making any prior assumptions about them. This translates into the fundamental research questions that this thesis tries to address:

1. What are the factors that affect the password behaviour of users in organisations?
2. How do these factors interact in order to cause specific behaviours?
3. How can knowledge about these factors and their interplay be used to improve user behaviour?

It is important to point out that these questions are meant to focus on those factors that directly affect user behaviour and do not take into consideration the factors that are further removed and might affect user behaviour indirectly through their effect on other factors. As an example, the questions aim to uncover factors such as the

usability of the password mechanism, but do not target any factor that is responsible for the poor usability of existing mechanisms and likely to be part of the wider organisational context. Identifying these ‘first-level’ factors and their interaction would also help to resolve the fourth issue discussed in the previous section by providing domain-specific knowledge about one particular security mechanism, in this case the password mechanism.

The research problem and the research questions as stated so far were not conceived at the beginning of the research process that led to the findings presented in this thesis, but are the result of a re-conceptualisation of the original research approach in the course of this process (see chapter 4). The data collected up to the point of this re-conceptualisation, and the methodology that was used to analyse it, also made it possible to provide partial answers to two further questions. The first is that of the true extent of users’ insecure password practices and of the specific insecure behaviours that users perform instead of the expected behaviours. The second is that of users possibly structuring their discourse about password security issues in order to justify their insecure password practices.

2.5 Chapter summary: The research agenda

This chapter has presented the background information that is necessary to understand the research presented in this thesis. Human-computer interaction has been presented as the research discipline within which the work is situated and a number of extensions to the traditional scope of this discipline have been suggested in order to enable the work presented here. It could be shown that for password security to be an effective way of protecting computing resources, users are required to perform a number of tasks (section 2.3.5.1). The research problem has been identified as the question of how the likelihood of users in an organisation performing the corresponding behaviours can be increased in a situation where some of these behaviours can be neither enforced, nor monitored adequately (section 2.3.5.2). Previous research that has tried to tackle this problem has been reviewed and its substantial and methodological contributions have been discussed. The weaknesses of this research have been pointed out (section 2.4.8.1), and it has been shown how they have motivated the research questions that have directed the research documented in this thesis (section 2.4.8.2):

1. What are the factors that affect the password behaviour of users in organisations?

2. How do these factors interact in order to cause specific behaviours?
3. How can knowledge about these factors and their interplay be used to improve user behaviour?

In addition, it has been pointed out that the research problem and the research questions as stated so far were not conceived at the beginning of the research process that led to the findings presented in this thesis, but are the result of a re-conceptualisation of the original research approach in the course of this process (see chapter 4). The data collected up to the point of this re-conceptualisation, and the methodology that was used to analyse it, also made it possible to provide partial answers to two further questions. The first is that of the true extent of users' insecure password practices and of the specific insecure behaviours that users perform instead of the expected behaviours. The second is that of users possibly structuring their discourse about password security issues in order to justify their insecure password practices.

The following chapter (chapter 3) will describe the research approach that was taken to implement this research agenda.

3 METHODOLOGY

3.1 Overview

Chapter 2 has developed the research problem and the research questions that have driven the work documented in this thesis. This chapter will introduce the methodologies that were used to address this problem and these questions. Background information for each methodology is given, and its use to approach specific parts of the research documented here is justified. Finally, the manner in which the methodology is applied is described. The summary at the end of the chapter (section 3.6) will recapitulate the way in which the different methodologies have been used in conjunction to create the specific research approach employed in this thesis. The following chapter (chapter 4) will describe the studies that have been performed on the basis of this research approach.

3.2 Data collection techniques

3.2.1 Questionnaires

3.2.1.1 Justification

A questionnaire was used in studies 2, 5 and 6 to assess the effect of different versions of a fear appeal (see section 3.5) on distinct groups of recipients (see chapter 4). This made it possible to collect large amounts of data in a relatively short timeframe. The use of closed questions in the questionnaire resulted in data that was easy to analyse and compare between groups using standard statistical techniques.

3.2.1.2 Application

Questionnaires make it possible to gather large amounts of data in a short timeframe. However, for them to be useful research tools the researcher has to ensure that they are well-designed (Dix et al. (1998)). The first question that should be answered is that of the purpose of the questionnaire, which will determine the information that is sought. It is also useful at this stage to decide how the questionnaire responses are going to be analysed. The next question then will be what types of question structure will be used (Preece et al. (1994)). *Open-ended* questions give respondents the freedom to provide their own answers and provide a rich source of data that can be difficult to analyse. *Closed* questions restrict the possible answers respondents can give to a selection of alternative replies, which may be provided in a number of ways.

Scalar questions, for example, ask the respondent to judge a specific statement on a numeric scale. *Multiple-choice questions*, on the other hand, offer the respondent a choice of explicit responses, and will ask him to select one or several of these. Finally, *ranked questions* ask the user to put an ordering on items in a list. Closed questions of this type will result in data that is easier to analyse and quantify than that obtained by open questions. Once the question structures that are to be used have been determined, the wording of the actual questions has to be chosen carefully, since even minor changes in it can alter respondents' answers (e.g. Loftus (1975)). Finally, the issue of sampling has to be considered carefully in order to avoid biased results.

3.2.2 Semi-structured in-depth interviews

3.2.2.1 Justification

Semi-structured in-depth interviews were used in study 1 to gain a detailed understanding of the participants' points of view with respect to password issues. The semi-structured agenda ensured that certain issues were covered in all interviews, whereas the in-depth nature of the interviews made it possible for other issues to be brought up by the interviewees and to be followed up in the course of the session.

3.2.2.2 Application

In-depth interviews make it possible to gain a detailed understanding of individual participants' points of view. Semi-structured elements used within the interview procedure give the interviewer the ability to maintain control about the direction that is being taken and can also increase the reliability of the data obtained and the speed with which this data can be analysed (Cooligan (1990)). At the same time, the interview needs to be kept flexible enough to make it possible to pursue key issues that are introduced by the interviewee into the conversation. The interview usually starts with the initial introductions, after which the interviewee's permission to record the session will be obtained and assurances about the confidentiality of the data will be given by the interviewer. After this, the interviewee will be asked some questions about his background, which in itself can provide valuable data but also has the function of getting the conversation started and putting the interviewee at ease. This is followed by the main part of the interview, in which the interviewer will ask the questions that he has determined beforehand to be relevant and will decide to follow up certain issues that the interviewee has brought up and which were not part of the

original interview plan. Towards the end of the interview, the interviewer will typically sum up the main issues that have been discussed in order to ensure that any necessary clarifications can be undertaken during the interview time that remains. Finally, the interviewer might debrief the interviewee about the purposes of the research and the ways in which the interview data will be used.

3.2.3 Focus groups

3.2.3.1 Justification

Focus groups, rather than interviews, were used as the data collection tool as of study 3 for three principal reasons:

1. A number of hypotheses that were slowly emerging as a result of the analysis of the data that was being collected needed to be both validated and refined. This led to a set of questions that was more specific and less tentative and exploratory than the ones used in the interviews of study 1. These questions could have been asked in interviews, but focus groups made it possible to gather the opinions of a larger number of people in the same amount of time that an interview (and its subsequent transcription) would have taken.
2. Focus groups provide a social occasion that allows for public opinion to develop through debate as in real-world situations (Lunt & Livingstone (1996)). This makes them an ideal setting in which to gather participants' opinions about scenarios that contain changes to the password mechanism (e.g. single-sign-on) or fictitious appeals by the organisation to improve password behaviour (e.g. fear appeals) that cannot be tested in real life.
3. Focus groups strongly promote interaction among participants and encourage the development of ideas that are important to them. This made it possible to identify and follow up issues that were not yet covered by the hypotheses that had been created.

3.2.3.2 Application

Focus groups typically consist of 4 to 10 individuals who come together to discuss a topic under the direction of a moderator. The aim of the moderator is to create a genuine and relaxed setting for the focus group, which provides a social occasion that allows for public opinion to develop through debate as in real-world situations (Lunt & Livingstone (1996)). The moderator will promote interaction among the participants while ensuring that the focus group agenda is adhered to. This agenda is designed to ensure both that the issues the moderator wants to discuss are covered in the focus group and that the development of ideas that are important to participants is encouraged. Focus groups usually start with the initial introductions, during which permission to record the session will be asked of the participants, who will also be assured about the confidentiality of the data. This is followed by a short round-robin

question, which typically pertains to the participants' background and is a way of getting everybody talking, thus encouraging the contribution of all members of the group. During the main part of the focus group, the moderator will ask a number of questions and encourage discussion of these among the group. Towards the end of the focus group, the moderator will typically sum up the main issues that have been discussed in order to ensure that any necessary clarifications can be undertaken during the remaining focus group time. Finally, the moderator might debrief the participants about the purposes of the research and the ways in which the focus group data will be used.

3.3 Grounded Theory

3.3.1 Justification

Grounded theory is an established social science methodology that provides a focussed and structured approach for the collection and analysis of data with the aim of creating empirically-based theory. It was originally conceived by Glaser and Strauss as the product of the close inspection and analysis of qualitative data (Glaser & Strauss (1967)), but was later developed further by Strauss and Corbin to incorporate the use of quantitative data (Strauss & Corbin (1998)). The latter also specify further what is meant when they use the term grounded theory by stating that this refers to theory

“...that was derived from data, systematically gathered and analyzed through the research process. In this method, data collection, analysis and eventual theory stand in close relationship to one another. A researcher does not begin a project with a preconceived theory in mind (...). Rather, the researcher begins with an area of study and allows the theory to emerge from the data.” (Strauss & Corbin (1998))

This grounding of concepts in data is the main feature of the method and is achieved by analysing the data in a standard grounded theory format. It is this focussed and structured approach which sets grounded theory apart from most other social science methodologies. The citation also makes clear another important feature of grounded theory, namely that no prior hypothesis is needed in order to focus the analytic process (Strauss et al. (1964)). This makes the method particularly suitable for complex subjects or phenomena about which little is known (Strauss & Corbin (1998)). The successful application of the methodology is assessed in terms of the final account's comprehensiveness and fit with the data.

Grounded theory was used in the work documented in this thesis for a number of reasons:

1. Grounded theory is ideally suited to the investigation of complex high-level phenomena about which little is known and makes it possible to create a theory of such phenomena in a systematic and empirically-based manner. Such theories are usually presented in the form of story-lines (see section 3.3.2), which makes it easy to relate them back as prescriptive knowledge to designers and other people involved in security. Grounded theory also makes it possible to tailor the data collection procedures to any time, ethical and practical constraints under which the work has to be carried out.
2. Section 2.4.8.1 has pointed out that one of the weaknesses of most research that has previously been carried out in this area has been that it has assumed certain factors to be responsible for users' password behaviour, which has led to other factors being ignored. Grounded theory makes it possible to approach data without a preconceived theory and instead allow a theory to *emerge* from the data. A conscious decision was made to use this approach in its most orthodox form by carrying out the grounded theory analysis without resorting to either the findings that had been generated by the limited amount of previous research in this field or to theoretical constructs that could be found in related disciplines. The data was to be examined without any hypothesis in mind, and only after this investigation had been completed would its findings be validated further by comparing them to theoretical constructs obtained from previous research or other disciplines.
3. Adams & Sasse (1999) have already shown that grounded theory is a suitable methodology for the specific research questions this thesis tries to address. In addition, there is a small but slowly growing number of examples of the successful application of grounded theory to problems in HCI that are not dissimilar to the one this thesis deals with (e.g. Bouch & Sasse (1999), Adams & Sasse (2001)) in that they require a method that makes it possible to create theories during the *discovery* stage of the scientific process (Henwood & Pidgeon (1992)).

3.3.2 Application

The basic operation of the grounded theory method is that of taking data, breaking it down, conceptualising it and putting it back together in new ways. Strauss & Corbin (1998) have devised 3 major coding stages (open, axial and selective) in order to ensure that this process takes place in a structured manner. However, the lines between these different forms of coding are somewhat artificial, as is the divide between data collection and analysis. In reality, all of these elements are tightly intertwined in a complex and iterative process of analysis and verification.

The first stage of the analytic process, *open coding*, begins with the identification of *concepts*, which are abstract representations of an event, object or action/interaction that the researcher identifies as being significant in the data. These concepts, which can be thought of as labelled phenomena, are then compared to determine whether

they can be grouped together under a *category*, which is a more abstract higher-order concept. In addition, subcategories are identified. These specify a category further by denoting information such as when, where, why and how a phenomenon is likely to occur. Finally, the properties (characteristics pertaining to a category) and dimensions (possible values of a property along a continuum) of the categories and subcategories are determined. Appendix A contains a list of all the grounded theory categories and properties that were identified during open coding, together with example quotes from the transcripts (which were initially coded as concepts).

The second stage of the analytic process, *axial coding*, begins the process of reassembling data that were fractured during open coding. It does this by relating categories to their subcategories along the lines of their properties and dimensions in order to form more precise and complete explanations about phenomena, integrating process with structure. Subcategories can take on the form of *conditions*, *actions/interactions* and *consequences*. *Causal conditions* represents sets of events that lead to the occurrence or development of a phenomenon. *Intervening conditions* mitigate or otherwise alter the effect of causal conditions on phenomena. *Contextual conditions* are the specific sets of both causal and intervening conditions that intersect dimensionally at a specific place and time to create the set of circumstances to which people respond through actions/interactions. *Strategic actions/interactions* are purposeful or deliberate acts that are taken by people in response to issues, problems, happenings or events that arise under the contextual conditions. *Routine actions/interactions* are more habituated ways of responding to occurrences in everyday life. *Consequences* are the outcomes of people either taking some action/interaction or of them not taking it. Table 1, page 64, shows an example of the axial coding that was undertaken during this research.

The third and final stage of the analytic process, *selective coding*, integrates and refines the theory. The core category, which is the central phenomenon around which all other categories are integrated, is chosen. The story line, a descriptive narrative about the central phenomenon of the study, is created. The subsidiary categories are related around the core category by means of its properties, and the categories are related at the dimensional level. Selective coding is an iterative process and is validated by continual comparison with the raw data to confirm or refute the conclusions that are being made. This validation process can also identify gaps that need to be filled by further research.

Axial coding example

Open coding had identified the categories **usability** and **proactive password checking mechanism** (see appendix A). During axial coding, the two categories were related to each other on the basis of repeated occurrences of statements such as the following:

- “It does not let me choose any password I want, it makes me, forces me to choose a really obscure one [proactive password checking mechanism]. It would be much easier to use [usability] if I could choose my own password.”
- “and I don’t think that forcing me to choose passwords with, ahm, with numbers and symbols in them [proactive password checking mechanism] makes it any more usable [usability].”
- “A system that forces me to spend ages on creating a password that satisfies the criteria [proactive password checking mechanism] is just not very user-friendly [usability].”

Statements such as these show that the existence of a **proactive password checking mechanism** reduces the **usability** of the system in question.

Table 1: Axial coding example

Once the selective coding stage is finished, it is possible to take the analysis one step further by integrating *process effects*, which describe the sequences of actions/interactions which can be traced back to changes in structural conditions and which might themselves change these structural conditions, possibly resulting in further actions/interactions.

3.4 Discourse Analysis

3.4.1 Justification

The term *discourse analysis* has been used by researchers from a wide variety of disciplines (e.g. psychology, sociology and linguistics) to label their specific research approaches, which often differ wildly in their philosophical framework, emphasis and levels and styles of analysis (Burman & Parker (1993)). This situation can be bewildering for a researcher who simply wants to pick up a useful set of analytical and theoretical tools, but who actually might end up with two different books on discourse analysis with no overlap in content at all (Potter & Wetherell (1987)). The specific brand of discourse analysis used for the work documented in this thesis is that of Potter and Wetherell (Potter & Wetherell (1987)), who popularised discourse analysis in social psychology in Britain from the end of the 1980s. For the remainder

of this thesis, the term ‘discourse analysis’ will carry the meaning they have assigned to it, unless stated otherwise.

Potter & Wetherell (1987) use the term ‘discourse’ in its most open sense “*to cover all forms of spoken interaction, formal and informal, and written texts of all kinds*”. They then go on to describe discourse analysis itself by introducing the interconnected concepts of *function*, *construction* and *variation*. The central tenet of discourse analysis is that function involves construction of versions, and is demonstrated by language variation. Discourse analysis argues that language constructs reality, rather than representing or reflecting it (this makes it a constructionist method). There is always more than one way to describe things, and our choice of how to describe particular aspects of reality has an immense power to shape the way we experience the world and behave in it. In other words, we construct different versions of reality, and the use of a specific version in a particular situation may serve a function. This function can go from the interpersonal (e.g. explaining, justifying, accusing) to the wider purposes the discourse might serve (e.g. creating an ideological effect in the sense of legitimating the power of one group in society). The construction of particular versions need not be deliberate or intentional, and their deployment may have repercussions of their own which may not have been formulated or even understood by the speaker. The primary method of revealing function from a study of discourse is the identification of variation in the use of language over time.

The construction of different versions of aspects of reality by individual speakers and writers draws on pre-existing and shared linguistic resources which Potter and Wetherell have named *interpretative repertoires* (Potter & Wetherell (1987), Wetherell & Potter (1988)). These form the analytic unit of discourse analysis, and can be thought of as the building blocks which speakers use to construct versions of actions, cognitive processes and other phenomena. Particular repertoires are composed of a restricted range of terms used in a specific stylistic and grammatical fashion, and will often be derived from one or more key metaphors. The notion of an interpretative repertoire makes it possible to investigate the organisation of phenomena which psychologists have traditionally understood in terms of attitudes, beliefs and attributions.

Discourse analysis has primarily been used in this thesis to identify the interpretative repertoires that participants in interviews and focus groups drew on in order to describe phenomena related to password security. This was originally inspired by a

loosely related use of discourse analysis in HCI (Rimmer et al. (1999)), but was later refined to incorporate ideas discussed in Willig (1999). Willig (1999) aims to examine critically the contribution discourse-analytic studies can make to processes of social and political change. This is motivated by a desire to move beyond critical commentary and towards active engagement with social and political practices. One way of moving towards this goal is to deconstruct discourse in order to ask how things could be different. Are there ways of talking about phenomena which are preferable? How could these be promoted? Asking these questions leads to a focus on the discursive resources people draw on rather than on the discursive practices they engage in (Potter & Wetherell (1995)). It stresses that the possibly unintentional use of specific interpretative repertoires can shape the way we experience the world and behave in it and might also have repercussions we might not have anticipated. This particular approach to discourse analysis was used in the work documented in this thesis for two main reasons:

1. The notion of interpretative repertoires that are culturally shared immediately brought up a number of intriguing questions. Is it possible that the use of some of these repertoires by individuals corresponded with their exhibiting desired behaviours (i.e. good password practices)? Could the use of other repertoires correspond with undesired behaviours? And if this is the case, would it be possible to promote the use of *resource repertoires* that correspond with good password practice and discourage that of *problem repertoires* that correspond with bad password practice? Trying to find answers to these questions seemed a worthwhile enterprise and was the initial focus of the research carried out as part of this thesis (see Chapter 4).
2. Discourse analysis takes a more critical view of the accounts people give than grounded theory, and makes it possible to investigate the possibility of discourse being constructed in order to achieve certain goals. This makes it an appropriate methodology to use when investigating whether participants in interviews or focus groups might construct their discourse in order to justify their bad password practice (see section 2.4.8.2).

Discourse analysis' approach to the collection and analysis of data is not as focussed and structured as that of grounded theory (see section 3.3.2). This was one of the reasons why the data collected for this research was analysed using grounded theory after a first discourse-analytic study had been carried out (see chapter 4). The results of the discourse-analytic study are still included in this thesis (chapter 8), since they do not only constitute a valid substantive contribution but also make it possible to derive a first evaluation of discourse analysis' applicability to problems within the field of HCI.

Rimmer et al. (1999) is the only example that has employed Potter and Wetherell's brand of discourse analysis to an HCI problem. The use of discourse analysis in this thesis and the subsequent evaluation of its suitability to problems within HCI is a methodological contribution of the work documented here.

3.4.2 Application

Potter & Wetherell (1987) stress the fact that there is no *method* to discourse analysis in the way we think of an experimental method or content analysis method. Instead, there is a broad theoretical framework concerning the nature of discourse, "*along with a set of suggestions about how discourse can best be studied and how others can be convinced findings are genuine*". However, they do suggest ten stages which a discourse-analytic study should go through, but qualify this strongly by calling these ten stages a springboard rather than a template. These stages are also not clear sequential steps, "but phases which merge together in an order which may vary considerably". The following is a list of the ten stages they propose, some of which will then be discussed in greater detail:

1. Formulation of research questions
2. Sample selection
3. Collection of records and documents
4. Interviews
5. Transcription
6. Coding
7. Analysis
8. Validation
9. The report
10. Application

The **coding stage** in discourse analysis is quite distinct from the subsequent analysis stage. It is an analytic preliminary which prepares the way for a much more intensive study of the material and its main goal is not to find results but to turn an unwieldy body of discourse into manageable chunks by collecting together instances for examination. This pragmatic goal means that coding should be done as *inclusively* as possible, with all borderline cases and instances that seem initially only vaguely related being included.

The **analysis stage** is principally made up of two closely related steps. The first step involves the search for pattern in the data, which will be in the form of either variability (differences in either the form or content of accounts) or consistency (features shared by accounts). The second phase consists of forming hypotheses about

the function and effects of people's talk. In the discourse-analytic study documented in this thesis, the search for pattern focussed on the identification of interpretative repertoires and was followed by the formulation of hypotheses about the effect that the use of specific repertoires by individuals has on their password behaviour. A further examination of the material then investigated the question of whether the choice of specific repertoires or discursive practices by individuals might have the function of justifying and excusing their poor password practice.

The **validation stage** uses four main criteria to validate the findings of the research. Firstly, the analytic claims being made should give *coherence* to a body of discourse. Secondly, the *participants' orientation* to the findings about the variability and consistency of their statements should confirm these findings. Thirdly, the analysis should have identified *new problems* that were not envisaged at the outset and should have provided solutions to them. Finally, the fourth and in many ways most powerful criterion of validity is *fruitfulness*, which refers to the ability of the analytic scheme to make sense of new kinds of discourse and to generate novel explanations.

The **report stage** does more than just present the research findings and constitutes part of the confirmation and validation procedures itself. The presentation should make it possible to assess the researcher's interpretations and should therefore include a representative set of examples from the area of interest along with a detailed interpretation which links analytic claims to specific parts or aspects of the extracts.

The **application stage** aims to put the research findings into practice but is often ignored in research of this type. One possible model for the application of discourse analysis is that of popularisation by giving away the knowledge as freely as possible. Another model is to open up a dialogue with the people who have been researched. The discourse-analytic study presented in this thesis has fed into the overall recommendations on how to implement persuasive password security (Chapter 8), which should ensure both the application and further validation of the findings.

3.5 Protection motivation theory

3.5.1 Justification

The *fear appeal* is the persuasive message form most commonly used in health campaigns, but is also employed by, amongst others, advertisers, politicians and even parents (Witte et al. (2001)). It is designed to arouse fear in the recipient by outlining the negative consequences that occur if a certain action is not taken. This is hoped to

change the behaviour of the recipient in ways that are deemed advantageous by the creator of the fear appeal. Rogers' protection motivation theory (Rogers, 1983) states that fear appeals will be effective if they convince the recipient that

1. the problem is serious (*perceived severity*);
2. it may affect him (*perceived susceptibility*);
3. it can be avoided by taking appropriate action (*response efficacy*);
4. the recipient is capable of performing the necessary behaviours required to avoid the problem (*self-efficacy*).

Chapter 4 describes how the research approach used in this thesis has been re-conceptualised during the research process. However, when studies 1 and 2 were carried out, protection motivation theory was still regarded as a possible cornerstone of the research that was to be conducted. There were 2 principal reasons for this::

1. The main application area in which fear appeals have been tested successfully is that of health campaigns (Witte et al. (2001)). Their use in this area to entice individuals to perform behaviours that are good for them mirrored the author's understanding at the time of the main task that would have to be performed as part of the research.
2. Protection motivation theory offers a set of variables (e.g. perceived severity) that appeared to provide a suitable initial focus for an investigation of mental constructs such as attitudes, beliefs and knowledge items that might influence password practice.

However, study 2 revealed a number of problems with the use of protection motivation theory for this research:

1. Restrictions imposed by security department in hosting organisation
A central element of every fear appeal is the threat part of the message (see section 3.5.2), which aims to convince the recipient of both the severity of the danger and his susceptibility to it. However, the design of this part of the fear appeal used in study 2 was restricted severely by constraints placed upon it by the security department of the organisation in which the field trial was conducted. As an example, it was not possible to mention previous security breaches since this would have tarnished the reputation of the organisation. These unforeseen restrictions led to the fear appeal being considerably less threatening, and, as a consequence, less convincing than was desirable. This point became even more poignant when the author discovered another model of how fear appeals work, the Extended Parallel Processing Model (EPPM) (Witte et al. (2001)). This model is an integration and extension of a number of behaviour change models, one of which is protection motivation theory. It stresses that fear appeals act as external stimuli to prompt some sort of message processing and that their appraisal by recipients goes through an initial stage in which the relevance and seriousness of the threat is evaluated, which determines whether there will be a response to the message. It became clear at this point that fear appeals were not a suitable means of changing user behaviour in this research, since most organisations were likely to put similar limitations on their use.

2. Factors other than fear present

The goal of fear appeals is to change the behaviour of recipients by arousing a fear of the consequences of not taking certain actions. It assumes that users will consider the variables proposed by the theory to determine their behaviour. However, it became clear that while these variables had an effect on user behaviour, there also were other factors that affected it strongly and which had to be taken into consideration.

3. Difference between health campaigning and security awareness campaigning

There is a fundamental difference between a health campaigner exposing a recipient to a health message in the form of a fear appeal and an organisation presenting its employees with a security message in the form of a fear appeal. The former will usually have the aim of convincing the recipient to do something which is good for him, whereas the latter will try to convince him to do something that is good for the organisation. As a result, there will be a number of factors present in the second case which are not necessarily there in the first (e.g. user resistance), and these factors can have a strong bearing on the effect of the fear appeal.

These issues led to a reappraisal of fear appeals and resulted in them being discarded as a means of changing user behaviour in this research. However, they still were used in studies 3 and 5, but now with the aim of stimulating a discussion in the focus groups which would make it possible to identify further factors that affect user behaviour, in particular with respect to the effects of punishment regimes that are introduced by an organisation. The use of fear appeals in this thesis, and the identification of their shortcomings, are methodological contributions of the thesis.

3.5.2 Application

Fear appeals typically have two components (Witte et al. (2001)):

1. The threat part of the message usually describes the negative consequences that will occur if the recipient does not do what is advocated.
2. The recommended response, which will make it possible to avoid the threat, is then described.

The design of the fear appeals used in this research was performed under constraints imposed by the organisations in which the fear appeals were to be used. This resulted in appeals having to be designed in an iterative fashion. First of all, a 'strong' fear appeal was created, which was based on the knowledge about the variables used by protection motivation theory which had been gained from the studies that had been conducted prior to the design of the appeal. This was then given to the relevant people in the organisation for approval and usually returned with the request for a considerable number of changes which in effect weakened the fear appeal. A new appeal incorporating these changes was designed and put forward for approval again, until a final version had been arrived at.

3.6 Chapter summary: The research approach

This chapter has introduced grounded theory as the principal methodology to be used in the thesis research (section 3.3). Discourse analysis was presented as a supplementary methodology that will make it possible to investigate certain issues which grounded theory does not lend itself to easily (section 3.4). The use of this methodology, which makes possible a first evaluation of its suitability for problems within HCI, is a methodological contribution of the thesis. Protection motivation theory was discussed as the methodology that was originally intended to guide the design of persuasive messages, and a number of its shortcomings that became obvious during its use in the thesis research were detailed (section **Error! Reference source not found.**). The identification of these shortcomings is a methodological contribution of the thesis. Finally, questionnaires, semi-structured interviews and focus groups were presented as the means by which data was to be collected (section 3.2).

4 RESEARCH CHRONOLOGY

4.1 Overview

Chapter 2 has developed the research problem and the research questions that have driven the work documented in this thesis. Chapter 3 has presented and justified the methodological approach that was taken to address this problem and these questions. However, the research problem and the research questions, as well as the research approach had not been defined in their current form at the beginning of the research process that led to the findings that will be presented in chapters 5 to 8. Instead, the research approach was re-conceptualised and refined during the research process. As a consequence, data that had been collected up to a certain point was now re-used and re-analysed in a manner that had not been envisaged at the outset. This has led to a situation where the connection between individual studies that have been carried out and specific findings that are presented in chapters 5 to 8 is not as clear-cut as it might have been had the final research problem and research approach been used from the beginning. For example, chapter 6 provides a grounded theory model of the way in which users choose password-related behaviours in the absence of any organisational efforts to affect this choice. The initial version of this model was developed on the basis of the qualitative data collected in study 1, but it was continually refined as more data was collected in studies 3, 4, 6 and 7. The extensions to the grounded theory model, which incorporate the effect of regulations and their associated punishment regimes on user behaviour and are presented in chapter 7, were primarily developed using the qualitative data obtained in studies 3, 4, 6 and 7. However, they were also refined by looking at the limited bits of relevant material within the data gathered as part of study 1. As a consequence of all this, it is felt that a thesis structure in which individual results chapters were preceded by the studies that have informed them could not only be misleading, but would also make the thesis cumbersome reading material. Instead, the description of all the studies will be put into this single chapter, and can then be referred back to in other chapters of the thesis. This chapter then serves two functions within the wider context of the entire thesis:

1. It puts together into one place a chronological description of all the studies that have been carried out as part of the thesis research.
2. It outlines the development of the thoughts that have led to the formulation of the research problem, the research questions and the research approach in their final form.

4.2 Study 1: Interviews

The first study was motivated by a simple set of questions. In large organisations, a lot of users have similar jobs to do, and access information with the same degree of confidentiality. How can it be that some of them are motivated to behave in a security-conscious fashion, and some are not? Is this due to general personality differences, or can it be traced back to mental constructs they hold, e.g. their knowledge, beliefs and attitudes? And if it can be traced back to these mental constructs, would it be possible to entice users who behave improperly to take on the constructs of users that behave well, thus changing their behaviour? Finally, would it be possible to use fear appeals as a tool to instigate these changes in users?

In order to investigate these questions, semi-structured in-depth interviews on password security were carried out with 17 participants. Ten of these worked for British Telecom, 6 were doctoral candidates at University College London, and one was a systems administrator working in a bank. The interviews lasted 30-60 minutes and were subsequently transcribed for analysis. They were kept as open as possible and allowed participants to introduce relevant new topics to the discussion. However, they also all contained a set of questions that were chosen to investigate the fear appeal variables (see section 3.5.2).

The interviews were initially analysed with the aim of identifying beliefs, attitudes and knowledge items. However, at this point the author became aware of discourse analysis and its potential to answer the questions that motivated this study in a manner that could lead more directly to the creation of interventions aimed at improving users' password behaviour (see section 3.4.2). This led to the decision of using discourse analysis in order to identify the interpretative repertoires that participants drew on to describe password-related issues. The results of this analysis and the recommendations based on it have been presented in Weirich & Sasse (2001b), Weirich & Sasse (2001a) and Sasse et al. (2001).

4.3 Study 2: Field trial 1

Study 2 was designed to investigate a way of improving users' password behaviour by changing password policies and the way they are enforced in a manner that highlights the danger of insecure password practices for the organisation rather than the user. This idea was a first result of the application of discourse analysis to the interview data, which at the time of study 2 was still on-going. The original aim of study 2 was

to test this suggestion in a field trial that was to be conducted with students in the computer science department of University College London. A one-page flyer was constructed which reminded students to protect themselves by protecting their computer science account (see appendices C and E). The flyer was designed as a fear appeal within the constraints imposed by the security department within University College London computer science. In particular, it was not possible to mention specific breaches to security that had occurred in the past or cases of students actually having been punished in particular ways because of their insecure password practices. In addition, it was necessary to keep the form potential punishment might take as vague as possible. The resulting flyer contained two short paragraphs that explained the importance of secure password practices to University College London's computer science department, arguing that it needs to be seen as highly security-conscious in order to guarantee continued support from industrial collaborators and research councils. This was followed by a brief list of the measures students should undertake in order to protect their account, which also contained the link to a website which students were encouraged to visit in order to find out more about how to create strong, yet memorable passwords (see appendix D). The first version of the flyer left it at that (see appendix C), whereas a second version included a short statement in bold which stressed that in the case of illegal activities being undertaken from a student's account and the student claiming they were performed by someone else, his past behaviour with respect to the password instructions would be investigated and any breaches would count strongly against him (see appendix E).

New students in the computer science department had to visit a designated room during enrolment week in order to obtain the username and password for their computer science account. They were separated into two groups of roughly equal size on the basis of the course they were studying. Students in group 1 were given version 1 of the flyer by the author immediately after they had been given their username and password by a member of the security department and were reminded to take its content to heart. Students in group 2 were given version 2 of the flyer. For practical reasons, it was not possible to create a control group from the enrolling students by giving some of them no flyer, since the room in which the usernames and passwords were given out was filled constantly by waiting students from all courses who would have been likely to notice a difference in procedure. It was decided to use the 2nd and 3rd year students in the department as the control group (group 3).

The field trial was followed up with a questionnaire (see appendix F) 5 months later. The questionnaire was designed to determine the effect of the flyers on the attitudes and behaviour of the students and was handed out at the end of lectures in which each of the groups was represented exclusively. Students were given 2 pounds for completing the questionnaire. 57 questionnaires were collected from students in group 1, 39 from students in group 2 and 36 from students in group 3. The responses were then analysed using standard statistical techniques (see appendix B). The results were disappointing. The differences in the responses of the individual groups to the questions were statistically not significant. In addition, only two of the students who had received one of the flyers had visited the website that was advertised in them. The only possible interpretation of these results was that the flyers had had no effect on the students. The field trial had not yielded the expected results.

The field trial was followed up with a reappraisal of the work that had been carried out so far. Informal conversations with students who had been part of the field trial, as well as a first tentative analysis of some of the interview data using grounded theory pointed to a possible reason for the inability of the fear appeal to change user behaviour: it seemed that a number of factors which strongly influenced the participants' password behaviour in general and their reaction to threats of punishment in particular had been ignored. The research up to this point had focussed on a number of factors, in this case the interpretative repertoires users hold and the variables proposed by protection motivation theory, and assumed them to be largely responsible for users' poor password practice, ignoring other factors that play a major role.

The recognition of this fact led to a major rethink of the approach to be taken in this research. As a result, the research problem and research questions were re-conceptualised and took on the form discussed previously in this thesis (see section 2.5). As importantly, it had become clear that discourse analysis would not make it possible to answer these questions adequately and grounded theory was chosen as the new, primary methodology to be used to analyse all the data that had already been gathered and all the data that would be collected during the remainder of the research. In addition, fear appeals were discarded as a means of changing user behaviour in this research (see section **Error! Reference source not found.**). This still left open the question of what was to be done with the data and results obtained so far. It was decided that the interview data that had been collected would form the basis of a first

grounded theory analysis. The quantitative results obtained as part of study 2 were to be considered as coming from one group and would form the basis of a description of user behaviour and attitudes (see chapter 5). The discourse-analytic study would be continued, but would now be considered ancillary to the grounded theory analysis, which was to be the main focus of the research.

Field trial one was followed up with a number of focus groups, which are described as part of study 6 (section 4.7).

4.4 Study 3: Fear appeal scenario

At the time of study 3, the re-conceptualisation of the research problem had only just begun, and discourse analysis was still the methodology to be used for the analysis of the data. The study aimed to investigate users' reactions to threats of punishment for insecure password practices.

A flyer (see appendix G) containing a fear appeal was designed, as had been done in study 2. However, this time the flyer was only going to be used to stimulate discussion and it was possible to make it slightly more threatening than the one used at University College London. The study took the form of focus groups, which were conducted in conjunction with another researcher into the human issues involved in password security, since this was the only way in which British Telecom would allow access to end-users at Adastral Park, where the groups were held. 3 focus groups with 4, 5 and 5 participants respectively were carried out. The participants had been recruited by the author's industrial supervisor by means of e-mail circulars, and consisted of receptionists, security researchers, system administrators, human factors specialists, engineers, contractors and postgraduate students on a placement.

The focus groups lasted for one hour, which was split equally between the two moderators. The first half of each group, which belonged to the author, began with the participants being given the flyer to read. This was then followed by a discussion of the flyer and its potential effect on participants' attitudes and behaviour. The second half of the group, which belonged to the other moderator, began with the participants being given a password manager to read. This was a paper-based tool designed to help users overcome usability problems typically experienced with respect to password usage. The password manager and its potential usefulness for the participants was then discussed.

The two researchers agreed to let each other use the data obtained through the focus groups in its entirety. The transcripts were initially analysed by the author using discourse analysis, but once grounded theory had been chosen as the primary methodology for the research, it was also applied to the data.

The format of focus groups employed in this study had shown itself to be useful for a number of reasons. Firstly, it made it possible to gain access to a larger number of participants than interviews would do in the same amount of time. Secondly, using a semi-structured focus group agenda allowed for the elicitation of targeted information which could confirm or disprove hypotheses that had been made on the basis of the analysis of previously collected data. Finally, the more open-ended parts of the focus group facilitated the expression of new ideas by the participants which then made it possible to develop further hypotheses. All this meant that focus groups were an ideal data collection tool for the refinement and extension of the grounded theory model, and for the identification of additional interpretative repertoires. It was decided to use focus groups as the central data collection technique for the remainder of the research process.

4.5 Study 4: Investigating a different user group

Study 4 took place shortly after study 3 and still was guided by the original approach to the research. A conscious effort was made in studies 1 to 3 to ensure that the participants would be recruited from user groups that were as varied as possible. However, the practical constraints the work was carried out under had clearly led to there being a bias towards users with an academic background. Study 4 was a piece of opportunistic research that aimed to investigate whether the findings that had been made so far in the research process would be applicable to a user group with a different background. The opportunity for this study came about because the researcher the author had co-operated with in study 3 had obtained access to end-users in one of British Telecom's operational arms through a contact of his within the company. As a result, two focus groups could be carried out, each containing 5 participants who had been encouraged to participate by the promise of drinks being bought for them after work. The participants consisted of administrative personnel, technical support engineers and their line managers, and all came into regular contact with sensitive customer data as part of their work. The focus groups were restricted to 40 minutes, which were again split equally between the two moderators. The first part

of each group was led by the co-operating researcher, who instigated a discussion about the number of passwords participants had and the ways in which they managed them. The second half was led by the author, and consisted of a condensed set of the questions that had been asked in studies 1 to 3 and which aimed to investigate the viability of as many aspects of the findings that had been generated so far for this user group as was possible.

The two researchers again agreed to let each other use the data obtained through the focus groups in its entirety. The transcripts were initially analysed by the author using discourse analysis, but were later also investigated using grounded theory.

4.6 Study 5: Field trial 2

At the time of study 5, the re-conceptualisation of the research approach had been finished and a grounded theory analysis of the data that had been collected so far had been begun. The field trial carried out in study 5 was an extension of the one carried out in study 2 and took place at the beginning of the following academic year. All new students to the computer science department of University College London were asked to sign a copy of version 2 of the flyer (see appendix E) when they were handed their username and password. This was not expected to incur serious changes in their attitudes and behaviour compared to the first field trial, where students were not asked to sign the flyer. Rather, it was meant to test an argument brought up in the focus groups carried out with participants in study 2 (described here as part of study 6 in section 4.7), where some participants had stated that the flyer would have had a stronger effect on them had they been forced to sign it. Testing this argument should either confirm its validity or make it possible to identify additional factors that influence password behaviour. In addition, a failure to confirm the argument would support a discourse-analytic study of the possibility of participants structuring their discourse in order to justify and excuse their poor password behaviour. The questionnaire used in study 2 (see appendix F) was employed in the same manner here to determine the effect of the intervention, and 33 responses were collected 3 months after the flyers had been handed out. These were compared to the responses obtained in study 2 using the same statistical techniques, and the result was again that the intervention had had no effect on the students.

4.7 Study 6: Field trial focus groups

The evaluation of both the first and the second field trial was followed up with focus groups that primarily aimed to refine the grounded theory model that had been developed so far. However, the fact that all of the participants had been exposed to the fear appeal also made it possible to investigate whether there were any additional variables that affect users' reaction to punishment regimes. Moreover, this exposure to the fear appeal, which some of the participants also had signed, allowed for a further study of the discursive practices that users might engage in to justify their poor password practice.

6 focus groups were carried out, 3 with participants in field trial 1 (with 3, 4 and 5 members), and 3 with participants in field trial 2 (all with 4 members). The focus groups encouraged interaction among the participants and the voicing of possibly controversial opinions, but were still guided by a semi-structured agenda. The participants were first asked to fill in the questionnaire used in studies 2 and 5 in order to ensure that no outliers were among the sample group. The first half hour of the session, which took one hour in total, was then dedicated to general issues about password security and was included to validate and refine the grounded theory model. The participants were then shown version 2 of the flyer and the remainder of the session was dedicated to discussing its effect on them and ways in which it could be improved. The transcripts of these focus groups were analysed using grounded theory and discourse analysis.

4.8 Study 7: Single-sign-on scenarios

The seventh and final study was carried out at about the same time as the focus groups with participants in study 5 and had the same goals: to refine the grounded theory model, to identify whether there are additional factors that affect users' reaction to punishment regimes and to investigate the possibility of any discursive practices users might employ to justify their poor password practice. The study consisted of 6 focus groups (with 6, 5, 5, 5, 4 and 4 members) that were carried out at British Telecom Adastral Park. As in study 3, the participants had been recruited by the author's industrial supervisor by means of e-mail circulars, and again consisted of receptionists, security researchers, system administrators, human factors specialists, engineers, contractors and postgraduate students on a placement. None of the people participating in study 3 were present in these focus groups.

The focus groups lasted one hour and were divided into 3 stages, each of which lasted about 20 minutes. The first stage consisted of a set of questions that further validated and refined a number of elements of the grounded theory model. During the second stage, the participants were shown on a blackboard a list of the 5 password-related behaviours that make up the guidelines discussed in section 2.3.5.1, and these behaviours were discussed in the group in order to determine whether there were any additional factors that needed to be considered in the corresponding part of the model. The third and final stage consisted of the discussion of three different variants of single-sign-on which were presented and explained to the participants. These scenarios were primarily designed to disarm arguments that had been brought forward in previous interviews and focus groups to justify poor password practice. Their use made it possible to refine the model further, but also allowed a clearer insight into whether participants might structure their discourse in order to justify insecure password behaviours.

The transcripts from these focus groups were analysed using grounded theory and discourse analysis.

4.9 Chapter summary

This chapter has provided a chronological description of the studies that were carried out as part of this thesis. In conjunction with this description, the development of the thoughts that have led to the formulation of the research problem, the research questions and the research approach in their final form has been outlined. There were four pivotal moments in this development:

1. The research was begun with a focus on the attitudes, beliefs and knowledge items users hold with respect to password security and fear appeals as a potential way of changing them (study 1).
2. Discourse analysis was chosen as the primary data analysis methodology (study 1), which meant a move away from mental constructs such as attitudes and towards interpretative repertoires.
3. A field trial did not yield the expected results and led to a re-conceptualisation of the research problem and the research questions (study 2). As a result, grounded theory was chosen as the primary methodology to be used and discourse analysis was relegated to being an ancillary methodology. Fear appeals were discarded as a means of changing user behaviour.
4. Focus groups were found to have a number of advantages over interviews and were chosen as the data collection procedure for the remainder of the work (study 3).

5 PASSWORD PRACTICES AND ATTITUDES TOWARDS PASSWORD SECURITY

5.1 Overview

In section 2.4.3, it was pointed out that one of the central tenets of research on human factors in computer security - that users regularly behave in a manner that undermines security - is primarily supported with evidence that is of an anecdotal nature. There is little published research that establishes users' insecure password practices in a quantitative manner. There is even less research that tried to identify the specific insecure behaviours users choose over secure ones, and the situational contexts in which these choices take place. In section 2.4.8.2, it was pointed out that these issues are not the main focus of this thesis, but that the data collected makes it possible to address them anyway. This chapter primarily provides the results of an opportunistic investigation of these issues, by analysing the data collected in studies 1 to 7. In addition, it summarises data on users' attitudes collected as part of studies 2 and 5.

While the aim of collecting quantitative data on user behaviour is to identify the extent of malpractice that occurs, it is the collection of qualitative data about specific user practices that will often allow organizations to design specific measures aimed at improving password practices and overall security. The most obvious ways of using such information are:

1. To improve the usability of password mechanisms, e.g. by making it easier and faster to set up new accounts for visitors who need short-term access, which means that there is no need for their host to give them access to theirs.
2. To change the password mechanism so that certain practices are not feasible, e.g. by introducing proactive password checking, which makes it impossible to choose weak passwords.
3. To monitor and punish specific insecure practices, e.g. by checking employees' desks on a regular basis.
4. To point out the danger of specific practices in security awareness and education programs, e.g. by running a campaign that explains to users how practices such as the storing of passwords in files on their computers system can get exploited by attackers.
5. To reduce the risk associated with specific practices which cannot be stopped, by offering more user-friendly alternatives, e.g. by providing a common place to store copies of passwords for groups of workers.

In the first part of this chapter (section 5.2), the password practices of the participants in the seven studies carried out for this thesis will be detailed. Part of the

questionnaire - handed out to students at University College London as part of studies 2 and 5 - will supply quantitative data about the extent to which the password rules were violated. In addition, the password practices reported in the interviews and focus groups of studies 1, 3, 4, 6 and 7 will be collated. The results that will be presented are obviously specific to the user groups that have been investigated, which means that any attempts to generalize them have to take into consideration the constraints discussed in section 1.4. They are also based on participants' self-reports and have to be taken with a certain degree of caution until they have been confirmed by observational studies, where this is possible (see section 3.4.1). However, they provide a first step towards a larger collection of such data which can guide future research and organizational interventions.

The second section of the chapter (section 5.3) will present the results of an analysis of the parts of the questionnaire from studies 2 and 5 which collected users' attitudes to password security. This information can be used to improve security awareness and education campaigns, which can target specific attitudes that need to be altered in order to improve password practices.

5.2 Password Practices

5.2.1 Password choice and password changing

5.2.1.1 Choosing an individual password

A number of authors (e.g. Anderson (2001), Schneier (2000), Winkler (1997)) have claimed that users will choose weak passwords unless they are prevented from doing so by mechanisms such as proactive password checking. The claim is often substantiated by demonstrations of how easily and quickly a large number of passwords from particular real-life password files can be cracked with standard cracking tools (e.g. Anderson (2001)). However, there has been very little research on the actual methods that users employ to generate their passwords (the main exception is Adams & Sasse (1999)).

In the questionnaire handed out as part of studies 2 and 5 (see appendix F), 73.1% of the students claimed that they had made a conscious effort to choose a departmental password that was difficult to crack. They were then asked about the technique that

they had used to select this password. They had a choice of three types of password stems:

1. A word.
2. A concatenation of words.
3. A password constructed from the first letter of each word in a song or poem.

They were also asked to tick any of the following two alterations to this stem:

1. Adding a number.
2. Replacing characters or numbers with a symbol.

In addition, they could describe any method that was different from the ones presented. The only alternative approach that occurred repeatedly was the choice of what was described as a 'random password' for the stem. Out of the 165 returned questionnaires, 5 had no answers to any of the questions concerning participants' password practices. A further 47 were ambiguous in their answers to the question about participants' method of password choice. This means that 52 had to be discarded, leaving a total of 113. The count and percentage of the various combinations of stem choice and alterations are summarized in Table 2.

The results confirm the claims that users tend to choose weak passwords when not prevented from doing so by measures such as proactive password checking. A considerable number of the students chose their password in a manner that is likely to lead to weak password content. In particular,

1. 9.7% of the students had chosen an ordinary word or a concatenation of words with no alterations, and
2. 58.4% used an ordinary word or a concatenation of words as a stem, which was then altered by adding numbers and/or replacing numbers and letters with symbols.

These are the methods of password construction that are targeted first by cracking tools, and consequently passwords generated this way are the first to be cracked (see sections 2.3.4.1. and 2.3.4.2).

The interviews and focus groups conducted with participants at University College London and at British Telecom further confirmed these findings. This will be discussed in the following section, which deals with the ways in which users handle multiple passwords that have to be changed regularly.

		ALTERATIONS					
STEM		No alterations	Numbers only	Replacement only	Numbers and replacement	Other	Total
	Word	7 (6.2%)	26 (23.0%)	3 (2.7%)	6 (5.3%)	2 (1.8%)	44 (39.0%)
	Concatenation of words	4 (3.5%)	14 (12.4%)	1 (0.9%)	5 (4.4%)	0 (0.0%)	24 (21.2%)
	Letters from song/poem	1 (0.9%)	8 (7.1%)	6 (5.3%)	2 (1.8%)	0 (0.0%)	17 (15.1%)
	Random	0 (0.0%)	4 (3.5%)	0 (0.0%)	6 (5.3%)	0 (0.0%)	10 (8.8%)
	Other	2 (1.8%)	3 (2.7%)	0 (0.0%)	0 (0.0%)	13 (11.5%)	18 (15.9%)
	Total	14 (12.4%)	55 (48.7%)	10 (8.8%)	19 (16.8%)	15 (13.3%)	113 (100%)

Table 2 : How respondents chose their password (n=113)

5.2.1.2 Choosing multiple passwords that have to be changed regularly

The previous section has shown how the students at University College London chose their departmental password. 48.8% of them stated that they were also using this password for other systems. This is a first indicator of the problems and security violations that can be caused when users have to use passwords for different applications. In British Telecom, the difficulties created by such multiple passwords were obvious. Three factors in particular stood out:

1. Users had to access not just one system within the organization, but a potentially large number of them, for each of which they were supposed to generate a unique password. In addition, a lot of participants accessed a number of non-organizational systems, which also needed passwords.
2. Many organizational systems employed proactive password checking, but the criteria for what constitutes an acceptable password differed between them.

3. Most systems forced their users to change the password regularly, most commonly every 30 or every 90 days. Some of them kept a history list of the most recent passwords users had chosen, making it impossible for them to pick them again.

Adams & Sasse (1999) have reported that 50% of respondents to their web-based questionnaires created memorable multiple passwords by linking them via some common element. The interviews and focus groups conducted in British Telecom made it possible to identify more clearly the various strategies users employ to deal with multiple passwords that have to be changed regularly.

The problems faced by users in British Telecom made it impossible for them to comply with all of the password guidelines. To do so would have meant to choose a large number of completely different strong passwords, memorize them, and change them for new strong passwords after relatively short periods of time. The creation of such a large number of passwords at regular intervals is in itself a major overhead, in particular since most participants had not been trained in how to choose strong and memorable passwords. In addition, there is obviously a limit to the number of passwords that a user can remember, especially when he has to change most of them again after 30 days (see section 2.4.4.2.1). Almost all participants responded by writing down at least some of their passwords (see section 5.2.2). In addition, they would try to reduce the number of passwords they were dealing with. For some, this meant that they attempted to reduce them to a single one, which they then would change slightly in order to accommodate the different criteria used by the various proactive password checking mechanisms. Others would employ a system of two or more tiers, which would be differentiated by how secure they wanted to keep the system(s) in each tier. The low-security tier would then typically have one password which would be used for a number of systems, whereas a high-security tier might use a different password for each system.

In order to deal with the change regimes employed by various systems, participants tended to connect subsequent passwords in some manner. Most commonly, this was done by altering elements of the password. Since some systems kept a history of the last few passwords used, this could take on rather ingenious forms:

“I've got a system that's pretty foolproof. My password say is twittwit5; the next month is uwit-after t is u - uwituwit5... then x y z, then awit bwit cwit, and so on. Most machines can't remember more than 26.”

The other way of linking passwords was through the use of some common theme, such as characters in a cartoon or even fruits:

“I mean, an example of a theme set would be fruit. [...] So the password then would be orange, apples, pears or something like that.”

Some participants used cues from their physical context to choose and remember such theme-based multiple passwords. These included books on the bookshelf next to the computer, or mouse-mats:

“Well, the NT one, it’s my, what I actually got on my, I’ve got a mouse mat, and it happens to be a phrase connected to the current mouse mat, and I’ve got about 4 or 5 different mouse mats.”

It is obvious that the lack of a consistent password policy at British Telecom made it impossible for the participants to follow all of the password guidelines. Instead, they had to develop their own ways of dealing with the large number of passwords, the inconsistent criteria used by the various proactive password checking mechanisms and the varying timeframes within which passwords would expire. The systems could be made more usable through integration. Single-sign on, despite some drawbacks, can significantly reduce users’ mental workload and anxiety. Ultimately, the number of passwords should be reduced, but in the meantime, increasing the consistency of the password policy on different systems would at least alleviate the problem. In essence, this would make it easier for users to reduce the number of passwords by using the same password for at least some of the different systems without being hindered by the proactive password checking mechanism and the password ageing mechanism.

5.2.2 Memorizing the password and writing it down

Adams & Sasse (1999) report that 50% of the respondents to their web-based questionnaire had written down their passwords in one form or another. The students at University College London stated in the questionnaire that 20.6% of them had initially written down their password, but destroyed the note once they had memorized it. 8.1% of them kept a permanent written copy of the password. The interviews and focus groups conducted with employees of British Telecom found not only that most participants wrote down at least some of their passwords, but also made it possible to identify the manner in which these written copies are typically kept. This information can be used to determine more clearly the security risks users take when writing down their passwords.

The participants stored copies of their passwords both electronically and in a paper-based format. Electronic copies were kept in a file on

1. their PDA, which might itself be password-protected;
2. on a computer at work or at home.

The file was usually given an inconspicuous name (i.e. it was not called 'password file'), and in some cases was encrypted.

Paper-based copies were kept

1. in a locked drawer;
2. in the participant's wallet;
3. in a diary or personal organizer;
4. at the back of a book which was locked away;
5. buried among other material:

"Yeah, I do, yes, I do have a, an innocuous way of, of, buried in loads of other stuff, so nobody would ever, ever be bothered to go through and find it."

One participant reported mislaying the copy on a regular basis:

"I tend to forget them. I tend to write them down and then forget where I have put them."

In addition, some parts of the organization had set up a routine whereby all members of a workgroup would hand an envelope with a written copy of certain passwords to their superior, who would lock these away. This was meant to be used in cases where an employee was absent and access to his computer was urgently needed (this dealt with users' individual accounts, not the ones they already shared as part of the workgroup).

Finally, not all participants kept copies of the passwords themselves. Instead, some wrote down hints that allowed them to reconstruct the passwords when needed:

"I use a little symbol or something to remind what the basic password is and then I use a little squiggle to remind me what the ending part is or something like that."

The information in this section can be used in three ways:

1. It makes it possible to determine more clearly the security risks users take when writing down their passwords.
2. It can be used to improve monitoring efforts by the organization.
3. It can be employed to design measures that reduce the risk. The previous section has shown that for a lot of users it is impossible to handle the multitude of passwords they require without writing them down. However, they could be shown ways of doing this that are more secure than the ones they are currently using.

5.2.3 Sharing the password

A number of authors (e.g. Anderson (2001), Schneier (2000)) have reported that users easily share their passwords with other people, which is a fact that can be exploited by social engineering attacks (see section 2.4.3). 21.3% of the students answering the questionnaire in studies 2 and 5 stated that they had shared their departmental password. 16.3% had been on the receiving end of such sharing practices. The interviews and focus groups showed that sharing was a common practice both at University College London and at British Telecom. They also made it possible to describe the manner in which such sharing typically takes place, and in particular the situational contexts in which it happens. The following list summarises the eight situations for which participants reported having disclosed their password to third parties in the past. It does not include information on how many participants mentioned each situation. Since it is based on data drawn from interviews and focus groups which did not necessarily cover exactly the same subject areas, such information could be misleading. However, the first situation was clearly the most common, with the following four occurring with lower and about equal frequency. The last three situations were only reported once.

1. Remote access

By far the largest number of cases of sharing occurred with people who needed to access their email or other resources while they were away from their office and could not log in remotely:

“Ah, when I’ve had to go away on holiday, and haven’t had a PC with me for mail, I’ve shared my user, I’ve shared my mail password, I’ve usually changed it to something temporarily while that person has access to it, so that it’s private once I get back.”

2. Co-working

Participants shared passwords when they worked together on some project and an officially shared password had not been set up:

“There’s been, when we do experiments, it’s often to set, to set up the computer in order to do an experiment, we sometimes have to give each other passwords, I mean I’ve had another colleague’s password as well.”

3. Access to resources

Passwords were shared with colleagues who did not have access to certain resources themselves, either because they were not supposed to or because it was considered too much work to set these up:

“I’ve used other people’s passwords to log on to systems, where officially you shouldn’t be, you know what I mean, but I think it’s

generally accepted that you do and I think within the same context it's still considered reasonably safe to do."

4. Emergency access provision

Participants would issue colleagues with passwords in case they themselves were absent when access to their resources was urgently needed.

"Yeah, my work colleague knows my, my NT logon, ah, just in case I go, my machines goes and locks up and there's something on there that he needs to get, or he needs to close down the PC and he needs to get on himself so, yeah, he does know my password."

As mentioned in the previous section, parts of British Telecom had also set up an 'official' way of doing this through users' issuing superiors with envelopes containing relevant passwords.

5. No account set up

British Telecom regularly had visiting students who did not have an account of their own. Since setting up an account would have taken a certain amount of time, these students were either given their superiors' passwords or they at least got logged in by them without receiving the password:

"I have done that for, yes, I had a student come on site and she needed to, ah, run a, access a computer, and I just said 'Here is my account. Use that.' and she was sort of, in effect sitting on the next, next, ahm, computer, and she just used that, just to log in, connect to the network."

In addition, new starters in the organization might receive passwords by their superiors or colleagues until their account is set up.

6. Work for superior

Participants had received passwords from their superiors so they could carry out work on their behalf:

"And then there's [application name], where a chap who I used to work with, he was actually a program manager, and a very busy man, and he would ask me to do his [application name] work, so I had to have his, have his access."

7. Lack of computing competency

Passwords can get shared so that users with less computing knowledge can receive help:

"Well, ahm, because I'm computer-illiterate, ahm, I have to have a trusted friend who can help me out, so, ahm, one of the young people in the team who is very, very good at jiggling around with PCs has sort of taken me under her wing."

8. Contractors

Contractors can find themselves in situations where, because they officially are given less rights than ordinary employees, they receive passwords from other members of the organization:

"Sometimes you have to give your password, like, for example the work we, I work on we have a lot of user ID's to get onto BT's whole system, they might be preset in and as a contractor I've not got the authority to reset that particular password, it's only my boss, but if

we've got 20 passwords to reset in one go it's not worth, it's mismanagement really that he should have to reset the passwords, it should be me, so he gives me his password to be able to reset them"

The aim of listing the circumstances under which the participants had shared their password in the past is twofold. Firstly, it can direct the development of measures that prevent these actions from happening, either by monitoring them more closely, or by making them impossible through changes to the mechanism. Secondly, it can be used to educate users as to the specific risks they are taking, and to show them alternative ways of achieving their goals which compromise security at least to a lesser extent. In cases where an organization cannot realistically hope to prevent all forms of sharing, users can at least be made aware that is usually more secure to log the other person in, rather than giving them the password. The password can also be changed before and after sharing, which some of the participants already did (see point 1 above). In addition, an organization might of course also provide 'official' means of sharing passwords, which provide a higher level of security.

5.3 Attitudes

5.3.1 The risk of malpractice

The first set of questions of the questionnaire (see appendix F) handed out to students at University College London as part of studies 2 and 5 aimed to establish the risks that were associated with specific insecure password practices (see Table 3). The first question established the students' overall level of concern with protecting their account, and also served as a base mark against which to evaluate the answers to the following questions about specific password-related behaviours.

Testing for statistically significant differences ($p < 0.05$) made it possible to come to the following conclusions (see appendix B):

1. Respondents considered not protecting their account from unauthorized use as a risk to themselves. Only sharing the password with a stranger was considered to be as risky as not protecting the account from unauthorized use. All the other insecure practices mentioned in the following questions were considered less risky.
2. Three practices (keeping a written copy of the password until it is memorized, sharing it with someone who is known and trusted, using it for other systems) were considered to incur a lower risk than the remaining two.
3. Two practices (keeping a permanent copy of the password, sharing it with someone who is unknown) were seen to incur a higher risk than the other three. Of these two, sharing the password with a stranger was seen to be even riskier than keeping a permanent copy of it.

This information can be used in security awareness and education campaigns to target the practices that are considered lower-risk by users and make them aware of the dangers associated with them.

I put myself at risk if I...	Mean	Median
(1)...do not protect my computer science account from unauthorized use.	5.45	6
(2)...write down my computer science password in order to remember it, even if I destroy the note once I have memorized the password.	3.84	4
(3)...keep my computer science password written down permanently.	4.88	6
(4)...give my computer science password to someone I don't know.	5.52	7
(5)...give my computer science password to someone I know and trust	3.78	4
(6)...use my computer science password for other systems.	4.10	4

Table 3 : The risks associated with breaking password rules (on a scale of 1 (strongly disagree) to 7 (strongly agree)) (n=165)

5.3.2 The likelihood of an attack

The following three questions established the students' estimation of the likelihood of different kinds of attack (see Table 4). The likelihood of any of these attacks was regarded as quite low. The first one shows a statistically significant difference to the other two ($p < 0.05$), which between them show no such difference (see appendix B). A further question established that 86.1% of students believed that someone determined enough would ultimately find a way to get into their departmental account, even if they chose a strong password and kept it perfectly secret. Again, this information can be used in security awareness and education campaigns to improve users' knowledge about the types of attacks that are likely to be carried out, and about the effectiveness of the measures they can take to protect themselves.

Someone tries to gain access to your computer science account...	Mean	Median
...in order to damage your work.	2.59	2
...in order to perform actions that would incriminate you.	2.98	3
...to use it as a base for carrying out criminal activities.	3.18	3

Table 4: The perceived likelihood of certain types of attacks (on a scale of 1 (very unlikely) to 7 (very likely)) (n=165)

5.3.3 The threat of punishment

The two questions in Table 5 show that respondents

1. did not expect to get punished for breaking password regulations unless this had caused a security incident;
2. did not expect to get punished for a security incident as long as they had followed regulations in the past.

	True	False
The department cannot penalize me if someone does get into my account and carries out criminal activities, as long as I have followed the departmental regulations on proper password practice.	85.5%	14.5%
The department can only penalize me for not following departmental regulations on proper password practice if my behaviour results in someone getting into my account and doing harm to other people or the department itself.	64.2%	35.8%

Table 5 : Punishment expectancy for breaking password regulations (n=165)

The respondents were also given the scenario in Table 6, which describes a case where a student who has broken password regulations in the past experiences a severe security breach. They were then presented with the four arguments in Table 7, each of which implies a decreasing severity of punishment. For each of them, respondents were asked

1. how fair they considered it to be;
2. how likely they thought it would be that this would become the position finally taken by the department.

A scenario
<p>The computer science account of an undergraduate student has been used to hack into the system of a major company in the UK. Information in that system was tampered with, causing considerable financial harm to the company. The department has approached the student, who clearly states that these actions were not carried out by him/her. Further investigations by the department were inconclusive as to whether the student or someone else carried out the criminal activities. However, it became evident that the student had chosen a password that was easy to crack, and had also shared it with other people, thus violating departmental regulations. It could not be established whether the failure to keep his/her password secure directly led to someone else being able to break into the account.</p> <p>The members of the department sit down together and discuss whether they should punish the student. The four lines of argument presented below dominate the discussion.</p>

Table 6 : Scenario used in the questionnaire

Argument	Fairness		Likelihood	
	Mean	Median	Mean	Median
(1) The student should be punished – ➤ S/he cannot prove that s/he did not carry out the attack on the company’s system.	2.59	2	3.07	3
(2) The student should be punished – ➤ the department cannot prove that the student carried out the attack on the company’s system, <i>but</i> ➤ the student cannot prove that her/his failure to keep the password secure did not directly lead to someone being able to break into the account.	3.47	3	3.81	4
(3) The student should be punished – ➤ the department cannot prove that the student’s failure to keep the password secure did lead directly to someone being able to break into the account, <i>but</i> ➤ the student broke the regulations on choosing a strong password and keeping it secure.	4.34	4	4.82	5
(4) The student should not be punished – ➤ S/he might have broken the regulations on choosing a strong password and keeping it secure, <i>but</i> ➤ so do a lot of other students, and s/he cannot be punished for being unlucky enough for someone to get into the account, which was the only reason the bad password practice was found out.	4.35	4	3.82	4

Table 7: Fairness (on a scale of 1 (very unfair) to 7 (very fair)) and likelihood (on a scale of 1 (very unlikely) to 7 (very likely)) of the four arguments (n=165)

In conjunction with tests for statistical significance (at $p < 0.05$), the conclusions that can be drawn from Table 7 are (see appendix B):

1. The respondents considered the first argument to be both the least fair and the least likely to be taken up by the department;
2. The second argument was considered to be fairer than the first, but less fair than the 3rd and 4th, which were considered equally fair;
3. The third argument was considered to be the most likely to be taken up by the department, while the second was considered as likely as the fourth.

The information in this section cannot be used in the same way as the results of the previous two questions. However, it does highlight two important issues:

1. It is obvious that respondents needed more, and clearer, information on the punishment they could expect for violating password regulations.
2. Most organizations would have problems introducing punishment regimes that were regarded as blatantly unfair by their employees. This means they either have to change their employees' perception of what is fair, or they have to abstain completely from employing such regimes.

5.4 Chapter summary

A number of authors have reported that users customarily violate the basic password rules (e.g. Anderson (2001), Schneier (2000)). However, very little research has been carried out to substantiate these claims. The first part of this chapter has achieved two things:

1. It has shown that a considerable number of participants in the studies conducted as part of this thesis did indeed violate the password guidelines (e.g. 80 of the 165 students at University College London who filled in the questionnaire also used their departmental password for other systems). This has substantially increased the knowledge about the extent of this problem.
2. The exact manner in which these violations were carried out, and the contextual circumstances in which they occurred have been identified. This makes it possible to design measures that aim to reduce password violations, be it by changes to the mechanism, improved security awareness and education campaigns or more effective ways of monitoring user behaviour.

The second part of the chapter analyzed users' attitudes towards various aspects of password security. This information serves two purposes:

1. It supports the argument that users' insecure password practices may often be the result of their attitudes.
2. It details some of the problematic attitudes which need to be changed in order to improve user behaviour.

This chapter has addressed a number of issues that are not the main focus of this thesis (see section 2.4.8.2), but which could be addressed anyway through the opportunistic use of data that has been collected during the various studies that have been carried out. The following chapter will move on to the grounded theory model of how users choose password-related behaviours, which lies at the heart of the research presented in this thesis.

6 HOW USERS CHOOSE PASSWORD-RELATED BEHAVIOURS: THE CORE MODEL

6.1 Overview

Users of organizational computing resources regularly find themselves in situations in which they carry out behaviours that are directly related to password security. From the point of view of the organization, it would be ideal if users could only perform actions that are in accordance with security policies and consequently maximize overall security. However, in reality users usually have a choice of behaviour: they are able to pick a specific one from a pool which is available to them in a particular situation. This pool contains all those behaviours which a user is aware of and which he is able to perform (available behaviours). In this sense, it is a subset of the larger pool of all possible behaviours in a particular situation (see figure 1).

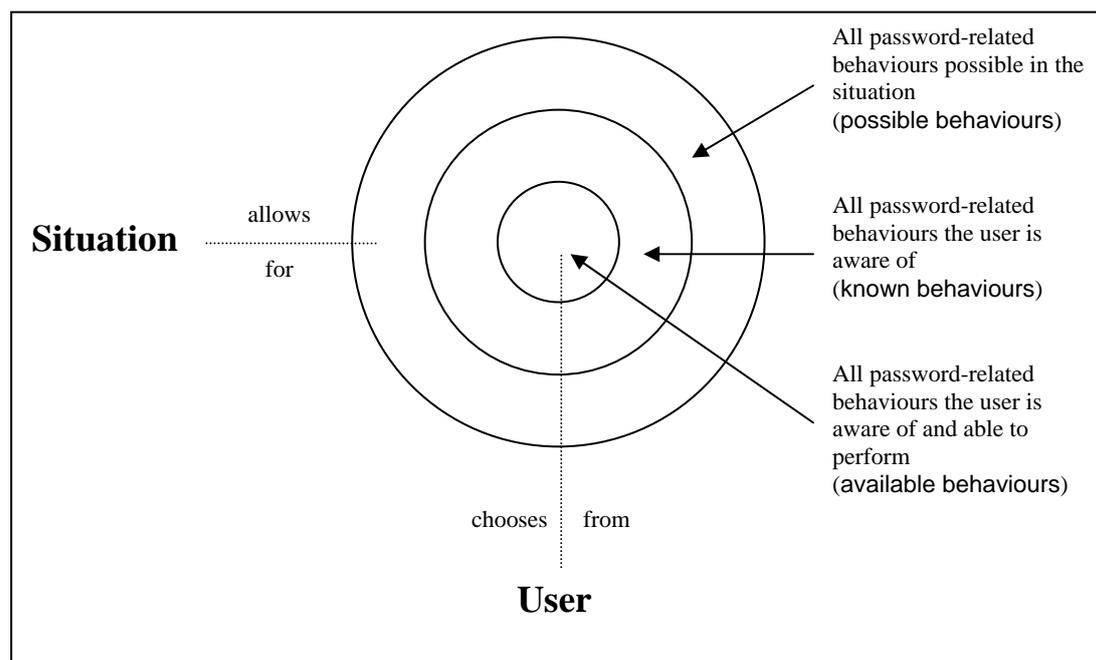


Figure 1: Pools of behaviours users can choose from in specific situations

One way in which an organization can try to improve users' password practices is by actively reducing the pool of available behaviours. This is usually done by changes to the password mechanism itself. Proactive password checking, for example, makes it impossible to choose certain weak passwords. A voluntary change regime allows users to choose non-action as an option, whereas an enforced change regime removes

this behaviour from the pool. Many effects of user education and training can also be explained by looking at their effect on the pools of possible and available behaviours. They often aim to move a behaviour, such as choosing a strong yet memorable password, from the former pool to the latter by making users aware of it (thus moving it to the pool of known behaviours) and by teaching them how to carry it out (moving it to the pool of available behaviours). However, in many situations, users still have a choice of a number of behaviours, some of which are secure, whereas others are not.

In this and the following chapter, a grounded theory model of the factors determining users' password practices will be presented. The central question any such model will have to answer is: why does a user choose one of several available behaviours over the others in a specific situation? Obviously, the ultimate aim of answering this question is to be able to influence this choice so that behaviours that are secure are chosen over those that are not.

The initial version of the grounded theory model presented in this chapter was developed using the interview data obtained in study 1. This data did not make it possible to determine the effect of organizational efforts to influence user behaviour sufficiently to include them at this point in time. This resulted in a core model that isolated the factors that determine users' choices if the organization makes no effort to influence them. Later studies collected data that made it possible to extend the model to account for the effect of such organizational efforts, and the results of analysing this data, which also developed and refined the core model, will be presented in the following chapter (chapter 7). The separation of the results of the grounded theory analysis of the qualitative data collected in this thesis into a core model and its extensions also has shown to have a number of conceptual and didactic advantages, which will be pointed out in the course of the presentation of the model.

The next section will present the core model of how users choose password-related behaviours in the absence of any organizational measures to influence this choice. Section 6.3 will discuss habitual choices and the recursive application of the model. In section 6.4, a specific user strategy that can be used to explain how password practices can deteriorate over time is examined.

6.2 Choice: the central category and the grounded theory model based on it

6.2.1 A grounded theory model of choice

The main aim of the model presented here is not to show how particular behaviours become part of the pool of available behaviours, even though this point will be touched on regularly. Instead, the goal is to isolate the factors that lead to users' choosing one of the behaviours in this pool over the others. The central category of the model therefore is **choice**, which from now on is understood to mean "choice of one of several available behaviours in a specific situation by a specific user" (see figure 1). A user will base this choice on his own estimates of the following four factors:

1. The level of security that the computing resource protected by the password in question requires (required security level).
2. The level of security provided by the various password-related behaviours that are available in the situation (provided security level).
3. The user cost incurred by performing any of these behaviours.
4. The user benefit obtained from carrying out any of these behaviours.

The required security level and the provided security level are continuous variables, but users generally treat them as discrete variables with a limited set of possible values. The number of these values varies between users, but a typical set would contain two: high-security and low-security. The use of these discrete values makes it possible to group computing resources and available behaviours and to match them easily according to the security level that is required and provided. It is important to note here that there are situations where even the behaviours that are in accordance with the guidelines set out in section 2.3.5.1, and which therefore provide the highest level of security that a user can achieve, might be seen by him as not to provide the required security level. However, this is a reflection of the users' perception of the security level that the password mechanism as such provides, and can be dealt with by him through the use of safeguards (e.g. backing up data regularly or taking sensitive information off company computers), which effectively reduce the required security level of the computing resource (see section 6.2.2.1.2). If the required security level cannot be reduced to a point where it can be met by the available password-related behaviours, the only option the user has is to choose behaviours whose provided security level comes as close to the required security level as is possible.

In an ideal world, users would choose only those behaviours that provide the level of security they believe is required. However, available behaviours can also incur a user cost, and this cost is often higher for more secure behaviours. Choosing a unique password, for example, requires a greater effort than re-using a password that is already in place for other computing resources. In addition, the various behaviours can also provide a user benefit. For example, sharing a password with a colleague who desperately needs access to a particular computing resource is likely to build trust and improve the relationship with that colleague. This means users choosing one of several available behaviours have to take into consideration not only the level of security each provides, but also the user cost that it would incur, and any user benefit that it might deliver (see figure 2).

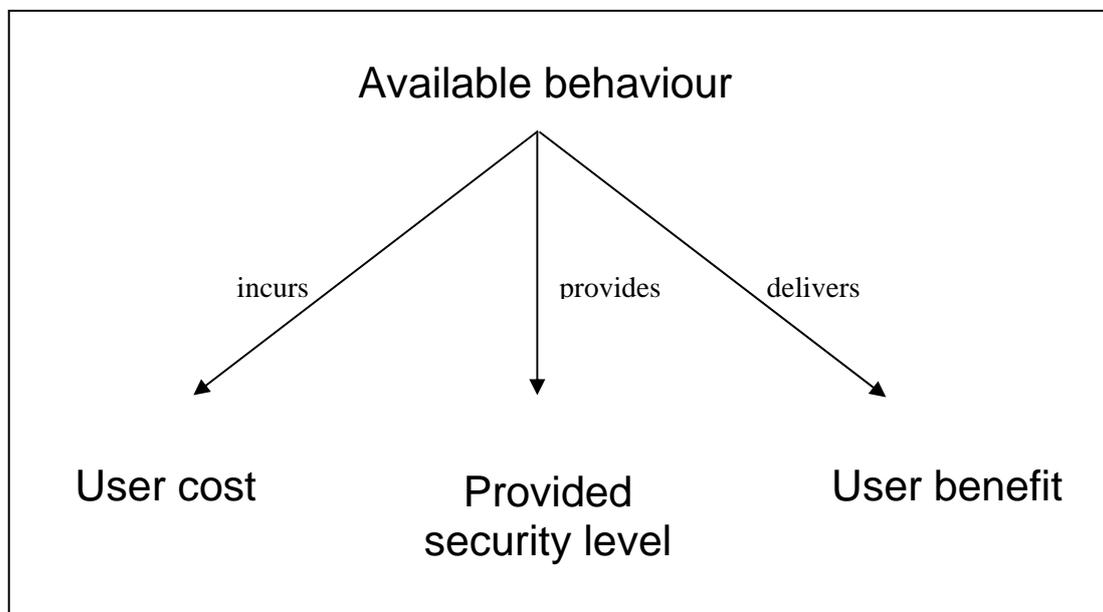


Figure 2: The three properties of available behaviours

Once a user has estimated the required security level of the computing resource, as well as the provided security level, user cost and user benefit of the various available behaviours, he will often find himself in a situation where he has to choose between behaviours that provide the level of security he requires, but come at a high cost, and behaviours that offer less security, but incur lower costs or grant certain benefits. An example is a user who is asked by a colleague to share a high-security password. If he shares the password, he will gain the benefit of improving his relationship with the colleague, while at the same time compromising security.

Refusing to share provides the high level of security he requires, but might come at the cost of souring his relationship with the colleague. The decision he will make is strongly influenced by his **risk personality**, in particular by the way in which he deals with low and uncertain risks. A user might be aware that choosing a less secure behaviour could lead to a security breach, but believes that *“it’s not gonna be me”*. Other users are willing to incur a higher cost, or forego a benefit, because they do not like taking risks for themselves. There are some indications from the data that this **risk personality** also can be observed in other areas of a user’s life. Some people write down the PIN number for their VISA card on a piece of paper in their wallet, whereas others do not. The fact that a user’s personality can play a part in how he chooses his password-related behaviours is vitally important, since it means that it is possible to deduce aspects of his personality from his behaviour. This means that password-related behaviours can have an effect on a user’s **self-image** and **public image** (see sections 6.2.2.3 and 6.2.2.4).

The storyline of **choice** can be summarized as follows (see also Figure 4, page 100 for a graphical representation of the **choice** process):

- 1) The user estimates the **required security level** of the computing resource protected by the password in question.
- 2) He makes an estimate of the **provided security level**, **user cost** and **user benefit** of each of the password-related behaviours that are available.
- 3) Based on his **risk personality**, he chooses either
 - a) the behaviour that provides the **required security level** (or comes as close as possible) and incurs the lowest **user cost** and/or the highest **user benefit**.

OR

 - b) a behaviour that provides a lower level of security, but also a lower **user cost** and/or a higher **user benefit**.

Figure 3: The storyline of choice

It is important to point out that some of the steps users undertake when making **choices** can themselves have become habitual. The best example of this is the **provided security level** of the various behaviours. Most commonly, users already have grouped these behaviours in the past, and refer back to this grouping when making a new **choice**, rather than creating it again from scratch. The amount of effort a user is willing to expend on the various steps that comprise the **choice** process can

itself be explained with reference to this very process, as will be pointed out in section 6.3.

The following section will discuss the factors that form the basis of choice in greater detail.

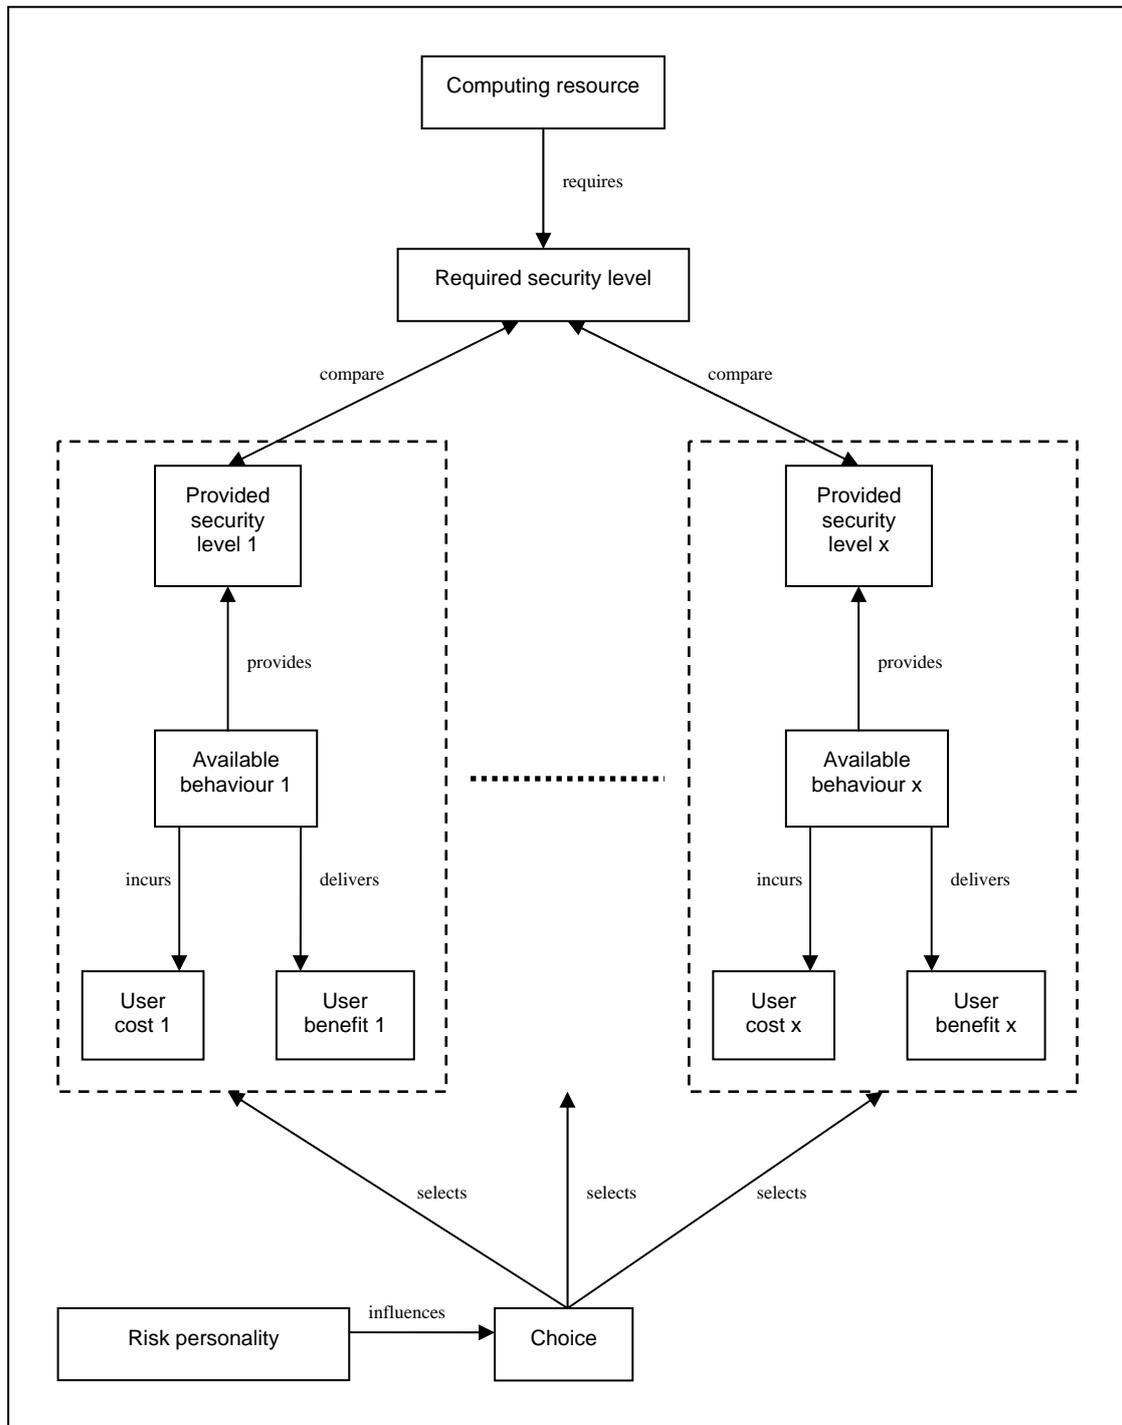


Figure 4: The factors influencing the choice process

6.2.2 The individual elements of the model

6.2.2.1 Required security level

Users assign a required security level to computing resources based on a combination of their estimates of two factors:

1. The likelihood of someone trying to attack the computing resource protected by a password (attack likelihood);
2. The possible consequences of a successful attack (attack consequences).

These will be dealt with in separate sections.

6.2.2.1.1 Attack likelihood

In order to explain how users estimate the attack likelihood, we first have to take a slight detour. Users perceive potential attackers of computing resources as having certain motivations (attacker's motivations). This determines their targets (attacker's targets), the methods they employ to gain access to computing resources (attacker's methods), and the actions they perform when an attack has been successful (attacker's actions) (see Figure 5).

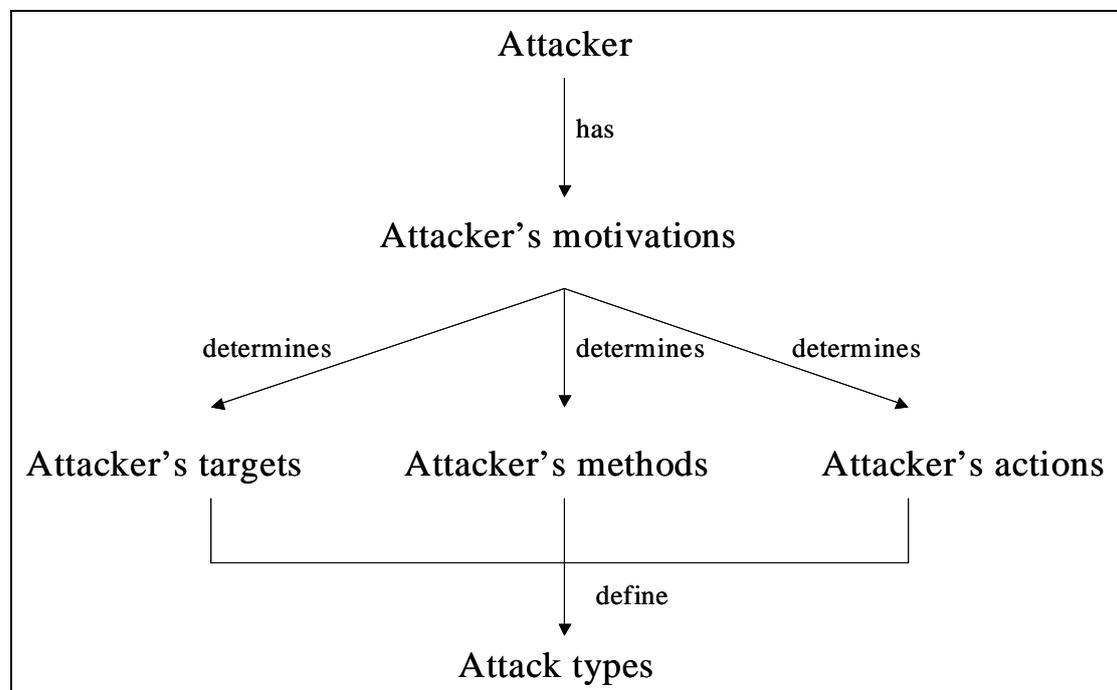


Figure 5: How the different attack types are constructed on the basis of users' perceptions of attackers

Chapter 8 will show in detail how these perceptions are connected to various interpretative repertoires, but for now it will suffice to say that there are five different attack types:

1. A **personal attack** targets one specific computing resource. The attacker might know there is something of value in that resource, or he may want to harm the owner of its password.
2. A **reward attack** targets any resource access to which offers the successful intruder some kind of benefit. An attacker might repeatedly try to find his way into various online banking accounts in order to obtain some financial gain.
3. A **group attack** tries to get into any one of a number of resources which belong to a group. The attacker wants to get access to one of them, but does not care which one. An example would be someone trying to get access to any account of a specific university department in order to hack his way into the system and obtain this year's exam papers.
4. A **cruising attack** targets just about any resource and is usually carried out by attackers with indiscriminate aims, such as young hackers trying to break into anything they can get access to.
5. An **opportunistic attack** is not pre-planned but simply exploits an easy opportunity that presents itself. The attacker finds a password written down on someone's desktop and uses it, even though he didn't go to the desktop with an attack in mind.

As can be seen, the first three forms of attack have specific targets, whereas the other two might be aimed at anybody. Users evaluate whether and to what degree a computing resource of theirs might be a potential target based on three factors:

1. The characteristics of the **computing resource** a specific password gives them access to (**resource characteristics**). These can be subdivided into **resource content**, **resource abilities** and **resource visibility**. Users assess whether the information accessible through the resource and abilities such as sending emails or signing off orders would be of value to a potential attacker. However, this in itself would not be enough to warrant an attack. The attacker also has to *know* that the resource has this content or these abilities (**resource visibility**). One of the reasons that users tend to be very concerned about their online banking password is the high visibility of the resource, which increases the likelihood of attacks.
2. The **personal characteristics** of the user of the password. The target of a **personal attack** is often seen to be somebody the attacker holds a grudge against. **Reward attacks** are frequently aimed at people who are "rich and famous", the reward being an increase in the reputation of the attacker.
3. The **group characteristics** of the user. **Group attacks** are often carried out against organizations that are seen to be unusually security-conscious, such as the Pentagon, since they present a larger trophy to an attacker breaching their defences. Alternatively, someone might want to gain access to a large company's central system and needs just one individual account connected to it as a starting point. Users often are part of many groups, and each one of them can increase the likelihood of them being targeted.

There is one additional factor which influences users' estimates of the **attack likelihood**. Most users do not consider themselves likely targets for a **personal**

attack, and tend to reduce the chance of an opportunistic attack by maintaining a minimum level of security for all their passwords. They can often guard themselves against a reward attack by employing safeguards, e.g. by taking sensitive information off a computing resource (see section 6.2.2.1.2). This leaves them open to one of the other two forms of attack, which are essentially based on a numbers game. Attackers are carrying them out, and some people *do* get hit by them. However, it is difficult even for seasoned security professionals to put an estimate on the likelihood of this happening to a specific individual. Users often consider any incident evidence they have of breaches that have occurred in the past together with the perceived security of their environment to get some kind of idea of the number of attacks being carried out. The perceived security is comprised of

1. physical security, such as guards on the premises;
2. software security, such as known weaknesses of operating systems;
3. and people security, meaning the extent to which other people in the organization are seen to make an effort to be security-conscious.

The logic here is quite simple. If the perceived security is low, it should be easy for attackers to breach it. A low number of known incidents then implies that the number of attacks must be low as well. It is important to note that hardly any of the users who participated in the studies had experienced any security breaches or had even heard of any having occurred in their close environment. Breaches reported in other organizations or in other parts of the users' organization do not necessarily increase the perceived attack likelihood. The fact that breaches that have occurred are often not made public by organizations does not help this situation.

The way in which users estimate the attack likelihood can be summarized as follows:

- | |
|---|
| <ol style="list-style-type: none">1) The user compares the resource characteristics, personal characteristics and group characteristics to his perception of attacker's targets and determines whether they make him a likely target.2) If the perceived security is low, and there is low incident evidence, he may reduce the attack likelihood. |
|---|

Figure 6: How users estimate the attack likelihood

6.2.2.1.2 Attack consequences

The damage that may occur from a successful break-in into a user's computing resource can affect

1. the user himself (consequences to user), such as when important files get deleted;

2. some other person or persons (**consequences to others**), e.g. someone whose job performance review the user stores or the other members of his workgroup with whom he shares access to a central file store;
3. the organization itself (**consequences to organization**), e.g. when large parts of the system are destroyed.

Obviously, the latter two are likely to affect the user indirectly through possible repercussions by the parties harmed. One finding is that users often are more concerned about **consequences to others** than about consequences to themselves. They may make a greater effort to protect personal information belonging to a colleague or friend than they do to protect their own. One possible explanation is that they consider it acceptable to take a risk on their own behalf, but not to do so where others are concerned. The extent to which concerns about **consequences to organization** preoccupy an individual user and make him willing to perform behaviours with a higher cost or lower benefit is determined strongly by his **responsibility personality** (see section 6.2.2.5). Some users regard it as their own responsibility to protect the organization as best as they can, whereas others believe that this cannot be expected of them. An example of the latter would be a user who chooses low-security behaviours to protect a certain **computing resources** even though he knows it would be a good starting point for an attack on the whole system. In his view, the organization should set up the system so that hacking from one resource to another is impossible, rather than expect him to incur overheads to protect it. Such a user will then estimate the **attack consequences** (and consequently the **required security level**) by only considering **consequences to users** and **consequences to others**, while completely ignoring **consequences to organisation**.

There are three factors that influence the severity of the damage that users expect to occur from a successful break-in into one of their resources:

1. The **resource characteristics** (see section 6.2.2.1.1).
2. The perception of **attackers' actions** (see Figure 4).
3. The **safeguards** which have been put into place by the user or his organization.

The **resource characteristics** obviously affect what an attacker actually can do. Focussing only on the possible **consequences to user**, the actions users typically are most worried about are:

1. personally sensitive information being divulged;
2. data being destroyed;
3. reputation-damaging actions, such as the sending of offensive emails to superiors;
4. financially damaging actions.

An interesting finding is that users are often aware that serious harm could be done to them and their organization once an attacker has gained access. However, they usually regard the most likely form of attack to be either a **group attack** or a **cruising attack** (see section 6.2.2.1.1). They would not expect such an attacker to make the effort that is needed to cause serious harm (**attacker's actions**). At worst, he might delete some information. In addition, a lot of users have **safeguards** in place. They might regularly back up their data, if the organization does not do so anyway. They might also keep data that they consider personal off their company computers. **Safeguards** effectively reduce the possible **attack consequences** and, as a result, the required **security level**.

The way in which users estimate the **attack consequences** can be summarized as follows:

- 1) The user determines which **attack types** he is likely to be subjected to (which is part of determining the **attack likelihood**).
- 2) He considers the actions that are typically undertaken in these **attack types** (**attacker's actions**).
- 3) He evaluates the possible consequences of a successful break-in to himself, others or the organization, based on the **resource characteristics**. Depending on his **responsibility personality**, he may ignore the possible consequences to the organisation.
- 4) If the possible **attack consequences** are unacceptably high, he may reduce them by employing **safeguards**.

Figure 7: How users estimate attack consequences

6.2.2.2 Provided security level

Users are fully aware that there is no perfect security. A sufficiently determined and skilled attacker will always get access to the **computing resource** he is targeting. The aim of making the effort to behave in a security-conscious fashion is not to make it impossible for an attacker to succeed, but to make it more difficult for him. The analysis has identified three strategies that are used independently or in conjunction to assign a **provided security level** to the various **available behaviours**:

1. Users often assign certain security levels to **available behaviours** without any explicit reasoning, based on **password regulations** and what they see happening in the **social context** they work in. The best example of this is the regular changing of passwords. Although almost all users consider this a high-security practice, few of them are able to explain why this is so.
2. As stated before (figure 5), users have perceptions of **attackers**, their motivation and the methods they employ to breach security. Each of these methods requires a

certain amount of effort, and security behaviours are also rated according to the effort that is required to exploit them. Writing down the password is a good example of this. For an attacker to exploit the various forms of this behaviour, varying degrees of effort are required. This makes it possible to assign increasing levels of security, e.g.:

1. putting down the password on a post-it note on the monitor;
 2. writing it down in a book close to the computer;
 3. putting it into the personal organizer, which is password-protected itself;
 4. keeping a copy of it at home.
3. Finally, users can observe the effectiveness of certain behaviours in their **social context**. If many people write down their passwords and there is no evidence of incidences occurring, then this behaviour may be seen to provide a high enough level of security.

These three strategies were clearly established from the data captured. The first one in particular warrants further research to identify the exact mechanisms by which users pick up their evaluations of security behaviours from their **social context** without understanding the reasoning behind it.

The estimate of the **provided security level** of various behaviours obtained by any of these strategies can then be altered by taking into consideration the **perceived security**. Interestingly, this can lead to less secure user practices both when the **perceived security** is low and when it is high:

1. When it is low, users might assign a low **provided security level** to all behaviours, and then choose the one with the lowest cost and/or the highest benefit, which in reality constitutes an insecure practice. For example, an organization that employs software systems that are known to have serious security flaws can be seen to make it easy for attackers to get access to its **computing resources**, which can result in password security as a whole, and all the behaviours associated with it, being regarded as a futile exercise by users.
2. When it is high, users might assign a higher **provided security level** to a number of behaviours than they would otherwise. For example, in an organization that employs security guards on the premises, users can regard it as unlikely that an attacker would get physical access to the offices. As a result, behaviours such as writing down a password on a note in a drawer next to the computer can receive a higher **provided security level** than they would in an environment where the **perceived security** is low. Obviously, this means that users can then meet what they consider the **required security level** with behaviours that are likely to come at a lower cost and/or higher benefit, but are in effect less secure.

The way in which users estimate the **provided security level** of a password-related behaviour can be summarised as follows:

- 1) The user makes a first estimate of the provided security level by employing one or several of the following strategies:
 - a) He bases the estimate on the password regulations and/or on what he observes in his social context.
 - b) He bases the estimate on the amount of effort that an attacker would have to make in order to exploit the behaviour in question.
 - c) He bases the estimate on behaviours that he has observed in his social context and which have not led to security breaches.
- 2) Depending on the perceived security, this estimate may be altered:
 - a) If the perceived security is high, the estimate of the provided security level may be raised as well.
 - b) If the perceived security is low, the estimate of the provided security level may be reduced as well.

Figure 8: How users estimate the provided security level

6.2.2.3 User cost

Password-related behaviours can incur four kinds of costs:

- 1) The primary cost of performing the actual behaviour, such as memorizing a password. This depends to a large extent on the usability of the mechanism and the abilities of the user (user's abilities). It is important to remember that the behaviour which the user wants to carry out may be one that his organisation does not want him to perform. The password mechanism may therefore consciously have been set up to enforce certain behaviours, which, from the user's point of view, makes it unusable with respect to alternative behaviours he might want to perform. An explanation of the aspects of password mechanisms that impact on their usability will make this point clearer:
 - a) A user choosing a memorable yet strong password may have his choice rejected by the proactive password checking mechanism, because the criteria used by that mechanism are either not explicitly stated or too unclear or difficult for the user to understand and satisfy. This increases the primary cost of choosing a strong yet memorable password, since the user will have to create and try out new passwords until one is accepted.
 - b) A user might want to choose a weak password, but the proactive password checking mechanism rejects his choice. This does not alter the primary cost of choosing a weak password, but it puts a high secondary cost (see point 2 below) on doing so – the user is not able to access the computing resource to achieve whatever goal he has in mind. This essentially is the effect of enforced password regimes on the primary and secondary cost of behaviours.
 - c) A user might want to choose a different password for each password-protected computing resource he accesses, but the sheer number of passwords needed to do so incurs a primary cost that is too high and might even make it impossible for him to do so. This situation is complicated further by the fact that passwords for systems that are used very infrequently (frequency of use) are more easily forgotten than passwords that are used frequently.
 - d) A user might want to choose the same password for different password-protected computing resources, but the proactive password checking mechanisms employed by these resources differ in their criteria (inconsistency of proactive password checking criteria). This means he

either has to choose different passwords for the systems (a different behaviour than the one he had originally chosen, which would have incurred a **secondary cost** of not being able to access certain resources) or create a password in a trial-and-error manner which satisfies the criteria of all systems (increasing the **primary cost** of this behaviour).

- e) A user might want to choose a strong password for the resource and memorise it. However, he knows that he often needs several attempts at entering this password, since he tends to confuse passwords for different systems. The mechanisms' **strike policy** blocks the account after a low number of strikes, and forces the user to write down the password to make sure he only needs one try to enter it. The **secondary cost** of staying with his original choice of behaviour would have been to be denied access to the resource on a regular basis.
- f) A user might want to keep his password for the resource the same over extended periods of time, but the mechanism's **change regime** forces him to change it at regular intervals. Again, the **secondary cost** of his original choice of behaviour would be not to have access to the resource. The fact that different resources may have change regimes with different change intervals in place, and that the timing of the changes may not be synchronised (**inconsistency of change regime**) makes the management of these changes even more difficult.

The user's abilities that impact on the primary cost of certain behaviours do not only include those abilities that make it possible for him to perform secure behaviours, but also those that allow him to perform insecure ones in the face of password mechanisms that aim to prevent just that. An explanation of the specific abilities users need will make this point clearer:

- a) A user wishing to use a password of a certain strength and memorability must be able to create adequate password content (**ability to create password content**).
- b) A user wishing to use different passwords for at least some of his resources and wanting to change them regularly has to be able to manage these passwords correspondingly (**ability to manage multiple passwords**).
- c) A user who is faced with a proactive password checking mechanism has to be able to create passwords that satisfy the criteria of that mechanism (**ability to satisfy proactive password checking criteria**), whether he wants his password to be strong or not. This ability can become even more stretched in the face of **inconsistency of proactive password checking criteria**.
- d) A user who is subjugated to resources that employ change regimes and who does not want to choose perfectly unconnected passwords has to find a way of connecting his passwords in some way (**ability to connect passwords**), e.g. by basing them on a theme or by indexing them.

The user's abilities can sometimes be improved, e.g. by the teaching of mnemonic techniques. However, there are cases where this is not possible. For example, some users are dyslexic and incapable of memorising any but the simplest passwords. The particular problems of this user group may seem obvious with hindsight, but had been so far overlooked in the research literature..

- 2) The **secondary cost** of a behaviour is that incurred as a consequence of performing it. For example, a user might be at a conference where he cannot get online access to his computer at work. However, he urgently needs access to his email system. Some of his options are

1. to drive back to his office and read the email, which incurs a high **primary cost** and a possible secondary one since he might miss parts of the conference;
 2. to phone in and give a colleague his password so he can read the email for him, which reduces both the **primary** and the **secondary cost** but provides a low level of security;
 3. to stay at the conference without accessing his email, which might incur an enormous **secondary cost** if he loses a large contract because he did not reply to the email.
- 3) The **social cost** of performing a behaviour that has an effect on his relationships with other people. An example would be a user who refuses to share one of his passwords with a colleague who urgently needs access to the **computing resource** in question. This might result in the colleague believing that the user does not trust him. Another important example is the **social cost** incurred by not complying with the **social norms** set in the **social context** users find themselves in. In an environment where a lax attitude toward password security is the norm, any user not participating in such practices can slowly become an outsider to the group.
- 4) The **image cost** of performing a behaviour that contradicts a user's **self-image** or the **public image** he wants to portray. Examples of this would be:
1. A user might have an image of himself as someone who is "non-technical" and "not a nerd". Any knowledge of security issues, or even interest in them, would clash with this image.
 2. Being particularly concerned about password security often carries the tag of "paranoia", which most users would refuse both as a **self-image** and a **public image**.
 3. Someone wanting to convey a **public image** of being a "team player" might find it difficult to refuse to share his passwords with the other members of his team.

6.2.2.4 User benefit

Password-related behaviours can provide two kinds of benefits to users:

1. The **social benefit** of performing a specific behaviour that has an effect on the relationship with other people. As an example, someone who freely shares his passwords with colleagues in need shows that he trusts them and can expect reciprocal trust.
2. The **image benefit** of performing a behaviour that reaffirms or improves one's **self-image** or **public image**. As an example, somebody who is outwardly seen to take care of his passwords can create a public image of being "professional".

6.2.2.5 Personality

There are two aspects of a user's **personality** that affect how he chooses his password-related behaviours, both of which have already been introduced and discussed in previous sections (see section 6.2.1 and section 6.2.2.1.2):

1. The way in which he deals with low and uncertain risks (**risk personality**), which plays a major part at the top-level of the model.
2. His willingness to take responsibility for protecting the resources provided by his organization (**responsibility personality**), which influences his estimate of the attack consequences.

6.3 Habitual choices and the recursive application of the model

The majority of password-related behaviours are not chosen by users consciously going through the individual steps described by the core model. Instead, users tend to choose habitual responses in situations that they have already encountered in the past. Even if they go through the **choice** process, they will often skip individual steps and rely instead on estimates they have arrived at previously. The best example of this is the **provided security level** of the various **available behaviours**. Most commonly, users already have grouped these behaviours in the past, and refer back to this grouping when making a new **choice**, rather than creating it again from scratch.

A user acting out of habit takes a risk, because the situation he finds himself in might actually be different from the one he was in when he established the habitual behaviour. At the same time, he reduces his cost, since he does not have to go through a process involving risk management and a cost/benefit analysis. The analysis of the data showed that this **choice** of not going through the **choice** process and of behaving out of habit instead can itself be described by the core model. However, the analysis only made it possible to ground a single recursion of the core model to itself in the data. This means that while further recursions are theoretically possible (i.e. users habitually behaving out of habit), they could not be found in the data collected in the studies.

As mentioned above, users behaving out of habit run the danger of maintaining habitual behaviours in situations which are actually different from the ones in which these behaviours were originally formed. The studies showed that this happens on a regular basis, but two particular contexts stood out:

1. A university graduate in his first job is likely to maintain at least some of the password-related behaviours he has employed in his years at university. The same could be observed for users changing jobs or even just moving department within one company.
2. A user who already has many passwords and is given access to a new **computing resource** with a new password is likely to treat it in a fashion that is similar to how he treats his other passwords.

The participants in the studies were clearly reluctant to change habitual behaviours. At the same time, on the few occasions when participants reported performing behaviours that required a relatively large effort on their behalf, they also realized that this had become easier once the behaviours had become habitual. The obvious conclusion is that good practices should be established from the very outset, and not only by companies who hire new employees, but also by the institutions that first

introduce users to passwords, such as universities. Chapter 9 will argue this point in greater detail.

6.4 Testing the waters

As discussed in the previous section, in most situations users choose their password-related behaviours out of habit. Obviously, new behaviours may be chosen whenever one of the factors of the model changes, and might then themselves become habitual responses. However, there is another way in which users might change their choice of behaviour over time. This process usually results in deteriorating password practices, and occurs for behaviours that users have chosen primarily because of their estimate of the **attack likelihood** and the **provided security level**. Common examples of this process are:

1. A user employs high-security behaviours for some of his passwords, since he considers the resources in question to be a potential target for attackers. However, he observes people in his **social context** and dealing with the same resources using low-security behaviours, with no security breaches occurring over a period of time. Consequently, he reduces his estimate of the **attack likelihood** or changes his perception of the **provided security levels**, and starts using low-security behaviours as well.
2. A user who is willing to take a certain degree of risk (**risk personality**) chooses to perform medium-security behaviours for a password which he believes requires a high level of security because of its potential attraction to attackers. If no security breaches occur, he decides that the **attack likelihood** is lower than he thought, and might even try out low-security behaviours for a while.
3. The previous case can also occur in a different form and with users who are not necessarily risk-takers. A user might employ high-security behaviours for a certain password. However, one day he has to employ a medium-security behaviour since the secondary **COST** of performing the high-security one would be exorbitant. Alternatively, he might use the weaker behaviour by accident. No security breach occurs, and he concludes that the **attack likelihood** is lower than he expected. As a result, he employs the medium-security behaviour in the future.

The pattern behind these examples is clear:

1. A behaviour of lower security than the user initially chose is carried out by him or people in his **social context**.
2. No security breach occurs for a period of time.
3. The user assumes this to mean that either the **attack likelihood** is lower than he thought, or that the new behaviour provides at least the same level of security as the one he originally chose.
4. He uses the new behaviour in the future.

The whole pattern has a quality of users either ‘testing the waters’ themselves, or observing other people doing so. The more blatantly security concerns are ignored

without negative consequences for the offenders, the faster this deterioration process is likely to take place.

6.5 Chapter summary

In this chapter, a grounded theory model has been presented which describes the factors users consider when choosing password-related behaviours. In the absence of any efforts by an organization to influence a user's choice, four factors have been identified: the required security level of the computing resource in question, and the provided security level, user cost and user benefit of the various available password-related behaviours. Users aim to employ behaviours that provide the levels of security they believe are required, while minimizing the user cost and/or maximizing the user benefit. In certain situations, they choose less secure behaviours in order to reduce costs or increase benefits. Their willingness to do so is strongly based on their risk personality. In most situations, users will at least partly choose their behaviour out of habit, and the way in which they make this choice can be described by a recursive application of the model to itself. An interesting strategy users reported to have used in the past involves them testing less secure behaviours in order to find out whether security breaches occur, and this strategy can be used to explain how password practices can deteriorate over time.

The model that has been presented can be used in three ways:

1. User behaviour in the absence of any efforts by an organization to influence it can be explained and predicted. This might seem to be irrelevant, since almost all organizations do employ measures that aim to improve user behaviour. However, the participants in the studies, and probably users in a large number of organizations, often ignored these measures and chose their password-related behaviours in a manner that can be described comprehensively by the model.
2. The model can form the basis of a larger, extended model, which incorporates the effect of any measures that organizations currently employ in order to improve user behaviour. The following chapter will provide one such extension, by incorporating the effect of password regulations and their associated punishment regimes into the model.
3. Finally, the model, either in its current form or with a number of extensions, can be used to design new, and more effective measures that are aimed at improving user behaviour. Chapter 9 will introduce the concept of *persuasive password*

security, which is an integrated approach that aims to improve user behaviour and which is based on the core model presented in this chapter and the extensions to it discussed in the following chapter.

7 EXTENSIONS TO THE MODEL

7.1 Overview

The previous chapter has presented a grounded theory model that describes the way in which users choose their password-related behaviours in the absence of any efforts by their organisation to influence this choice. However, almost all organizations *do* implement a number of measures to improve the password practices of the users of their computing resources. The most common ones are

1. Changes to the password mechanism itself, which indirectly affect users' choices by restricting the pool of available behaviours or by changing the user cost of specific behaviours (e.g. forced change regimes).
2. Security awareness and education programs, which aim to increase users' knowledge of security, as well as their ability to perform the required behaviours.
3. Regulations and their associated punishment regimes.

In this chapter, the core model presented in chapter 6 will be extended to incorporate the effect of regulations and their associated punishment regimes on users' choices.

This serves three purposes:

1. University College London and British Telecom both had regulations on how users should protect their passwords. However, these regulations were habitually violated (see chapter 5). Extending the model will make it possible to identify why the regulation-based approach was rather ineffective in University College London and British Telecom, and to develop ways of improving it. This information can be used by other organizations who face similar problems.
2. The aim of developing the model is to provide organizations with a tool that makes it possible not only to improve existing measures, but also to devise new and more effective ones. These will often have to work in conjunction with a regulation-based approach, which is the most common way in which organizations try to improve user behaviour. This means that the effect of such regulations, and their associated punishment regimes, needs to be understood before additional measures can be devised. It also is necessary to comprehend how past exposure to regulation-based approaches might have affected users' attitudes towards password security in general.
3. The core model has been designed in part with possible future extensions in mind. These include both existing measures, such as security education and awareness campaigns, and new approaches, which have yet to be developed. Extending the model to incorporate the effect of regulations and their associated punishment regimes will give a first indication of how easy it will be to integrate such extensions into the core model.

Section 7.2 will present the extended grounded theory model, which incorporates the effect of regulations and their associated punishment regimes on user behaviour. It will be shown that the model can be extended by adding just a limited number of additional elements, without the need for any structural changes to it. Chapter 9 will

discuss how this extended model can be used to improve regulation-based approaches, and how any new measures can work in conjunction with them.

The participants in the studies had not been exposed to any concerted security awareness and education efforts by their organization. While this means that the qualitative data gained in the studies does not make it possible to extend the model to incorporate these, it is still possible to give some first insights into the subject. These will be outlined in section 7.3. While the main function of this information will be to guide future research, it could also be used to improve existing security awareness and education campaigns in a limited manner.

In order for any measures aimed at improving user behaviour to be successful, users first of all have to be made aware of them. This can be difficult, because the strategy that users have employed to choose their weak password practices will also be used to determine their response to any such measures. This will be discussed in section 7.4.

The extensions to the model presented in this chapter are primarily based on the qualitative data collected in studies 3, 4, 6 and 7. However, they do also fit the limited amount of relevant data present in the interviews conducted in study 1.

7.2 Extending the model to incorporate the effect of password regulations and their associated punishment regimes

7.2.1 Extending the model

Current password regulations tend to take the form of a set of rules which the users are expected to follow for all their passwords and in all situations they might encounter. The most common rules are (see section 2.3.5.1):

1. Users need to choose cryptographically strong passwords.
2. They need to choose a unique password for each password-protected resource they access.
3. They need to memorise their passwords (as opposed to keeping a physical copy).
4. They must not share their passwords with third parties.
5. They need to change their passwords at regular intervals.

All of these behaviours provide a level of security which most users would deem to be high. Problems occur when the behaviour a user would choose in the absence of any regulations (as described by the core model) is different from the one prescribed by these rules. This will be the case when one or more of the following occur:

1. The user believes the **required security level** of the resource in question to be lower than the **provided security level** of the regulation-based behaviours.
2. He believes the behaviour he would choose in the absence of regulations to provide as much security as the regulation-based behaviours.

3. He considers the cost of the regulation-based behaviours to be too high, in particular compared to the behaviour he would choose in the absence of regulations.
4. He considers the benefit provided by the behaviour he would choose in the absence of regulations to be too high, in particular compared to the regulation-based behaviour.

In any of these cases, a user has to decide whether he will follow the regulations or his own judgement. This decision will be determined by the elements already present in the model, and two more which will have to be added to it:

1. The user's estimate of the risk of punishment he creates for himself by performing a behaviour which violates password regulations. This punishment risk is an additional quality which each available behaviour possesses (see Figure 9). Obviously, it is zero for behaviours that are in line with regulations.

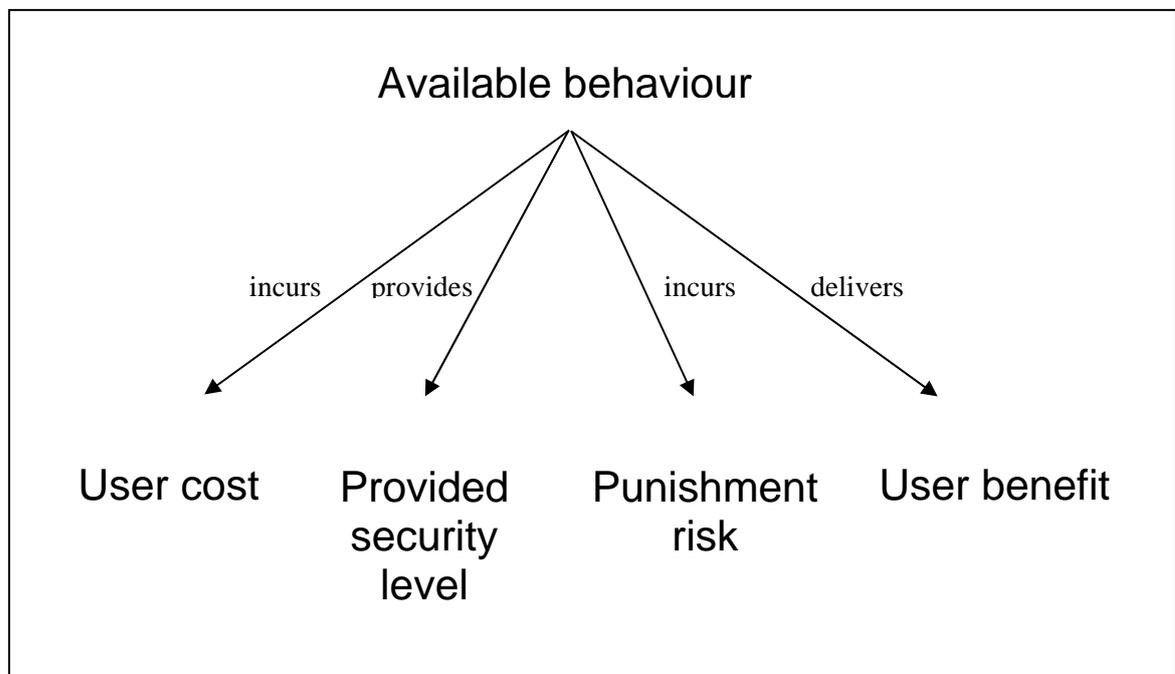


Figure 9 : The four properties of available behaviours

2. The aspect of a user's personality which determines how he responds to orders, in particular those that contradict his own, independent judgement. The participants in the studies took up three different stances: they complied with the orders, ignored them and followed their own judgement, or rebelled against them out of principle (see section 7.2.2.2). This authority personality works in conjunction with the risk personality and the responsibility personality to influence users' choices.

As can be seen, the model does not have to be changed structurally in order to accommodate the effect of password regulations and their associated punishment regimes. Instead, it is only necessary to add two more elements (and their sub-

elements) to it and to change the storyline as follows (see also Figure 11, page 118, for a graphical representation of the choice process in the extended model):

- 1) The user determines the required security level of the resource protected by the password in question.
- 2) He estimates the provided security level, punishment risk, user cost and user benefit of each of the password-related behaviours that are available.
- 3) Influenced by his authority personality,
 - a) he chooses the behaviour prescribed by the password regulations, as long as the user cost incurred by it and/or the user benefit of choosing an alternative behaviour are not too high.

OR
 - b) he rebels against the regulations out of principle, as long as the required security level is perceived to be low, and chooses a behaviour carrying an acceptable punishment risk which violates the regulations while incurring the lowest user cost and/or providing the highest user benefit.

OR
 - c) he relies on his own judgement rather than the password regulations, and, based on his risk personality,
 - i) chooses the behaviour that provides the required security level, carries an acceptable punishment risk and incurs the lowest user cost and/or provides the highest user benefit.

OR
 - ii) chooses a behaviour that provides a lower level of security and carries an acceptable punishment risk, but which also has a lower user cost and/or a higher user benefit.

Figure 10: Storyline of the extended model

There is an important case that this storyline does not capture. Many participants in British Telecom frequently found themselves in situations in which it was impossible for them to follow regulations, simply because the user cost incurred would have been too high. For example, one participant in the studies needed to access over thirty password-protected systems, most of which employed a monthly change regime. Obviously, it is impossible to create unique and strong passwords for all of these systems on a monthly basis and to memorize them. However, the participants in such situations believed the punishment risk to be low, and responded by violating the password regulations (point 3)b) or 3)c) of the storyline). What would have happened if the punishment risk had been high? In other words, how would users react if they were given regulations that are unusable and impossible to follow, while at the same time being punished severely for violating them? This question cannot be answered

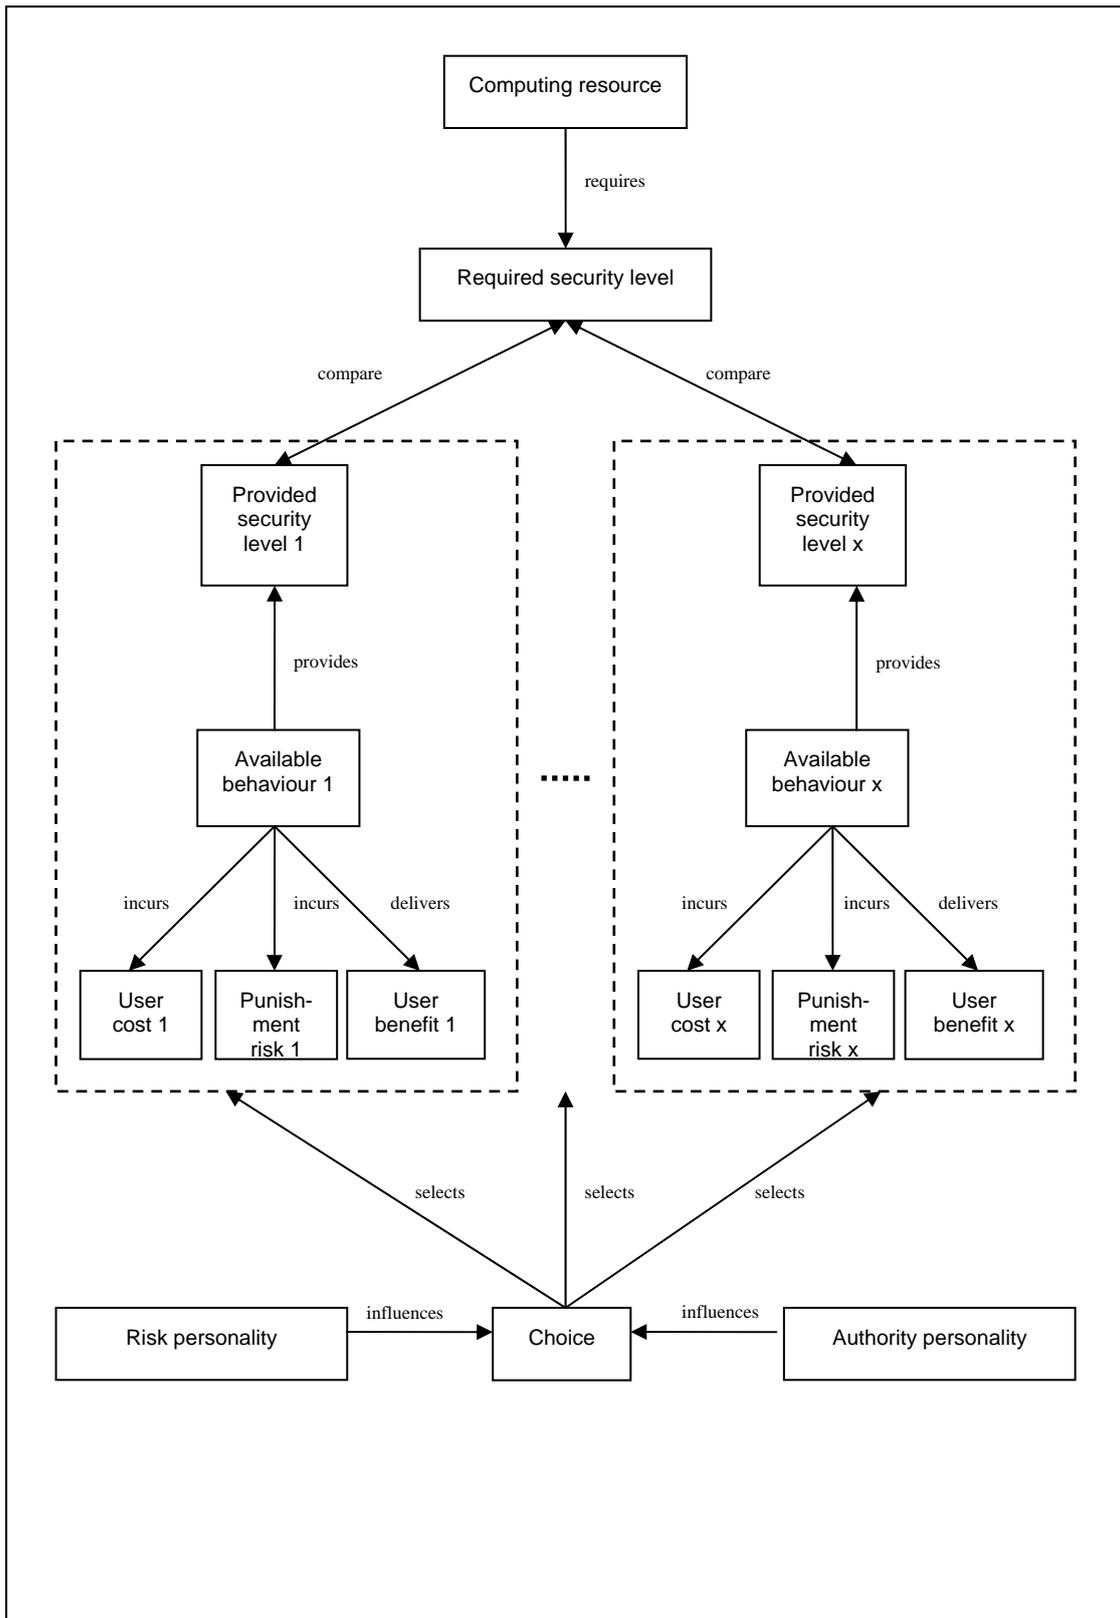


Figure 11: The factors influencing the choice process in the extended model

on the basis of the data captured, but it points to a fundamental problem in the relationship that has developed between security departments and the users of password systems. This will be discussed further in chapter 9.

The extended model is no different from the core model in that users might still behave out of habit, without going through the whole process described in the storyline. In the same manner, they might go through some of the steps, while relying on estimates arrived at in the past for others. They might also try out new behaviours that they believe to carry a higher **punishment risk** than the behaviour they have used so far in order to see whether any repercussions will materialise. These points will not be discussed further here, since the extension of the model has not changed the fundamental points made in section 6.3 and section 6.4.

The following section will discuss the two additional model elements, **punishment risk** and **authority personality**.

7.2.2 The additional model elements

7.2.2.1 Punishment risk

The perceived **punishment risk** associated with carrying out a behaviour that is not in accordance with regulations is based on users' estimates of two factors:

1. The perceived likelihood of actually being caught performing this behaviour (**detection likelihood**).
2. The severity of the punishment that is expected once one has been caught (**severity of expected punishment**).

7.2.2.1.1 Detection likelihood

In order to be punished for violating password regulations, users first of all have to be caught doing so. There are two ways in which this can happen. Firstly, a security breach can occur, with the subsequent investigation bringing to light previous misdemeanours. The perceived likelihood of this happening depends primarily on the **attack likelihood** of the resource in question (see section 6.2.2.1.1). Secondly, the organization might catch users in the act of performing behaviours that are against the rules. The perceived likelihood of this is determined by two factors:

1. The **ability to monitor** the behaviour in question.
2. The **detection effort** made by the organization.

Users are fully aware that a lot of the behaviours set out in password regulations are difficult, if not impossible, to monitor (**ability to monitor**). This means that various behaviours can be ranked according to the effort necessary to detect them, e.g.:

1. writing down a password on a post-it note attached to the monitor can be discovered by accident;
2. putting it into a book on the shelf next to the monitor requires someone consciously looking for it;
3. putting it onto a piece of paper in the wallet requires a strip-search.

There are three factors that determine how much effort users expect their organization to make in order to detect password violations (**detection effort**):

1. **Visible detection efforts**, such as when cameras are placed in computer clusters;
2. Any evidence of users having been punished in the past (**evidence of punishment**);
3. The perceived importance of security to the organization.

The last factor depends on the physical and software security provided by the organization, as well as the visible detection efforts and the awareness efforts undertaken by it. Organizations deploying resources on these, e.g. by running a lot of awareness and education campaigns, are seen to assign a high degree of importance to security, and might be expected to expand more means on detecting password violations as well.

Evidence of punishment plays a role similar to the one performed by incident evidence, which has been described in section 6.2.2.1.1. Users are often aware of how well people in their social context follow regulations, because they have either observed their behaviour or talked about it to them. If many people are seen to break regulations on a customary basis, and there is little or no evidence of anybody being punished for this, users tend to assume that the detection likelihood is low.

The way in which users estimate the detection likelihood can be summarised as follows:

- 1) The user estimates the risk of being found out due to a security breach, which depends on the **attack likelihood** associated with the resource in question.
- 2) He estimates the risk of being found out in the act by determining the effort necessary to detect the various password-related behaviours (**ability to monitor**) and by setting this off against the effort he expects his organization to make (**detection effort**).
- 3) He compounds these two risks.

Figure 12: How users estimate the detection likelihood

7.2.2.1.2 Severity of expected punishment

Users violating password regulations usually consider the danger of being punished for

1. breaking regulations as such;
2. a break-in due to the violation of regulations;
3. the consequences of such a break-in for the organization.

The extent to which they expect to be penalised for each of these depends on the following four factors:

1. Regulation characteristics;
2. Social context;
3. Importance of security;
4. Evidence of punishment.

There are two characteristics of regulations that have an effect on users' estimates of the severity of expected punishment:

1. Regulations that are unclear about the exact behaviours that users are expected to perform (**clarity of regulations**) leave room for interpretation, which in the eyes of some users means that they cannot be punished for actions that fall into grey areas. Even more important than this are exact statements about the kind of punishment the various violations will incur (also part of **clarity of regulations**). In the absence of these, users tend to assume that they will be let off "with a slap on the wrists" when caught for the first time, and that more serious repercussions will only be incurred by repeat offenders.
2. Regulations that vary between systems or are changed regularly (**consistency of regulations**) can undermine the belief in their own importance and authority, which reduces the **severity of expected punishment** as well. Users do not expect to be punished gravely for what may be a password violation now, when the same behaviour was acceptable only a short while ago or is still acceptable on other, seemingly identical systems.

The **social context** in which a user finds himself in also plays an important role. In an environment in which violations are common, users do not expect to be made scapegoats for behaving in a manner that is no different from anybody else's, just because they were "*unlucky to get caught*". In particular, this rules out being punished for any security breaches or their consequences, since they "*could have happened to anyone*". However, if an organization is seen to take security seriously (**importance of security**), it will be regarded as more likely to punish violations gravely. Any past **evidence of punishment** also gives an indication of what kinds of punishments are to be expected by offenders.

The way in which users estimate the **severity of expected punishment** can be summarised as follows:

- 1) The user takes the regulations and the punishment they threaten as a first yardstick for the **severity of expected punishment**. However, if they are unclear or inconsistent they have less effect on his estimate or might even reduce it.
- 2) He determines how important security seems to be to his organization (**importance of security**). If it is seen as important, he increases his estimate.
- 3) He determines whether there is any evidence of punishment taking place (**evidence of punishment**). If there is not, he reduces his estimate, in particular if people in his **social context** customarily violate regulations.
- 4) He observes how other people in his **social context** behave and does not expect to be punished severely for behaving like the majority of them.

Figure 13: How users estimate the severity of expected punishment

7.2.2.2 Authority personality

The authority personality of a user determines the way in which he responds to orders in general, even though it becomes particularly acute when these orders contradict his own judgement. Three types of responses were encountered in the studies:

1. Relying on one's own judgement to choose an appropriate behaviour and following this judgement, as long as the **punishment risk** is seen to be low.
2. Following the regulations where the **user cost** of the regulations-based behaviour or the **user benefit** of the alternative behaviours is not too high.
3. Rebelling against the regulations by violating them consistently and out of principle, as long as the **punishment risk** and the **required security level** are perceived to be low.

The last response is important and may give cause to serious problems in the future. It seems that organizations are losing the support of an increasing number of their users in security matters, simply because they have bombarded them with regulations that are unclear, change constantly, do not make sense to them and incur high **user costs**. It might be difficult to win these people back, a point which will be discussed in chapter 9.

7.3 Remarks on security awareness and education campaigns

University College London did not expose its students to any security awareness and education campaigns. A small number of the participants from British Telecom had seen security videos or participated in security seminars. The qualitative data gained in this fashion does not make it possible to extend the model to incorporate the effect of security education and awareness campaigns with any confidence. What can be done at this point is to present some first insights that might serve as pointers for future research:

1. Organizations usually try to influence a wide variety of their employees' behaviours. Members of large companies, for example, regularly are sent emails concerning issues ranging from fire and safety regulations to security problems with the parking lot. Any campaign aiming to change password-related behaviours competes for the attention of its targets with other campaigns that are being carried out. In order to be noticed and taken in by users, it has to find a way of sticking out in some fashion.
2. The **social context** of users informally provides them with a lot of their knowledge about security issues, and also guides some of them in choosing their password-related behaviours. Messages that come from outside of this environment have less effect than those coming from within. In particular, line managers seem to be the most potent messengers for their workers.
3. The more personal the message and its presentation, the stronger is its effect on user behaviour. Ideally, line managers would sit down with their workers individually or in small groups and present the message to them, pointing out its relevance to their specific situation. Impersonal circulars, on the other hand, are in most cases not even read by the recipients.
4. It seems to be important that users are reminded regularly about password-related issues. One-off events, such as the annual one-hour lecture on the subject, seem to have little long-term effects. At the same time, reminders should not be too frequent, since users then give them and their message less attention.

The limited effect of security awareness and education campaigns can in part be explained by their violating these basic points. However, there is an additional explanation, which will be discussed in the following section.

7.4 Reaching out to users

In order for measures such as security awareness and education campaigns or regulations and their associated punishment regimes to be successful, users first of all have to be made aware of them. An organization might, for example, put into place a new set of regulations and send these out to its employees via email. Alternatively, it might run a number of seminars on security issues, explaining their importance and teaching the participants the various behaviours that are required to behave in a security-conscious fashion. For these efforts to have an impact on user behaviour, the employees need to take in the information presented to them in a conscious fashion. They have to read the email sent to them. They need to pay attention in the seminar, and integrate the learned material into their working life. However, reading an email about password security or paying attention in a security seminar are themselves password-related behaviours. This means that the strategy that users employ to choose password-related behaviours, as described by the grounded theory model, will also be used to determine the response to any efforts by the organization to change this very strategy. This can lead to catch-22 situations:

1. A user who is not concerned about password regulations because he believes the punishment risk of most behaviours to be low might not read any emails from the security department, since he considers the **primary cost** of doing so to be too high. As a result, he never reads the email which explains to him that new regulations have been introduced.
2. A user who is not concerned about password issues because he believes the **attack likelihood** to be low might mentally “drop out” of the security seminar and think instead about the work he still has to do afterwards. He regards the **primary cost** of consciously attending to the seminar as too high, and also does not want to incur the **social cost** of being regarded as a ‘swot’ by his colleagues. As a result, he never takes in the information that might change his estimate of the **attack likelihood**.

These examples show that organizations have to plan the measures they employ to improve user behaviour carefully. It is not good enough just to supply users with the necessary information, since the ones who are most in need of it are also the most likely to ignore it. Chapter 9 will discuss ways in which this problem can be addressed.

7.5 Chapter summary

In this chapter, the core model presented in the previous chapter has been extended to incorporate the effect of password regulations and their associated punishment regimes on users’ choices of available behaviours. This has served three purposes:

1. The new model can be used to improve regulation-based approaches.
2. New measures aimed at improving user behaviour that have to work in conjunction with regulation-based approaches can be developed using the model.
3. It has been shown that it is possible to extend the model without the need to make any structural changes to it.

This chapter has also presented first insights into how security awareness and education campaigns can be improved. These can be used both by the designers of such campaigns and by researchers wishing to investigate the subject further.

Finally, it has been pointed out that users will employ the same strategy that they use to choose their password-related behaviours, as described by the model, to determine their response to any efforts by the organization to influence that very strategy. This can lead to a catch-22 situation, where the users who most need to be reached by such efforts are also most likely to ignore them.

8 A DISCOURSE-ANALYTIC STUDY OF USERS' ACCOUNTS OF PASSWORD SECURITY ISSUES

8.1 Overview

In chapter 5, it has been shown that the participants in the studies habitually violated the most common password rules. Chapters 6 and 7 have presented a model that describes the way in which users choose their password-related behaviours. This model can be used to explain the ubiquity of weak password practices, and it can also be employed to develop organizational measures that are aimed at improving these practices. It is based on a grounded theory analysis of the accounts of password security issues users have given in the interviews and focus groups conducted as part of this research. In the current chapter, the results of analysing the same accounts using discourse analysis will be presented. The purpose of carrying out such an analysis is twofold:

1. Section 3.4.1 introduced the idea that the unintentional use of specific interpretative repertoires can shape the way we experience the world and behave in it, and might also have repercussions we might not have anticipated. A discourse-analytic study of the data collected in this thesis makes it possible to determine whether the interpretative repertoires participants drew on to describe certain aspects of password security might have an effect on their estimates of the various factors of the grounded theory model which strongly determine their choice of password behaviour, such as the **user cost** or the **provided security level**. In order to improve user behaviour, it might not be enough to change the reality users find themselves in, e.g. by improving the usability of the password mechanism. A change in the interpretative repertoires users draw on to describe those aspects of reality that influence their choice of password-related behaviours will also be necessary.
2. The grounded theory model is based on an analysis of the accounts users gave of various password-related issues. However, these accounts cannot necessarily be taken at face value. It is possible that they were constructed using specific interpretative repertoires and discursive techniques (see section 3.4.1) in order to justify malpractice. This is not necessarily a conscious deception, but could happen unconsciously, and in a manner which is shared within social groups. In this research, discourse analysis has been applied to identify cases of this.

An exposition and discussion of all the interpretative repertoires and discursive techniques identified in the data is not possible in the space of this chapter. It also would not be desirable, since the analysis has found a considerable number of repertoires and discursive practices that, while being of interest, only have a limited discernible effect on user behaviour, and cannot be conglomerated to address larger issues. The focus here will be on three particular issues:

1. A user's perception of potential **attackers** of computing resources strongly influences his estimates of the **attack likelihood** (see section 6.2.2.1.1). Section 8.2 presents the interpretative repertoires that the participants employed to describe potential attackers, and shows how they contribute to their estimates of the **attack likelihood**. This is a good example of how the exclusive use of impoverished repertoires can lead to estimates that result in the **choice** of insecure password-related behaviours. It opens up avenues for improving user behaviour by introducing valid repertoires in security awareness and education programmes.
2. In some situations, it is useful to distinguish between what I have named *resource repertoires* and *problem repertoires* (see section 3.4.1). Employing the former in a specific situation will lead to less secure password practices than employing the latter does. Security awareness and education programs, as well as the general discourse about password security should encourage users to utilize *resource repertoires* in order to improve their password-related behaviours. Section 8.3 will give examples of this.
3. Section 8.4 will ask whether participants did indeed construct their accounts using specific interpretative repertoires and discursive techniques in order to justify their malpractice. It is impossible to answer this question conclusively on the basis of the data collected for this research, but it will be argued that there are indications of such discursive practices taking place.

Some of the results presented in this chapter have been published in Weirich & Sasse (2001b), Weirich & Sasse (2001a) and Sasse et al. (2001).

8.2 Impoverished repertoires used to describe potential attackers

It was possible to identify eight clearly distinguishable interpretative repertoires that were used by the participants to describe potential **attackers** of computing resources. The exclusive use of some or all of these repertoires directly implies that attacks are based on **resource, personal and group characteristics**, which most of the participants did not think they possessed. As a result of this, their estimate of the **attack likelihood** and subsequently the **required security level** was low. One way of improving password practices then might be to educate users in a fashion which introduces and reinforces repertoires that describe attackers who might target users with resource, personal and group properties similar to those of the participants. The following sections will briefly present the eight repertoires found. The most prominent one was *kids*, with all the other ones coming a distant second.

8.2.1 Kids

Kids were described either as 'sad little kids' ...

*"spotty little s***s, basically, that have nothing better to do than to fascinate themselves by writing programs about how to get into things which they're not supposed to."*

...or as 'curious kids':

“Curiosity, just saying ‘This is secret, can I break into it, that would be fun.’ Like, basically, kids playing around.”

Some technically-minded participants even expressed a certain degree of admiration for them:

“Very technically literate, very capable technically, with a devious mind. [laughs]”

Kids’ motivation is to prove they can do it, to get a buzz, to get a sense of achievement, or to be better than someone else and impress their friends. They attack organizations that are high-profile or known to be very security-conscious...

“They’re high-profile. Some of those are supposed to be very secure, like the Pentagon is supposed to be unbeatable, so if you can get into that, it’s like a big thing, a big macho-thing, but also if you tell your mates you hacked into some system that nobody has ever heard of, they won’t be very impressed, however secure it was. They won’t be very impressed. You hacked into X, who is X? No-one’s heard of them, so it’s not very impressive.”

...or the rich and famous:

“If, say, my name was well-known, it’s a case of, if, say, my name was Bill Gates, a lot of people no doubt would want and try to break into Bill Gates’s account, to say ‘Yes, I’ve done it.’”

In addition, they might target people indiscriminately just to hone their skills:

“Ahhmm, I think a lot of it goes on in computer science departments, there’s a lot of guessing of other people’s passwords, so it’s, ahm, it’s like playing with the concepts that you’re learning about, isn’t it.”

Once they have broken into a system, they might deface a web page, or leave a message, but they do not do any serious harm.

8.2.2 Vandals

Vandals are seen as ‘abnormal’:

“I don’t know how to describe them. They’re obviously not normal people.”

They want to have a pop at the establishment or are just plain mad:

“... but all the destructive stuff is like a cat burglar that’s just having an episode in a place, you know, they lose their rag, they go completely mad and start wrecking the place, that seems a bit unnecessary to me.”

They have the same targets as kids, but unlike them, they do serious damage in the systems they break into.

8.2.3 Criminals

Criminals were mostly seen as trying to gain a financial advantage by getting access to bank accounts (which none of the participants had direct access to from their organizational accounts):

“...other people presumably have the intention of remaining, you know, at the other end of the spectrum, remaining completely undetected to some probably financial advantage of their own. Get into some bank account and shift money around.”

Some participants from the university environment also saw the possibility of someone trying to copy coursework or to get access to exam papers. A few people considered the threat of their computing account being used as a base to commit fraud undetected. The targets of these criminal acts are people who have access to the resources that are to be exploited.

8.2.4 Vengeful people

Vengeful people are vengeful against a specific individual:

“... basically, some thing, some point, some personal reason as well, like, they didn't like me, basically. [laughs]”

In order to be a target for such an attack, a user has to believe that there are people in his environment who dislike him or bear a grudge, which was not the case for any of the participants.

8.2.5 Disgruntled employees

Disgruntled employees want to get revenge on an organization:

“They might target organizations that they have a relationship with and if, for example, they feel begrudged in some way. So they can do some damage to somebody's system, because they didn't pay them enough, consistently underrated their work, consistently told them that they were foolish.”

These people would not target specific individuals to harm them, but might just use access to their resources in order to do damage to their organization.

8.2.6 Industrial spies

Industrial spies try to obtain very specific information:

“Ahm, there maybe opportunities for getting information such as, you know, industrial espionage, like designs and so forth. Ahm... I, I think anybody doing that, they would want to target exactly what they were going for, they'd have to know what they were looking for, you know, so I suspect that would be a particular, that isn't, that isn't gonna be casual, casual, ahm, a casual, ah, person. That would be a, a pretty

determined criminal act, and I think there might be better ways of doing, than getting a password. [laughs]”

Such attackers might target anybody in an organization as a starting point into other parts of the system. However, they usually know exactly what they are after and will tailor their attacks to reach those specific resources.

8.2.7 Terrorists

Terrorists target political organizations or those that commit what they consider dubious acts:

“If I was working, for example, if I was working, say, you know, in a department where we did animal rights experiments, sorry, animal experiments and we were the subject of malicious attack by animal rights investigators, and I think that would increase my chances of being exposed to the mass, you know.”

Again, such attacks are not directed against an individual but might use access to a user’s resources to harm his organization.

8.2.8 Jokers

Jokers play practical jokes on their acquaintances and friends:

“And they might even know the person that they’re targeting, where it’s just a joke, where they then send of an email purporting to come from some individual, saying outrageous things.”

The nature of a practical joke implies that no serious harm will be done to either the individual target or his organization.

8.3 Resource and problem repertoires

Resource repertoires are those which, when held and applied by a user in a specific situation, will increase the likelihood of him behaving in a secure fashion. *Problem repertoires* have the opposite effect. Sometimes, a *resource repertoire* can act as an antidote to a related *problem repertoire*, and occasionally the two might be mutually exclusive. Security campaigns should aim to introduce and reinforce *resource repertoires* while inoculating users against *problem repertoires*. In this section, a few examples of these two types of repertoires will be presented.

A good example of a repertoire that users employ and which almost inevitably has a detrimental effect on their security efforts is passwords are not infallible:

“Ahm, I mean, most, we work on, ah, in a company like BT, in such a big company, ahm, that sort of stuff may happen, and people are

aware that passwords are not infallible and therefore there is kind of a, a trust among people, and if I said 'I didn't do it', then I would expect people to trust me, because, ahm, basically, the, the, it should be clear that systems are not totally infallible and some systems can be compromised."

Users holding this repertoire can easily use it as an excuse for weak password practices, both to themselves and to their organization. A security breach that occurs is not necessarily due to their misbehaviours, but might be the result of the weaknesses that passwords inherently possess. In addition, incurring a possibly high user cost or forfeiting a user benefit by following regulations might not be worthwhile when the underlying security provided by passwords is quite low to begin with. These excuses are more difficult to use when the *passwords are secure* repertoire is employed instead, which does not claim passwords to be perfectly secure but stresses their relative merit:

"I think an eight-digit, sort of arbitrary password is, I think it's pretty secure, to be honest with you [...]"

An interesting point here is that security education might actually have a detrimental effect if users believe passwords to be fairly secure and are then taught how easy it can be to crack them, as is often the case in such programs. On the other hand, security efforts might also be undermined by stressing the security of the systems:

"Ah, perhaps not so important is, to me, is the passwords dealing with computing security in terms of files, file storing places because, mostly, because we're inside an intranet, it's mainly secure from outside."

These sorts of dilemmas facing security campaigners have first been identified by Adams et al. (1997), and will be discussed in greater detail in chapter 9.

Users often state that they violate password regulations because they are impossible to follow in certain situations, implying that there is no alternative behaviour which complies with regulations. However, some users hold the *can always find an alternative way* repertoire instead:

"[...] so, I wouldn't, if I needed really to read my email or something then I would find another way to do it, not by giving somebody my password to access the system."

Often users observe password violations in their social environment and assume that nothing can happen to them if they behave in the same manner, i.e. they believe the *punishment risk* to be low. The *just because everyone else does it* repertoire weakens that belief:

A *“I don’t know how I could defend myself. The fact that I’ve written them down for anybody to find... I, I just don’t know, don’t know. I’d have to think quite hard about that one.”*

Q *“Okay. I think what a lot of people just say is ‘I would just say: okay, that’s what everybody does.’, so...”*

A *“Oooh. Yeah, but it, just because everyone else does it...”*

An interesting repertoire is I don’t want to become a suspect. Most security campaigns tell users not to share their password, but sharing would be impossible if nobody was willing to receive a password:

“... and if something happens to that other person’s account, then you could be somebody who would become, would be a suspect in that situation, so I don’t try to get information about other people’s security information or password information”

Finally, one of the more obvious ways in which repertoires shape users’ reality and affect their password behaviours is in their use to describe people who behave in a certain fashion. Somebody might use the I am a trusting person repertoire to explain why he easily shares passwords, which offers him a clear image benefit:

“I guess, that’s pretty close to what I do. Somebody who is either trusting of other people and assumes other people at work will not mess things up without a very good reason. Somebody who wants to either help other people or at least not hinder them by saying ‘If somebody really needs to access my system, they probably can do.’”

In the same way, someone refusing to share their password might be seen as paranoid and incur an image cost:

“A rather overcareful, overcautious person, rather paranoid person, I think.”

As long as users hold these repertoires, the effect of password-related behaviours on their self-image and public image will increase the likelihood of insecure practices. Social marketing campaigns might be a fruitful way of connecting positive images with proper password practices.

8.4 Questioning users’ accounts

In the previous two chapters, a grounded theory model of how users choose specific password-related behaviours has been presented. In the preceding two sections of this chapter, discourse analysis has been used to identify interpretative repertoires that can influence this choice. Both analyses have used the same corpus of data, namely the transcripts of the interviews and focus groups carried out as part of this research. In this section, this data will be investigated again, but now with a different question in mind.

One of the central tenets of discourse analysis is that people construct accounts in order to perform actions, such as accusing, denying or justifying (see section 3.4.1). This action-orientation of discourse means that it cannot always be taken at face value, but might need to be questioned as to its underlying purpose. This section will ask the question of whether participants did indeed construct their accounts using specific interpretative repertoires and discursive techniques in order to justify their improper password practices. It is impossible to answer this question conclusively on the basis of the data collected for this research, but it will be argued that there are indications of such discursive practices taking place. This makes it possible to give first pointers for future research. It also encourages a more sophisticated approach to the use of interpretative repertoires in attempts to change user behaviour than would be taken if the accounts had not been questioned this way (see chapter 9).

The analysis in the remainder of this section will focus on three exemplary points:

1. The exclusive use of impoverished repertoires by participants and the way in which this can both cause and justify insecure password practices (section 8.4.1).
2. The use of *problem* rather than *resource repertoires*, which causes insecure password practices and justifies them at the same time (section 8.4.2).
3. The way in which participants in studies 3, 4, 6 and 7 reacted to any proposals to introduce punishment regimes (section 8.4.3).

8.4.1 Using impoverished repertoires

The exclusive use of the repertoires presented in section 8.2 directly implies that the threat of an attacker targeting one of the participants is low. In particular, it is highly unlikely that he would do so to get access to the resources provided by the organization. At the same time, the repertoires imply that an attack on personal resources, such as online banking, is more likely. This is not the place to discuss the validity of these assumptions, but it is clear that they can be used to justify weak password practices where organizational resources are concerned. In study 7 various single-sign on mechanisms were proposed to the participants, who showed serious concern when what could be considered personal resources, such as employment details, were included. They wanted to protect them with a separate password. In a way, large passages of the accounts given read like a justification of weak password practices by claiming that there is no real threat to the organizational resources. This means that an education campaign aiming to introduce repertoires of **attackers** that imply that the participants can realistically expect to get targeted could fall on stony ground with many users, who might just refuse to employ these new repertoires.

8.4.2 Using problem repertoires

In section 8.3, the related concepts of *resource* and *problem repertoires* have been introduced, and examples of some of these have been given. These repertoires can be thought of as being virulent within specific environments, be it a user's immediate work surroundings, his whole organization, or even his country or cultural sphere. Users pick up the repertoires they employ - both consciously and unconsciously - from these environments. As a result, they behave in a certain fashion. An obvious way of improving password practices would then be to educate users in a way that introduces *resource repertoires* and inoculates against *problem repertoires*. However, this might not work. The use of problem repertoires often does not only lead to weakened password practices, but also justifies them. Users who are not concerned about the resources provided by their organization, or who are willing to take risks with them, which they would not take with personal resources, might not want to admit this openly, and could instead construct their accounts using *problem repertoires*, which imply that there is no point in being too concerned about password security. Adams et al. (1997) have pointed out that password practices might become less secure both when the **perceived security** of the organization is high *and* when it is low. Their explanation for this is that in the former case users perceive the threat to the data to be greater, while in the latter case they reduce their estimate of its sensitivity. However, there is another possible way of explaining this. It might be possible that some users justify their weak password practices by claiming that, since the overall security of the system is low, there is no point in making the effort of keeping their password particularly secure. In an environment where the **perceived security** is high, they might use this fact to justify their behaviour by asserting that this makes them more susceptible to attacks from hackers. In conjunction with the use of the `hackers will always find ways in repertoire...`

“Ahm, I think if somebody is determined enough to break into a system, they will expand the effort, either guessing the passwords or rampaging through bins to find those torn-up envelopes or, or whatever. I think if somebody is determined enough, they'll break in. “

...this as well means that there is no point in being security-conscious. In other words, any level of the **perceived security** can lead to insecure password practices and might be used to justify these. This sort of dilemma will be discussed further in Chapter 9.

Adams et al. (1997) have also pointed out that users often see confidential information about individuals (such as personnel files) as sensitive and worthy of protection, whereas commercially sensitive information can be regarded as less sensitive. This has been confirmed by the research conducted for this thesis (see section 6.2.2.1.2). However, the way in which users classify information does not only have an effect on their behaviour, but can also be used to justify malpractice. There are numerous examples of users classifying organizational information in a manner which can at best be described as inconsiderate to the organization's concerns. The following quote also shows the use of the they wouldn't know how to exploit it repertoire in order to play down the importance of the information talked about:

“... but, but I mean, the sort of information that you, that's passworded is not of any interest to anybody. The number of man-hours that have been working on my project, who cares? There are items there that are important to me, and that I would know how to exploit them, but if somebody had a look at them, I think they would have great difficulty, first of all, in understanding them, and secondly, finding a market for them.”

It is clearly possible that users choose repertoires to justify their malpractice. This means that education campaigns introducing *resource repertoires* might have little effect on users, since they can always refuse to employ these repertoires. However, this would be more difficult if the whole discourse about password security within an organization were consistent in its use of repertoires, and in particular in its condemnation of those that lead to weak password practices. This will be discussed in more detail in chapter 9.

8.4.3 Resistance to punishment

The participants in all the studies customarily broke password regulations and had done so for considerable amounts of time. Studies 3, 6 and 7 were partially designed to find out how they would react to attempts by their organization to change their behaviour by introducing harsh punishment regimes. The results have contributed strongly to the extended grounded theory model (chapter 7). Here, the data will be examined using discourse analysis to focus on how participants might have constructed their accounts in order to reduce the chance of punishment regimes being introduced. Two points will be dealt with:

1. The proposals some participants claimed would improve password practices (section 8.4.3.1).
2. Participants' reaction to punishment threats (section 8.4.3.2).

The conclusion that can be drawn from the analysis is that users' accounts of the effect of punishment regimes should not necessarily be taken at face value. It seems that the best way of testing such regimes is in an experimental fashion.

8.4.3.1 Users' proposals to improve password practices

In the focus groups conducted as part of study 6, participants were asked what would need to be done to improve password practices. Their proposals included:

1. lectures on security;
2. workshops on security with a limited number of participants;
3. regular reminders of the importance of good password practice;
4. clear password regulations which are placed visibly in cluster rooms;
5. information sheets containing the regulations to be signed by students regularly;
6. punishment.

The interesting point about the first five proposals is that they all suggest what could be called a 'soft' approach to changing students' behaviour. There is no mention of punishment, increased monitoring of students' behaviour, or any other invasive techniques. In other words, they have in common that they rely on students' co-operation to be successful. A university - taking these statements literally - might put some of these proposals into practice, assuming that they should improve password-related behaviours since the students have claimed so themselves. However, participants in the second field trial (study 5) had actually signed off a very clear description of the regulations during enrolment, and they did not behave better than participants in the first field trial (study 2). Similarly, the university at the time ran an anti-plagiarism campaign, which made students sign a form whenever they handed in coursework saying that this was their own work. Still, this did not necessarily rule out plagiarism:

"I am sure if [lecturer's name] or whoever is in charge of plagiarism, if they saw us all doing that, say if they came to the labs on a Thursday night they would in theory expel half a dozen."

It is unlikely that just following the first five proposals would seriously improve password practices. From a discourse-analytic point of view, it is possible that they were made in order to avoid the fifth proposal, punishment. When this was brought up by participants, it was often accompanied by a demand for weak punishment:

"I think they should send out a global email to all the students reminding them [...]"

At the same time, some participants clearly stated that only harsh punishment to someone in their immediate environment would have an effect:

A *“I think the only way people are gonna really listen is through other people being punished, I can’t see, generally, constant reminders isn’t gonna affect everyone.”*

D *“I think it doesn’t hit home until it hits right near you then.”*

B *“Yes unless you see some action taken, then they’ll realise these guys mean business”*

This will be dealt with further in the following section.

8.4.3.2 Punishment

The previous section has shown that some participants in study 6 clearly believed that only harsh punishment would have a serious effect on students’ behaviour. However, this immediately met with another problem. Nobody expected to be the first person to be punished in an environment where nearly everybody violated regulations:

“Okay, so, say your account got hacked and because they found out your password was weak or not changed often enough, it wouldn’t be fair to get punished then because it was just bad luck that you were caught.”

Similar accounts can be found in the data obtained from study 3. However, here the resistance to any form of punishment was even stronger. To begin with, it was often remarked that it was highly unlikely that employees’ would get severely punished for violating regulations if this did not incur a security breach:

B *“It would be to be embarrassing to them [...]”*

C *“Yeah, I mean, it looks good in the press, doesn’t it. A person gets sacked for using all lower-case characters or something.”*

In addition, participants were reluctant to comply with regulations that seemed at times non-sensical to them, and preferred to follow their own judgement:

“I don’t think there would be any punishment for most it. Well, I partly oppose the whole idea of having a central set of rules for all systems that require passwords, because passwords are there for so many different reasons, I think it has to be on a per system basis, so the idea of having a central set of rules and punishments doesn’t seem to work.”

This is different from study 6, where participants might not have liked the idea of punishment, especially if it were severe, but admitted that it would change their behaviour. Here, participants wanted to be convinced that the regulations were sensible and well thought-out before they followed them:

“I think it’s perfectly reasonable to be given rules from the company that maybe even they say you get sacked if you don’t obey the rules because we think it’s important you do this, but if, but if they were then following up with something that was clearly badly thought out, and

wasn't actually gonna do anyone any favours, I thought that I understand it better than they did, and they clearly got it wrong, then I'd ignore it."

The question is what these participants would do if serious punishment regimes were actually employed by their organization, and they could see people being punished in their environment. Their discourse is at times structured to create the impression that this would not influence their behaviour, but this might just be done to decrease the likelihood of such measures being introduced.

8.5 Chapter summary

In this chapter, discourse analysis has been used to analyse the corpus of data which also formed the basis of the grounded theory model presented in the previous two chapters. This analysis served two purposes:

1. To capture the interpretative repertoires that participants drew on to describe those aspects of reality that influence their choice of password-related behaviours within the framework of the grounded theory model.
2. To determine whether it is possible that the accounts of password security issues which the grounded theory model is based on might have been constructed using specific interpretative repertoires and discursive techniques in order to justify malpractice.

The analysis has focussed on three points:

1. It has been shown that users employ a limited set of impoverished interpretative repertoires to describe potential **attackers** of computing resources. The exclusive use of these repertoires directly leads to insecure password practices. This information can be used in security campaigns to introduce repertoires that improve user behaviour.
2. The concepts of *resource repertoires*, the use of which improves user behaviour, and *problem repertoires*, which lead to insecure password practices, have been introduced. Examples of these have been given. Security campaigns can use this information to introduce and reinforce *resource repertoires*, while inoculating users against *problem repertoires*.
3. The accounts given by the participants in the studies have been investigated with respect to their action-orientation, and indications of them at times being constructed in order to justify malpractice and avoid punishment have been identified.

9 PERSUASIVE PASSWORD SECURITY

9.1 Overview

Organizations that rely on password mechanisms to secure their computing resources also rely on their users to behave in a manner that maximises the overall security of the systems in question. Often, it is possible to ensure that users employ secure behaviours by making them the only behaviours available to them, which is what measures such as proactive password checking achieve. However, in many situations, users will have a choice of a number of available behaviours. Some of these will increase the overall security, whereas others will reduce it, at times compromising it completely. Organizations who rely on password mechanisms for their security need to find ways in which they can influence users' choices so that they result in secure password practices. In this chapter, it will be shown how the research results presented in the previous four chapters can be used to guide this endeavour.

The core model introduced in chapter 6 describes the strategy that users employ to choose their password-related behaviours in the absence of any organizational efforts to influence this choice. When carrying out this strategy, users take into consideration a number of factors, such as the primary cost of the various available behaviours, any evidence of previous security breaches they are aware of (incident evidence), or their perception of the potential attackers of computing resources. The organization has a varying degree of control over these factors, and can configure them to influence the outcome of the choice process, e.g.:

1. The primary cost associated with performing the various behaviours available can be manipulated by changing the password mechanism.
2. Evidence of any security incidents in the past can be kept secret by the organization or made known to users.
3. A user's perception of the potential attackers of computing resources can be altered by security awareness and education campaigns.

There is no single factor that organizations can change to ensure that users behave properly. Instead, those factors that play a part in the choice process and which the organization has at least some degree of control over should be configured in an integrated fashion to achieve the following five goals:

1. The user cost of secure behaviours should be decreased, while that of insecure ones should be increased.

2. The **user benefit** of secure behaviours should be increased, while that of insecure ones should be decreased.
3. Users should assign the correct **provided security level** to the password-related behaviours available to them.
4. Users should assign the correct **required security level** to the resources they are accessing.
5. Users should be encouraged to secure the organizational resources they use in a responsible manner (**responsibility personality**). In particular, they should not take unnecessary risks with them (**risk personality**).

Section 9.2 will detail how organizations can use the core model to identify what factors they have to change, and how they have to change them, to achieve these five goals.

The extended model introduced in chapter 7 incorporates the effect of regulations and their associated punishment regimes into the core model. This is an example of an organizational measure which alters the strategy users employ to choose their password-related behaviours by introducing additional factors which have to be taken into consideration by them. Organizations who employ this particular measure need to manipulate a number of factors to achieve the following two goals:

1. Users should perceive the **punishment risk** associated with behaviours that violate regulations to be high.
2. Users who possess an **authority personality** that leads to them violating regulations should be discouraged to do so.

Section 9.3 will detail how the extended model can help organizations to identify the factors they have to change, and how they have to change them, to achieve these goals. It will also briefly discuss the possibility of other measures which change users' **choice** strategy by introducing additional factors that need to be taken into consideration by them.

Section 7.4 has pointed out that users employ the same strategy that they use to choose their password-related behaviours to determine their response to any efforts by their organization to influence this very strategy. This can make it difficult to get their attention when security-related information is being conveyed to them. Section 9.4 will discuss this issue.

The discourse-analytic study in chapter 8 has indicated the importance of the interpretative repertoires users employ for their choice of password-related behaviours. Section 9.5 will show how organizations can construct the discourse about password security in their communications with users so that *resource*

repertoires are introduced and reinforced, while the use of *problem repertoires* is discouraged.

Examples of how individual factors can be changed to alter user behaviour will be given throughout this chapter. However, changing just one of the factors that influence users' choices will not result in secure password practices being employed throughout the organization. Section 9.6 will summarise the recommendations made throughout this chapter and propose an integrated approach to improving user behaviour. This approach is based on the knowledge about individual factors, their relative importance, and the way they interact that the research presented in this thesis has provided. It aims to manipulate these factors in order to persuade users to choose behaviours that are secure over those that are not, resulting in *persuasive password security*.

9.2 Recommendations based on the core model

9.2.1 Required security level

Users assign required security levels to the resources they access with their passwords, based on their estimate of both the **attack likelihood** and the **attack consequences**. However, it is the former that most commonly leads to their assigning low and often inadequate required security levels to organizational resources. They just do not believe that they themselves, the groups they belong to or the resources they access will get targeted by attackers. Often, they realize that a successful attack could have dire consequences, but they deem this so unlikely to happen that they see no need to expand any serious effort on protecting the resources. This means that any organization wishing to improve user behaviour first of all needs to increase users' estimates of the **attack likelihood**. Once this has been achieved, users' estimates of the **attack consequences** should be amplified as well.

In order for users to increase their estimate of the **attack likelihood**, they need to believe that they themselves, the groups they belong to, and the resources they access are likely targets for attackers. There are three measures that organizations can employ to achieve this:

1. Users have to change their perceptions of potential **attackers** to include those that carry out regular attacks that target people with their **personal, group and resource characteristics**. As has been pointed out in chapter 8, users' perceptions of **attackers** are mediated through the interpretative repertoires they

use to describe them. At present, users do not hold repertoires that include attackers who would target them. Security awareness and education campaigns, as well as the general discourse about password security, need to introduce such repertoires and reinforce them constantly.

EXAMPLE 1	
PROBLEM	Users do not believe they will be targeted because they do not think their personal, group or resource characteristics warrant an attack.
ORGANIZATIONAL MEASURE	Security awareness and education campaigns, as well as the general discourse about password security, consistently mention attackers who target <i>any</i> ill-protected resource in order to cause damage.
RESULT	Users perceive themselves as potential targets, and increase their estimate of the attack likelihood.

2. Currently, organizations provide users with many different passwords to access different types of systems. Users believe some of these systems, such as those containing personal financial information or commercially sensitive company data, to have a higher attack likelihood than others. Organizations can bundle different types of systems together and protect them with a single password. In this case, the attack likelihood of the combined systems would be at least as high as the highest attack likelihood of all the elements. Similarly, users could be made parts of groups that they believe to be potential targets, or be provided with personal characteristics that increase their estimate of the attack likelihood. Any such measures have to be considered carefully, since they might also have unwanted side effects on the overall security of the systems. As an example, a password that protects multiple systems also puts all of them into danger, once it gets compromised.

EXAMPLE 2	
PROBLEM	Users do not protect a resource which the organization deems under threat of attack, because they do not regard it as a likely target.
ORGANIZATIONAL MEASURE	The organization sets up the system so that users' company-related financial information, which is currently protected by a different password, is accessible directly from the system it wants to protect.
RESULT	Users rate their financial information as a likely target for attackers, and subsequently assign a high attack likelihood to the password that protects both this information and the system the organization wants to secure.

3. Users also take into account the perceived security of their environment and any incident evidence they are aware of to determine the attack likelihood. Organizations should obviously increase the perceived security by providing physical and software security and by undertaking measures to improve people security. In addition, it is vitally important that any security breaches that do occur are published.

EXAMPLE 3	
PROBLEM	Users do not protect their resources because they assign a low attack likelihood to them, since they are not aware of any security breaches ever occurring. At the same time, such breaches occur regularly, but are not published by the organization.
ORGANIZATIONAL MEASURE	The organization for the first time publishes regular updates on security breaches.
RESULT	Users increase their estimate of the attack likelihood .

Once users have a sufficiently high estimate of the **attack likelihood**, the organization can also undertake a number of measures to increase their estimate of the **attack consequences**:

1. Users' perceptions of potential **attackers**, which has a strong effect on their estimate of the **attack likelihood**, also affects their estimate of the **attack consequences**. Security awareness and education campaigns, as well as the general discourse about passwords should introduce repertoires that describe attackers who do not only target typical users, but also do serious damage once they have successfully breached security. In particular, these repertoires should include the various ways in which such attackers can cause harm both to the user, his immediate colleagues and the organization.
2. Users are often more concerned about putting other people at risk than they are about endangering themselves. Depending on their **responsibility personality**, they may also want to protect their organization. Obviously, introducing repertoires that stress how successful attackers can use their account as a starting point to attack their colleagues and the organization will already increase their estimate of the **attack consequences**. In addition, resources can be bundled to increase the estimate even more.

EXAMPLE 4	
PROBLEM	Users do not protect their resources, because they believe that they can live with the consequences to themselves that a successful attack would incur.
ORGANIZATIONAL MEASURE	The organization sets up the system so that a number of resources that are shared within workgroups are directly, and without an additional password, accessible from users' systems.
RESULT	Users increase their estimate of the attack consequences .

3. Organizations commonly provide their users with **safeguards**, such as automatic backup. While these reduce the possible consequences of a security breach, they also reduce users' estimate of the **attack consequences**, and consequently the **required security level**. Organizations need to consider carefully whether there might be situations in which removing such **safeguards** might actually increase overall security.

EXAMPLE 5	
PROBLEM	Users do not protect their resources, since an automatic backup is in place, which means that they are not concerned about the consequences of a successful security breach.
ORGANIZATIONAL MEASURE	The organization stops the automatic backup.
RESULT	Users increase their estimate of the attack consequences.

Some of the recommendations made in this section seem at first counter-intuitive and go against common practices currently employed by organizations. In effect, they propose to increase the potential consequences of a security breach in order to improve user behaviour. Organizations will have to consider carefully the benefits and pitfalls of these proposals for each individual case. They also will need to make any such measures part of an integrated approach which aims to maximise overall security, as will be discussed in section 9.6.

9.2.2 Provided security level

Once users have assigned an adequately high required security level to the resources they access, they need to choose behaviours that provide this level of security. In order to do this, they have to be able to assign the correct provided security level to the password-related behaviours available to them. Organizations can undertake a number of measures to ensure that this is the case:

1. Users need to be informed about the provided security levels of the behaviours available to them. Currently, they often assign them based on the effort they believe attackers would need to make to circumvent them. Security awareness and education campaigns, as well as the general discourse about password security, need to introduce repertoires describing attackers that correctly reflect the ways in which password-related behaviours are circumvented. This way, they can then point out the provided security level of the relevant behaviours with an explanation of why they are deemed secure or insecure.

EXAMPLE 6	
PROBLEM	Users commonly share their password with students who are on a short-term visit to the organization, since they do not believe attackers to be able to exploit this behaviour.
ORGANIZATIONAL MEASURE	Security awareness and education campaigns, as well as the general discourse about password security, consistently mention attackers who disguise themselves as visiting students in order to obtain passwords.
RESULT	Users assign a low provided security level to the behaviour of sharing passwords with visiting students.

- Users commonly assign provided security levels based on what they see happening in their social context. Obviously, all the measures aimed at improving user behaviour will also change what users observe in their environment. However, it is likely that there will always be users who exhibit weak password practices. This will be particularly true in the transitional period following the initial deployment of the measures proposed in this chapter. Organizations should stress to users that they should follow organizational guidelines, rather than the people around them. Social marketing campaigns are one way of doing this.

EXAMPLE 7	
PROBLEM	Users, in particular those with little computing knowledge, commonly assign provided security levels to behaviours based on what they observe happening in their environment.
ORGANIZATIONAL MEASURE	The organization runs a social marketing campaign which specifically targets users with little computing knowledge, and encourages them to follow organizational guidelines.
RESULT	Users start following organizational guidelines.

- The perceived security of a user's environment can lead to incorrect estimates of provided security levels, whether it is high or low. Section 8.4 has pointed out that users might use the perceived security as an excuse for their misbehaviours, both to themselves and to their organization. The best strategy that organizations can employ is to have high levels of physical and software security, since this obviously protects their resources. In addition, security awareness and education campaigns, as well as the general discourse about password security (see section 9.5), should stress that high levels of security do not mean that users can relax their own security efforts.

EXAMPLE 8	
PROBLEM	Users assign low provided security levels to all password-related behaviours, since the perceived security is low. As a result, they choose insecure behaviours that incur low user costs and/or provide high user benefits.
ORGANIZATIONAL MEASURE	The organization improves physical and software security. In addition, it consistently mentions in its security awareness and education campaigns, and the general discourse about password security, that this does not mean that users can reduce their own security efforts.
RESULT	Users start following organizational guidelines.

9.2.3 User cost

9.2.3.1 Primary cost

It is obvious that organizations must reduce the primary cost to users incurred by performing secure behaviours. Currently, it is often not possible to behave in a correct manner, simply because the mechanisms are so ill-designed that doing so would come at an exorbitant cost. As a result, users have developed a number of coping strategies, some of which have been described in section 5.2. Organizations should aim to reduce the primary cost of secure behaviours, but this might not always be possible. In situations where users realistically need coping strategies to manage their passwords, these should be supported by the organization, so that overall security remains as high as possible. The primary cost of secure behaviours can be reduced by improving the usability of the password mechanism, and by increasing users' abilities. The most pressing usability problems that need to be solved are:

1. Users have to deal with too many resources that need separate passwords. Organizations should reduce this number by introducing single-sign on solutions for all or at least a larger number of these. In particular, systems that are used very infrequently should be bundled with systems that are used frequently. Resources that do not need protection should not be password-protected, or users should at least be encouraged to treat their passwords with less care than the ones used for sensitive resources. Finally, if proactive password checking is in place, the criteria for what constitutes acceptable password content should be synchronised across all systems.
2. The change regimes should be altered so that users have to change their passwords less often. They should also be synchronised, so that all passwords need to be changed at the same time.

3. The number of strikes that a user is allowed to incur before his account is blocked needs to be considered carefully to achieve a balance between security and usability (see Brostoff & Sasse (2003) for a recent update on this subject).
4. In cases where the usability of the password mechanism cannot be improved to an acceptable and manageable level, the coping strategies that users employ should be supported by the organization. The aim is to reduce the workload on the users, while ensuring that overall security remains as high as possible.

EXAMPLE 9	
PROBLEM	Users commonly write down their passwords, since too many systems need separate passwords for them to remember them by heart.
ORGANIZATIONAL MEASURE	The organization provides each small workgroup with a secure safe in which they can store their lists of passwords when they do not need them.
RESULT	Users still write down their passwords, but at least they are stored securely.

In addition, users have to be trained to cope with the tasks the password mechanism demands of them in an adequate fashion. In particular, this training should teach them the following:

1. The ability to choose strong yet memorable passwords.
2. The ability to manage large numbers of passwords which have to be changed regularly in a secure fashion. Again, where the number of passwords cannot be reduced, user should be taught coping strategies that are secure. Otherwise, they will develop their own coping strategies, which might severely undermine security.

It is difficult to think of ways in which the primary cost of insecure behaviours can be increased to make them less attractive to users. This is an area for future research.

9.2.3.2 Secondary cost

A secondary cost is incurred by users as a consequence of performing a specific password-related behaviour. Organizations need to identify the most common situations in which this problem occurs. A lot of these will involve users sharing their password, which actually incurs a higher (although negligible) primary cost than not sharing it. However, the secondary cost of not sharing can be too high. Section 5.2.3 has listed the situational contexts in which participants reported having shared their passwords. Such a list makes it possible to improve the password mechanism and its surrounding support system so that additional behaviours are available to users, which incur a low secondary cost, while maintaining overall security. Two examples make this point clear.

EXAMPLE 10	
PROBLEM	Users regularly phone a colleague in the office and give them their password, so that he can read their email for them while they are away.
ORGANIZATIONAL MEASURE	The organization sets up a helpline which users can call to have their email read to them.
RESULT	Passwords are no longer shared with colleagues in these situations.

EXAMPLE 11	
PROBLEM	Users regularly share their passwords with colleagues so that these can access specific resources.
ORGANIZATIONAL MEASURE	The organization makes it possible for users to set up guest accounts for their resources.
RESULT	Users do not share their own passwords any more.

It is difficult to think of ways in which an organization can systematically increase the secondary cost of insecure behaviours. This is a matter for future research.

9.2.3.3 Social cost and image cost

The social cost and the image cost associated with both secure and insecure behaviours are strongly influenced by the interpretative repertoires users hold (see section 8.3) and by the social norms of the environment they find themselves in. Organizations should use repertoires that lead to a high social cost and image cost for insecure behaviours, and no costs for secure ones, in their security awareness and education campaigns and the general discourse about password security. These repertoires can also be introduced and re-enforced by social marketing campaigns. However, none of these measures are likely to give overnight results, and they will need to be deployed consistently over extended periods of time. They are also unlikely to have a strong effect on user behaviour on their own, and should be employed in conjunction with other measures recommended in this chapter.

EXAMPLE 12	
PROBLEM	Users customarily write down their passwords.
ORGANIZATIONAL MEASURE	Security awareness and education campaigns, and the general discourse about password security, consistently stress that someone who writes down his password puts not only himself, but also his colleagues at risk.
RESULT	Users incur a social cost by writing down their password, since they are seen to endanger their colleagues.

9.2.4 User benefit

9.2.4.1 Creating primary benefits

The model currently does not contain any primary benefits obtained from carrying out secure behaviours, simply because they were not found in the data obtained. However, organizations can try to create such benefits artificially. One future area of research that seems promising is the use of pleasure-based approaches (e.g. Jordan (2000)), which might make it possible to reward users while performing a secure behaviour, or immediately afterwards.

9.2.4.2 Social benefit and image benefit

Organizations should change the social and image benefit associated with password-related behaviours in the same manner in which they change the social and image cost.

EXAMPLE 13	
PROBLEM	Users customarily share their passwords with colleagues, since this is a sign of trusting them.
ORGANIZATIONAL MEASURE	The organization runs a social marketing campaign which stresses that anybody who shares their password with a colleague puts his other colleagues and the organization at risk.
RESULT	Users obtain a social benefit from sharing their password, but also a social cost.

9.2.5 Personality

Organizations need not necessarily change their users' risk personality and responsibility personality in order to suppress any negative impact these might have on their choice of behaviour. Section 6.2.1 has already pointed out that, since these aspects of a users' personality influence his behaviour, they can also be deduced from it. Organizations can associate a negative image with overt risk-taking, or neglect of company concerns, by using appropriate interpretative repertoires in their security awareness and education campaigns, and the general discourse about password security. In addition, they can use social marketing to introduce and re-enforce these repertoires.

EXAMPLE 14	
PROBLEM	Users do not protect organizational resources, since they are not concerned about them.
ORGANIZATIONAL MEASURE	The organization runs a social marketing campaign which stresses that anybody who puts organizational resources at risk is 'unprofessional' and also endangers his colleagues.
RESULT	Users protect organizational resources, since they do not want to incur a public image of being 'unprofessional' and 'uncaring about their colleagues'.

9.2.6 Choosing out of habit

Users usually do not go through the whole strategy described by the core model to choose their password-related behaviours. Instead, they rely for some, or even all of the steps on estimates which they have arrived at in the past. This can lead to insecure choices when the situation they find themselves in is different from the one their original estimates were based on. To avoid this from happening, organizations should undertake three measures:

1. New members of an organization should be informed about the way password security is implemented. They should also be made aware of the password practices that are expected of them, and of the reasoning behind them. This is particularly important if they come from an environment in which password security was not taken seriously. Ideally, all this information should be conveyed on a personal basis, rather than in the form of a paper handout (see section 9.4). This should happen before they first use their computers in the new environment. This way, they are more likely to start with a new set of behaviours, which can then become secure habits.
2. Major changes to any aspect of password security should be accompanied by an awareness campaign, which should run along similar lines as the introduction to password security given to new members of the organization.
3. Minor changes should be accompanied by a warning to users that they should reassess their password practices, since they might want to choose different behaviours in the new situation.

At a societal, rather than organizational level, it would also be desirable for those institutions who first introduce users to password security (e.g. schools and universities) to take special care that good habits are developed from the very beginning.

9.3 Recommendations based on the extended model

9.3.1 Regulations and their associated punishment regimes

9.3.1.1 Punishment risk

Organizations wishing to increase users' estimates of the punishment risk associated with those behaviours that violate password regulations have to convince them that it is likely that malpractice will be both detected (detection likelihood) and punished (severity of expected punishment). In order to increase users' estimates of the detection likelihood, the following measures can be undertaken:

1. The organization has to be seen to make serious efforts to identify users who violate regulations. **Visible detection efforts** can be used to target specific violations, and will also make users more cautious about breaking other regulations. However, it is important to ensure that such measures fit in with the overall organizational culture.

EXAMPLE 15	
PROBLEM	Users customarily write down their passwords on a piece of paper, which they put into their drawer.
ORGANIZATIONAL MEASURE	The organization starts sending security officers around the building at irregular intervals. They check people's drawers for passwords that are written down.
RESULT	Users stop this particular malpractice, and are also more cautious about violating other regulations.

2. Organizations have to be seen to take security seriously (**importance of security**). This is an additional reason to improve both the **physical** and the **software security**, and to run security awareness and education campaigns on a regular basis.
3. Users need to be made aware that most improper behaviours which they believe are difficult to monitor (**ability to monitor**) can ultimately be found out, if the organization is willing to make a serious effort.

EXAMPLE 16	
PROBLEM	Users customarily share their passwords with close colleagues, since they trust them never to make this known to the organization.
ORGANIZATIONAL MEASURE	The organization stresses in its security awareness and education campaigns, and the general discourse about password security, that in a case of a security breach, close colleagues of the victim will be questioned about his behaviour in the past.
RESULT	Users do not share their passwords any more.

4. Users know their malpractice might be found out if a security breach occurs. However, they usually perceive the **attack likelihood** to be low, and are not worried about this. This means that measures that increase users' estimates of the **attack likelihood** will also raise their estimate of the **detection likelihood**.

5. Organizations have to be seen to punish offenders. This will be dealt with in the remainder of this section, which discusses the severity of expected punishment.

In order to increase users' estimate of the severity of expected punishment, organizations can undertake the following measures:

1. First and foremost, organizations must be seen to punish offenders. This should not be done behind closed doors, but as openly as is possible. Obviously, the organizational culture has to be taken into consideration, but as long as users cannot observe punishment when rules are breached, they assume the risk of punishment to be non-existent or negligibly small. The extent of the punishment is a sensitive issue. It should act as a deterrent, but not be considered unfair by users. It is for this reason that users need to be convinced about the importance of security before punishment regimes are implemented. In addition, the organization has to be seen to play its part in protecting computing resources before it starts punishing users for their malpractice. The situation will become particularly problematic for organizations which have neglected password security for a long time, and in which most users customarily break the password regulations. This will be discussed further in section 9.6.
2. The regulations have to be clear about the way users are expected to behave, and the punishments they will incur for violating password regulations. In addition, they should be consistent between systems and over time. All this not only ensures that users are aware of what they are supposed to do, but also clarifies their estimate of the severity of expected punishment.

9.3.1.2 Authority personality

Section 7.2.2.2 has shown that a user's authority personality can lead to his violating regulations. Organizations can undertake the following measures to stop this from happening:

1. The obvious way of dealing with users who rebel against the regulations out of principle, or who ignore them because they do not seem to make sense, is to increase the punishment risk to the point where they comply. However, this could also lead to severe dissatisfaction on their behalf, which is something a lot of organizations do not want to happen. For this reason, the following two recommendations should be implemented in conjunction with an increase of the punishment risk.
2. The measures discussed so far in this chapter should increase users' acceptance of the need for password security. They should also convince them that the organization is playing its part in securing its computing resources. As a result, a lot of users who have hitherto violated regulations because of their authority personality are likely to stop doing so and will accept that offenders need to be punished.
3. Users should be actively involved in the creation of secure password practices. They should be encouraged to inform the organization about the problems they encounter with the mechanism, instead of ignoring the regulations or even rebelling against them. This point will be discussed further in section 9.6.

9.3.2 Other possible extensions

Regulations and associated punishment regimes are an example of organizational measures which aim to influence the strategy users employ to choose their password-related behaviours by introducing new factors which they have to consider. An interesting area of future research would be to identify other factors that can be introduced to improve user behaviour, and which might be more successful than punishment-based approaches. Section 9.2.4.1 has proposed that research should be conducted into identifying primary benefits users can obtain from choosing secure behaviours. In the same manner, factors that reward users over longer periods of time need to be identified. A simple example of one such scheme would be to reward users for periods of time in which they have not experienced security breaches and have not been found out to behave insecurely by visible detection efforts. Such rewards could even be made part of employees' annual job performance reviews, as Parker has suggested (Parker (1998)).

9.4 Getting users' attention

This chapter has proposed a number of measures that rely on an organization's ability to convey security-related information to its users. Security awareness and education campaigns, social marketing initiatives and any similar measures usually have no problem getting their messages to the intended user group. However, users have to take in this information in a conscious fashion for it to have any effect on their behaviour. Section 7.4 has pointed out that users will employ the same strategy that they use to choose their password-related behaviours to determine how they deal with security-related information presented to them by their organization. This can lead to a situation where users who are not concerned about security also do not take in any security-related information. The following is a list of recommendations that organizations should take into account when wishing to convey information to their users:

1. Information should be conveyed person-to-person(s), whenever possible. This makes it possible for the messenger to react to the feedback that he gets from his audience and to ensure that they take in the information. Information that is presented by email, paper or similar media will often not be attended to by users.
2. The information itself should be personalised as much as possible. In particular, it should be made clear why it is relevant to the users it is targeted at.
3. Ideally, the information should be presented by members of the close social environment users find themselves in. Line managers, in particular, should be involved in this process.

4. Users should be reminded regularly of the relevant information that needs to be conveyed to them. It is advisable that the format in which this is done is changed occasionally to avoid users getting bored and disinterested.
5. Once such campaigns begin to have an effect on users, they will start to increase their estimates of factors such as the **attack likelihood** or the **punishment risk**. This then makes it possible to convey the information in manners that are less expensive, since users who are concerned about security are also more likely to read emails or circulars about it.

9.5 Recommendations based on the discourse-analytic study

This chapter has proposed a number of measures that aim to change the interpretative repertoires that users employ with respect to a number of password security-related issues. The overall aim is to introduce and reinforce *resource repertoires*, and to discourage the use of *problem repertoires*. Chapter 8 has given examples of these, as used by participants in the studies. These examples have been chosen because they relate to repertoires that have a particularly strong effect on user behaviour and should therefore be targeted first. It is hoped that future research will widen this pool of repertoires to make it even easier for organizations to construct the discourse about password security in a fashion that improves user behaviour. However, organizations using such information need to keep in mind that it is possible that users might refuse to employ new repertoires, since the ones they are currently using can serve as an excuse for malpractice (see section 8.4.2). One way in which this issue could be addressed is by creating what would in effect be meta interpretative repertoires, which condemn the use of specific repertoires by users and their refusal to employ others. In general, the whole discourse about password security within an organization also should be consistent in its use of interpretative repertoires, and in particular in its condemnation of those that lead to weak password practices. The best way of ensuring that this is the case is to issue recommendations for the way in which the discourse about password security should be constructed from interpretative repertoires. These should be employed organization-wide by all people who are responsible for communicating with users about these issues.

9.6 Persuasive password security: Recommendations

This chapter has presented a number of measures that organizations can employ to manipulate the factors that influence users' **choice** of password-related behaviours. The factors themselves differ in a variety of ways:

1. Some of the factors are interconnected, which means that changing one will affect several others. Users' perceptions of potential attackers, for example, influence their estimates of the attack likelihood, the attack consequences, the provided security level and the detection likelihood.
2. Different factors influence the choice process to different extents. Exaggerated high primary costs, for example, make the required security level less relevant.
3. The organization has varying degrees of control over different factors. The password mechanism in particular is under complete control of the organization, and can be changed easily, e.g. to reduce primary costs. Users' perceptions of potential attackers, on the other hand, can only be altered by measures such as security awareness and education campaigns. These are likely to alter this factor, but are not guaranteed to do so.
4. The timeframe that is needed to change a factor also varies. Those factors that are under complete control of the organization, for example, can usually be changed faster than those that are not.
5. The measures that are needed to change a factor differ in the costs they incur to the organization.

Persuasive password security is in place when all the factors that influence users' choice of password-related behaviours, and which the organization has at least some degree of control over, have been configured by the organization so that users are persuaded to choose secure behaviours over insecure ones. In reality, organizations will achieve degrees of *persuasive password security*. They might choose to implement only some of the recommendations, for a variety of practical reasons, cost being the main limiting factor. It is also likely that organizations will introduce persuasive password security gradually, thus spreading the cost and limiting any possible disturbances caused by major changes. In the remainder of this section, a high-level strategy which contains the main measures that would need to be undertaken to achieve a sufficient degree of *persuasive password security* will be outlined. This strategy can be followed completely, in parts or gradually by organizations.

Organizations should start out by altering those factors they have complete control over, and which have the maximum impact on user behaviour. In particular, the password mechanism and its support system (such as helpdesks) should be changed to achieve the following goals:

1. The primary costs of secure behaviours should be reduced to an absolute minimum (see section 9.2.3.1).
2. The coping mechanisms users need to employ should be supported in a manner that incurs low user costs and achieves an acceptable level of security (see sections 9.2.3.1 and 9.2.3.2).

3. Resources should be bundled and provided with a single password in a manner that increases users' estimates of the **attack likelihood** and the **attack consequences** (see section 9.2.1).
4. **Safeguards** should be assessed as to their effect on user behaviour, and dropped, if necessary (see section 9.2.1).

Improving the **physical and software security** in itself increases overall security, and also has beneficial effects on user behaviour:

5. The **physical and software security** should be improved, and users should be notified at the same time that this does not mean that they can relax their own security efforts (see section 9.2.2).

Incident evidence has a strong effect on user behaviour:

6. Any security breaches that occur should be communicated to users (see sections 9.2.1 and 9.3.1.1).

Security awareness and education campaigns will improve user behaviour in the long run, if they manage to get their attention:

7. New members of the organization should undergo an induction lesson on the way password security is implemented before they start using their computers (see section 9.4).
8. Regular security awareness and education campaigns should be conducted with users, in particular if aspects of password security have been changed (see section 9.4).
9. Users should be notified of minor changes, and advised to reassess their password practices (see section 9.4).

The core factor that needs to be changed by these campaigns is users' perception of **potential attackers**:

10. Introducing interpretative repertoires that improve users' perceptions of **potential attackers** should be a main prerogative of all security awareness and education campaigns (see sections 9.2.1, 9.2.2 and 9.3.1.1).

The campaigns should also train users:

11. Security awareness and education campaigns should teach users the skills necessary to use password mechanisms adequately (see sections 9.2.3.1 and 9.2.3.2).

The discourse about password security in general should be structured so that it positively influences user behaviour:

12. Organizations should issue recommendations for the way in which the discourse about password security should be constructed from interpretative repertoires, so that *resource repertoires* get introduced and reinforced, while the use of *problem repertoires* is discouraged (see section 9.5). These should be employed organization-wide by all people who are responsible for communicating with users about these issues.
13. Regular security awareness and education campaigns, as well as social marketing efforts, should be undertaken to introduce and reinforce *resource repertoires* and to discourage the use of problem repertoires (see section 9.5).

Organizations wishing to implement regulations with associated punishment regimes should take the following recommendations into consideration:

14. When introducing punishment regimes, it may be advisable to have a transitional period, during which punishment will be less severe and an amnesty for past offences will be granted. During this period, efforts should be undertaken to educate users about the need for password security, and to convince them that offenders have to be punished (see sections 9.3.1.1. and 9.3.1.2).
15. Regulations should be clear and consistent across systems and over time (see section 9.3.1.1).
16. Offenders should be punished as openly as possible (see section 9.3.1.1).
17. Visible detection efforts need to be undertaken on a regular basis (see section 9.3.1.1).
18. Users need to be warned that even behaviours that are difficult to monitor might be found out, if the organization has reason to put in the necessary effort (see section 9.3.1.1).

And finally, a recommendation which should be followed throughout the process of implementing *persuasive password security*:

19. Users should be encouraged to contribute to the efforts that are made to improve password security.

9.7 Chapter summary

In this chapter, it has been shown how the research results presented in the previous four chapters can be used by organizations to improve users' choice of password-related behaviours. Measures that can be undertaken to manipulate those factors that influence the choice process and which organizations have at least some degree of

control over have been listed. An integrated approach which aims to achieve *persuasive password security* has been presented. This approach is based on the knowledge about individual factors, their relative importance, and the way they interact that the research presented in this thesis has provided. It aims to configure these factors in order to persuade users to choose behaviours that are secure over those that are not, and it can be implemented by organizations as a whole, in parts or in a gradual fashion.

10 CONCLUSIONS

10.1 Thesis contributions

10.1.1 Overview

An organization that gives its members access to its resources via a password mechanism needs to ensure that they perform certain secure behaviours if it wants those resources to be protected adequately. The research problem this thesis has addressed is the question of how the likelihood of users in an organization performing the expected password-related behaviours can be increased in a situation where some of these behaviours can neither be enforced nor monitored adequately.

The critical review of previous research that has tried to tackle this problem has shown that the majority of this research has assumed one factor to have a strong influence on user behaviour and has then tried to find ways in which this factor can be changed in order to improve security practices. The factors that have been investigated by previous research are the usability of specific security mechanisms, users' knowledge and skills and users' motivation. It is interesting to note that research that has been published since the completion of the work presented in this thesis has not altered this approach. The papers presented at the CHI 2003 Workshop on Human-Computer Interaction and Security Systems focussed either on specific usability issues (Brostoff & Sasse (2003), Cranor (2003), Just (2003), Paul et al. (2003), Smith (2003), Wilson & Zurko (2003), Yurcik et al. (2003)) or suggested novel methodologies that will help in the creation of usable security (Grinter & Smetters (2003), Yee (2003)). Two papers looked at factors other than usability, namely the mental models users employ to understand security (Dourish et al. (2003)) and a methodology that can be used to train users of security mechanisms (Whitten & Tygar (2003)). Only one paper has explicitly pointed out that a number of factors will have to be taken into consideration *simultaneously* to improve user behaviour (Sasse (2003)).

The approach of assuming one factor to have a strong influence on user behaviour, followed by an investigation of how this factor can be altered in order to improve security practices, has two immediate drawbacks. Firstly, it makes it less likely that any other factors that influence user behaviour will be identified. Secondly, it does not make it possible to determine the way in which the interaction of different factors

causes specific behaviours. The research documented in this thesis has taken a different approach. Inspired by the work presented in Adams & Sasse (1999), it set out to answer the following 3 questions:

1. What are the factors that affect the password behaviour of users in organizations?
2. How do these factors interact in order to cause specific behaviours?
3. How can knowledge about these factors and their interplay be used to improve user behaviour?

The research has answered the first two questions by developing a grounded theory model of the way in which users choose their password-related behaviours and of the factors that influence this choice process. One particular factor that has been identified – the interpretative repertoires users draw on to describe aspects of password security – has been investigated further by using discourse analysis on the qualitative data collected. The opportunistic use of data that had been collected prior to a re-conceptualisation of the research approach (see chapter 4) has made it possible to examine the extent to which users violate password regulations, and an examination of all the qualitative data that has been gathered allowed for a first insight into the specific insecure behaviours users choose in certain situations. A discourse-analytic investigation of the discursive practices users engage in when creating their accounts of password security issues has made it possible to find a partial answer to the question of whether they might structure their discourse in a manner that makes it possible for them to justify their malpractice. Finally, the third question above has been answered by creating an integrated approach to improving user behaviour which is based on the findings presented in this thesis. This approach, named *persuasive password security*, is applicable by the people who have been entrusted by their organisation with the task of ensuring the effectiveness of password security.

The following subsections will elaborate further on the specific substantive and methodological contributions the thesis has made.

10.1.2 Substantive contributions

10.1.2.1 User behaviour and attitudes

The investigation of user behaviour in Chapter 5 has shown that students at University College London regularly violated password regulations. This substantiates a claim which so far has mostly been supported by anecdotal evidence (see section 2.4.3). The collection of the specific insecure practices participants engaged in, and of the

situational contexts in which they did so, is the first of its kind. It makes it possible for organisations to target these specific behaviours in the measures they undertake to improve password practices among their users. The analysis of users' attitudes towards various aspects of password security supports the argument that users' insecure password practices may often be the result of their attitudes and also highlights some of the problematic attitudes that need to be changed.

10.1.2.2 The core model

The core model of how users choose their password-related behaviours, and of the factors which influence this choice, lies at the heart of the thesis research. It includes a number of factors which had hitherto been overlooked, such as the **social cost** and **social benefit** of password-related behaviours, or individual users' **risk personality** and **responsibility personality**. Knowledge of these factors and of their interplay provides a fuller understanding of the way in which user behaviour is determined in specific situations, and opens up new routes of influencing it.

The model was consciously created without resorting to either the findings that had been generated by the limited amount of previous research in this field or to theoretical constructs that could be found in related disciplines (see section 3.3.1). A subsequent comparison of the model to the three factors that previous research had focussed on (see section 2.4.2) shows that they are incorporated in the model. **Usability** is a factor in the model which contributes to the **primary cost** and the **secondary cost** of password-related behaviours. Users' motivation is not an explicit factor in the model, but can in fact be explained by it in great detail by referring to factors such as the **required security level**, or the **user benefit** of password-related behaviours. User's knowledge and skills are specified by the model in the form of **user's abilities**, and factors such as a user's perception of **attackers** or the **provided security level** of password-related behaviours. As can be seen, the model provides detailed information about the previously under-specified factors of users' motivation, knowledge and skills, which makes it possible to design more specific measures to change them.

Comparing the model to the findings presented in Adams & Sasse (1999) (see section 2.4.4.2.2), it can be seen that the factors proposed there also appear in this model. "Multiple passwords", "change regimes" and "password content" appear as part of the **usability** of the password mechanism. "Users' perception of organisational security"

appears as **perceived security**, and “users’ perception of information sensitivity” as the **resource characteristics**. The “perceived compatibility between password practices and work procedures” shows up in the **secondary cost** of password-related behaviours. This validates the findings by Adams & Sasse (1999).

10.1.2.3 Extensions to the model

Chapter 7 has extended the core model to incorporate the effect of regulations and their associated punishment regimes on users’ **choices**. This was possible by simply adding two elements to the model, without a need to change its fundamental structure. This raises confidence that other extensions, incorporating further measures aimed at improving user behaviour, can be added just as easily. The extended model can be used to enhance regulation-based approaches to improving user behaviour, but can also be employed as a guide when designing different measures that have to work in conjunction with such regulation-based approaches.

The extended model can explain why users often ignore security awareness measures undertaken by their organisation: they determine their response to such measures using the same strategy that has led to their (often insecure) password practices in the first place.

The effect of security awareness and education campaigns is not included into the model, but it was still possible to give first insights into how these can be improved.

10.1.2.4 Interpretative repertoires

Chapter 8 presents the interpretative repertoires that were found to have the strongest effect on user behaviour. This information makes it possible to encourage the use of repertoires that improve user behaviour and discourage the use of those that worsen it.

10.1.2.5 Discursive practices

Chapter 8 also puts forward an argument that participants in the studies might at times have structured their discourse in a manner that made it possible for them to justify their malpractice and to reduce the likelihood of punishment regimes being introduced. This sheds light on the importance of investigating qualitative data in this field carefully, and of testing any findings based on this data in an observational or experimental fashion.

10.1.2.6 Persuasive password security

Chapter 9 has shown how the findings in this thesis can be used by organisations to improve user behaviour. The integrated approach aimed at achieving this, named *persuasive password security*, makes the findings thus easily accessible and applicable to people who have been entrusted by their organisation with the task of ensuring the effectiveness of password security.

10.1.3 Methodological contributions

10.1.3.1 Extension of HCI scope

Section 2.2 has presented three extensions to the traditional scope of HCI which are necessary to address the research problem and the research questions that have driven the work documented in this thesis. It is suggested that these extensions should be considered for incorporation into the canon of what constitutes HCI research in order to equip the field with the ability to deal with problems of a similar nature.

10.1.3.2 Fear appeals

The use of a fear appeal to change user behaviour in this thesis is the only example of this approach being used in this field that could be found in the literature. However, the failed fear appeal in the first field trial (see section 4.3) has shown up a number of issues that will make it difficult to use fear appeals in this field. This thesis has suggested that organisations should publish security breaches and cases of users having been punished for breaking regulations. It has to be seen whether such changes will make the use of fear appeals feasible again.

10.1.3.3 Discourse analysis

Only one previous example of the use of Potter and Wetherell's brand of discourse analysis to approach an HCI problem (Rimmer et al. (1999)) could be found in the literature. The notion of interpretative repertoires as culturally shared linguistic resources that individuals draw on and which have an effect on their behaviour in sometimes unintended ways has shown itself to be useful in this research. It has made it possible to make recommendations for the way in which organisations should structure the discourse about password security in their communications with their users to improve user behaviour. Discourse analysis' ability to question discourse for its action-orientation has also made it possible to put forward an argument that participants in the studies at times structured their discourse in order to justify their

malpractice and to reduce the likelihood of punishment regimes being introduced. On the one hand, both of these issues would have been difficult, if not impossible, to address taking a pure grounded theory approach. On the other hand, discourse analysis was replaced by grounded theory as the primary methodology in this research because it could not be used to create a theory of the phenomena under investigation. The work presented here indicates that these two methods might complement each other well, but it also supports an argument for the need to test the findings arrived at by their combination in an experimental fashion (see section 10.2.3).

10.2 Critical review of the thesis

10.2.1 Research scope

The research documented in this thesis has provided a grounded theory model of the way in which users choose their password-related behaviours, and of the factors influencing this choice. One particular factor – the interpretative repertoires users draw on to describe aspects of password security – has been investigated further by using discourse analysis. The almost exclusive use of qualitative analysis methods was motivated by the fact that, prior to this research, no model of the phenomena under investigation was in existence, which means that quantitative research questions would have restricted and pre-defined the phenomena from the outset. A side-effect of this use of qualitative methods was that it limited access to participants in two ways. Firstly, the number of participants that could be taken on was considerably smaller than it typically would have been in research using quantitative methods, simply because the qualitative analysis of data is extremely time-consuming. Secondly, participants had to be willing to expend one hour on an interview or a focus group, which reduced the number of willing volunteers and required the use of incentives or contacts who were able to convince people to participate. This, in turn, led to a restriction on the organisational backgrounds participants could realistically be recruited from.

As a result, the findings in this thesis are restricted to the scope defined by the user population and organisation population (see section 1.4) and need to be expanded to other contexts in future research (see section 10.3.2). At the same time, the user and organisation population in this research are representative enough to make the findings immediately applicable to a large number of organisations.

10.2.2 Thesis complexity

The grounded theory model that has been presented in this thesis is complex both with respect to the number of factors that it incorporates and the interaction of these factors. However, this is simply the result of the model being concerned with a complex phenomenon and cannot be changed without losing valuable information. At the same time, the model does have a hierarchical structure, and the top-level storyline should most certainly be accessible to the envisaged user of the model. The problem is that these users will only be able to use the model properly if they also understand the more complex parts of it below the top-level storyline. In order to use the findings of the discourse-analytic study as well, they do not only have to grasp them in their full complexity, but also have to understand where they fit in with the grounded theory model. In effect, the complexity of the model, the discourse-analytic findings and the interplay of the two makes it very unlikely that the thesis findings will be accessible to a lot of the people in organisations who have been entrusted with the task of ensuring the effectiveness of password security. Instead, these parts of the thesis are more likely to be of value to researchers who want to validate and extend the findings. Chapter 9 is aimed strongly at actual beneficiaries of the findings presented here, who want to employ them to improve password practices in their organisation. It is hoped that the integration of the findings into the approach named *persuasive password security* will make it possible for potential users of the thesis findings to get an overview of them which will allow them to come to a first evaluation of its suitability for their purposes and can then be followed up by a further examination of other relevant parts of the thesis which chapter 9 points to.

10.2.3 Discourse analysis invalidating grounded theory findings?

The investigation of the discursive practices participants in the studies engaged in has resulted in the realisation that it is possible that some of them might at times have structured their discourse in order to justify their malpractice or to reduce the likelihood of punishment regimes being introduced. While this is in itself an important finding of this research, it also puts a question mark over the results gained by applying grounded theory to this data. However, it can be argued that the use of both grounded theory and discourse analysis on the same corpus of data has sensitised the author to the possible action-orientation of some of the discourse and has made it possible to identify elements of it that need to be treated with caution. At the same

time, the discovery of the possible action-orientation of some of the discourse does reduce the degree of confidence one can have in the findings of a grounded theory analysis of qualitative data in this field, where participants can have a vested interest in structuring their discourse to achieve certain goals. This does, however, not invalidate the use of grounded theory as a tool that makes it possible to create a theory of complex phenomena that are not yet sufficiently understood. Instead, it strongly encourages the testing of such theories in an experimental fashion (see section 10.3.1).

10.3 Further research

10.3.1 Testing the model's predictions

Grounded theory was used as the primary methodology in this research because it made it possible to create a theory of a complex phenomena about which very little was known at the time. The use of an experimental approach would have relied on the assumption of certain factors influencing user behaviour strongly and could have led to other, relevant factors being ignored. However, now that the model has been developed, it is possible to use it as the basis for the design of exactly such experimental studies. As the discussion in section 10.2.3 has shown, such experiments will further increase our confidence in the results of the grounded theory analysis. As importantly, they will make it possible to begin quantifying the relative importance of the factors proposed by the model.

10.3.2 Expanding the scope of the results

Section 10.2.1 has pointed out that the scope of the research findings is currently limited to similar user populations and organisation populations as the ones it is based on. An important area of future research would be an investigation of other populations in order to generalise the findings. In theory, such research could apply the same methodological approach as taken in this thesis to create models of other populations and then compare them to the model presented here. However, it will be more practical to test the findings presented here on other populations by conducting focus groups that specifically ask questions aimed at validating the existence of factors and their interplay, while leaving room for the participants to bring up their own issues. Alternatively, an experimental approach, as suggested in section 10.3.1,

would make it possible to validate the model for larger populations in a smaller amount of time.

10.3.3 Adding further extensions to the model

The grounded theory model presented in this thesis includes an extension that incorporates the effect of password regulations and their associated punishment regimes on user behaviour. Further research is needed to integrate the effect of other measures organisations are already undertaking to improve user behaviour. The first candidates for this would be security awareness and education campaigns and social marketing initiatives. However, the structure of the model also points to means of improving user behaviour which currently are not being employed by organisations. An example of this would be the use of pleasure-based approaches to provide users with a primary benefit while performing a secure behaviour (see section 9.2.4.1). Such measures would have to be developed first, and could then be integrated into the model.

REFERENCES

- Adams, A., & Sasse, M. A. (1999). Users are not the enemy: Why users compromise security mechanisms and how to take remedial measures. *Communications of the ACM*, 42(12), 40-46.
- Adams, A., & Sasse, M. A. (2001). *Privacy in multimedia communications: protecting users not just data*. Paper presented at the People and Computers XV - Interaction without frontiers. Joint Proceedings of HCI2001 and ICM2001, Lille.
- Adams, A., Sasse, M. A., & Lunt, P. (1997). *Making passwords secure and usable*. Paper presented at the HCI '97 - People and Computers XII, Bristol.
- Anderson, R. J. (1994). Why Cryptosystems Fail. *Communications of the ACM*, 37(11), 32-40.
- Anderson, R. J. (2001). *Security engineering: a guide to building dependable distributed systems*: John Wiley & Sons.
- Belgers, W. (1993). Unix password security. Retrieved February 11th 2004, from www.ja.net/CERT/Belgers/Unix-password-security.html
- Bouch, A., & Sasse, M. A. (1999). *Network quality of service - an integrated perspective*. Paper presented at the RTA's'99, Vancouver.
- Brostoff, S., & Sasse, M. A. (2000). *Are passfaces more usable than passwords? A field trial investigation*. Paper presented at the HCI 2000, Sunderland, UK.
- Brostoff, S., & Sasse, M. A. (2001). *Safe and Sound: a safety-critical design approach to security*. Paper presented at the New Security Paradigms Workshop 2001, Cloudcroft, New Mexico, USA.
- Brostoff, S., & Sasse, M. A. (2003). *"Ten strikes and you're out": Increasing the number of login attempts can improve password usability*. Paper presented at the CHI 2003 Workshop on Human-Computer Interaction and Security Systems, Ft. Lauderdale.
- Bunnell, J., Podd, J., Henderson, R., Napier, R., & Kennedy-Moffat, J. (1997). Cognitive, associative and conventional passwords: Recall and guessing rates. *Computers and Security*, 16(7), 629-641.
- Burman, E., & Parker, I. (Eds.). (1993). *Discourse analytic research. Repertoires and readings of texts in action*. London: Routledge.
- Carroll, J. M., & Campbell, R. L. (1989). Artifacts as Psychological Theories: the Case of Human-Computer Interaction. *Behaviour and Information Technology*, 8(4), 247-256.
- Carstens, D. S., McCauley-Bell, P., & Malone, L. C. (2000). *Development of a model for determining the impact of password authentication practices on information security*. Paper presented at the International Ergonomics Association 2000/Human Factors & Ergonomics Society 2000 Congress, San Diego, California.
- Cooligan, H. (1990). *Research methods and statistics in psychology*. London: Hodder & Stoughton.
- Cranor, L. F. (2003). *Designing a privacy preference specification interface: A case study*. Paper presented at the CHI 2003 Workshop on Human-Computer Interaction and Security Systems, Ft. Lauderdale, Florida.
- Davis, D., & Price, W. (1987). *Security for computer networks*. Chichester: John Wiley & Sons.

- Deane, F., Barrelle, K., Henderson, R., & Mahar, D. (1995). Perceived acceptability of biometric security systems. *Computers & Security, 14*(3), 225-231.
- Dhillon, G., & Backhouse, J. (2001). Current directions in IS security research: towards socio-organizational perspectives. *Information Systems Journal, 11*(2), 127-153.
- Dix, A. J., Finlay, J. E., Abowd, G. D., & Beale, R. (1998). *Human-Computer Interaction. Second Edition*. Hemel Hempstead: Prentice Hall Europe.
- Dourish, P., Delgado de la Flor, J., & Joseph, M. (2003). *Security as a practical problem: Some preliminary observations of everyday mental models*. Paper presented at the CHI 2003 Workshop on Human-Computer Interaction and Security Systems, Ft. Lauderdale.
- Dufft, C., Espey, J., Neuf, H., Rudinger, G., & Stapf, K. (1999). Usability and Security. In G. Mueller & K. Rannenber (Eds.), *Multilateral Security in Communications, Volume 3 - Technology, Infrastructure, Economy*: Addison Wesley.
- Foley, S., & Jacob, J. (1995). *Specifying security for CSCW systems*. Paper presented at the 8th IEEE Computer Security Foundations Workshop, Kenmare, Co. Kerry, Ireland.
- Garfinkel, S. (1995). *PGP: Pretty good privacy*: O'Reilly and Associates.
- Garfinkel, S., & Spafford, G. (1996). *Practical Unix and Internet Security*. Cambridge, UK: O'Reilly & Associates.
- Glaser, B., & Strauss, A. (1967). *The discovery of grounded theory*. New York: Aldine.
- Gollmann, D. (1999). *Computer Security*. New York: Wiley.
- Greening, T. (1996). Ask and ye shall receive: A study in social engineering. *SIGSAC Review, 14*(2), 9-14.
- Grinter, R. E., & Smetters, D. K. (2003). *Three challenges for embedding security into applications*. Paper presented at the CHI 2003 Workshop on Human-Computer Interaction and Security Systems, Ft. Lauderdale.
- Henwood, K. L., & Pidgeon, N. F. (1992). Qualitative research and psychological theorizing. *British Journal of Psychology, 83*(1), 97-111.
- Hitchings, J. (1995). Deficiencies of the traditional approach to information security and the requirements for a new methodology. *Computers & Security, 14*, 377-383.
- Holmstroem, U. (1999). *User-centred design of secure software*. Paper presented at the Human Factors in Telecommunications, Copenhagen, Denmark.
- Jendricke, U., & von Markotten, D. G. (2000). *Usability meets Security - The identity-manager as your personal security assistant for the internet*. Paper presented at the 16th Annual Computer Security Applications Conference, New Orleans, Louisiana.
- Jordan, P. W. (2000). *Designing pleasurable products. An introduction to the new human factors*: Taylor & Francis.
- Just, M. (2003). *Designing secure yet usable credential recovery systems with challenge questions*. Paper presented at the CHI 2003 Workshop on Human-Computer Interaction and Security Systems, Ft. Lauderdale.
- Kabay, M. E. (1993). *Social psychology and InfoSec: Psycho-social factors in the implementation of security policy*. Paper presented at the 16th National Computer Security Conference, Baltimore, MD.

- Karat, C.-M. (1989). *Iterative usability testing of a security application*. Paper presented at the Human Factors Society 33rd annual meeting.
- Kim, H.-J. (1995). Biometrics, is it a viable proposition for identity authentication and access control. *Computers & Security*, 14(3), 205-214.
- Kirby, M. (1991). *Custom Manual* (Technical Report DPO/STD/1.0). Huddersfield: HCI Research Centre, University of Huddersfield.
- Klein, H., & Myers, M. (1999). A set of principles for conducting and evaluating interpretive field studies in information systems. *MIS quarterly*, 23(1), 67-93.
- Leyden, J. (2003). Office workers give away passwords for a cheap pen. Retrieved 14th August, 2003, from <http://www.theregister.co.uk/content/55/30324.html>
- Loftus, E. (1975). Leading questions and the eyewitness report. *Cognitive Psychology*, 7, 560-572.
- Lunt, P., & Livingstone, S. (1996). Rethinking the focus group in media and communications research. *Journal of Communications*, 46(2), 79-98.
- McLean, K. (1992). *Information security awareness - selling the cause*. Paper presented at the IFIP TC11 (Sec'92).
- Mitnick, K. D., & Simon, W. L. (2002). *The art of deception: Controlling the Human Element of Security*: Wiley Publishing, Inc.
- Mosteller, W. S., & Ballas, J. (1989). *Usability analysis of messages from a security system*. Paper presented at the Human Factors Society 33rd Annual Meeting.
- National Institute of Standards and Technology. (1995). NIST Special Publication 800-12, *An Introduction to Computer Security: The NIST Handbook*. Available from <http://csrc.nlst.gov/publications>.
- Parker, D. B. (1998). *Fighting computer crime: a new framework for protecting information*: John Wiley & Sons, Inc.
- Paul, N., Evans, D., Rubin, A., & Wallach, D. (2003). *Authentication for remote voting*. Paper presented at the CHI 2003 Workshop on Human-Computer Interaction and Security Systems, Ft. Lauderdale.
- Perrine, T., & Kowatch, D. (2003). Terracrack: Password cracking using TeraFLOP and PetaByte Resources. Retrieved April 15th 2005, from <http://security.sdsc.edu/publications/terracrack.pdf>
- Petrie, H. (2002). Password clues. Retrieved February 14th 2004, from www.centralnic.com/page.php?cip=77
- Pinkas, B., & Sander, T. (2002). *Securing Passwords against Dictionary Attacks*. Paper presented at the 9th ACM Conference on Computer and Communications Security, Washington, DC.
- Potter, J., & Wetherell, M. (1987). *Discourse and social psychology. Beyond attitudes and behaviour*. London: Sage Publications Ltd.
- Potter, J., & Wetherell, M. (1995). Discourse Analysis. In J. Smith, R. Harre & R. v. Langenhove (Eds.), *Rethinking Methods in Psychology*. London: Sage.
- Poulsen, K. (2000). Mitnick to lawmakers: People, phones and weakest links. Retrieved 15/06, 2000, from <http://www.politechbot.com/p-00969.html>
- Preece, J., Rogers, Y., Sharp, H., Benyon, D., Holland, S., & Carey, T. (1994). *Human-Computer Interaction*. Wokingham: Addison-Wesley.
- Reason, J. (1990). *Human Error*. Cambridge, UK: Cambridge University Press.

- Rimmer, J., Wakeman, I., Sheeran, L., & Sasse, M. A. (1999). *Examining users' repertoire of Internet applications*. Paper presented at the Human-Computer Interaction INTERACT '99.
- Saltzer, J. H., & Schroeder, M. D. (1975). *The protection of information in computer systems*. Paper presented at the IEEE 63.
- Sasse, M. A. (2003). *Computer security: Anatomy of a usability disaster, and a plan for recovery*. Paper presented at the CHI 2003 Workshop on Human-Computer Interaction and Security Systems, Ft. Lauderdale.
- Sasse, M. A., Brostoff, S., & Weirich, D. (2001). Transforming the "weakest link": a human-computer interaction approach to usable and effective security. *BT Technical Journal*, 19(3), 122-131.
- Schneier, B. (2000). *Secrets and Lies*: John Wiley & Sons.
- Schultz, E. E., Proctor, R. W., Lien, M.-C., & Salvendy, G. (2001). Usability and Security. An Appraisal of usability issues in information security methods. *Computers & Security*, 20(7), 620-634.
- Shen, H., & Dewan, P. (1992). *Access control in collaborative systems*. Paper presented at the ACM Conference on Computer Supported Cooperative Work, New York.
- Siponen, M. (2000a). A conceptual foundation for organizational information security awareness. *Information Management & Computer Security*, 8(1), 31-41.
- Siponen, M. (2000b). Critical analysis of different approaches to minimizing user-related faults in information systems: implications for research and practice. *Information Management & Computer Security*, 8(5), 197-209.
- Siponen, M. (2001). Five Dimensions of information security awareness. *ACM Computers & Society*, 31(2), 24-29.
- Smith, S. W. (2003). *Position paper: Effective PKI requires effective HCI*. Paper presented at the CHI 2003 Workshop on Human-Computer Interaction and Security Systems, Ft. Lauderdale.
- Spector, Y., & Ginzberg, J. (1994). Pass sentence - a new approach to computer code. *Computers and Security*, 13(2), 145-160.
- Spruit, M. (1998). *Competing against human failing*. Paper presented at the IFIP 14th International Conference on Information Security (SEC '98), Vienna/Budapest.
- Spurling, P. (1995). Promoting security awareness and commitment. *Information Management & Computer Security*, 3(2), 20-26.
- Stallings, W. (1995). *Network and internet security*: Prentice-Hall.
- Strauss, A., Bucher, R., Ehrich, D., Schatsman, L., & Sabshin, M. (1964). *Psychiatric ideologies and institutions*. IL: Free Press.
- Strauss, A., & Corbin, J. (1990). *Basics of qualitative research: Grounded theory procedures and techniques*. London: Sage.
- Strauss, A., & Corbin, J. (1998). *Basics of qualitative research: Techniques and Procedures for developing grounded theory*. Thousand Oaks, CA: Sage Publications.
- Thomson, M. E., & Solms, R. v. (1997). *An effective information security awareness program for industry*. Paper presented at the WG 11.2 and WG 11.1 of the TC11 IFIP.
- Thomson, M. E., & Solms, R. v. (1998). Information security awareness: educating our users effectively. *Information management & computer security*, 6(4), 167-173.

- Viega, J., & McGraw, G. (2001). *Building secure software: How to avoid security problems the right way*: Addison Wesley.
- Wang, p., Kim, Y., Kher, V., & Kwon, T. (2005). *Strengthening password-based authentication protocols against online dictionary attacks*. Paper presented at the Applied Cryptography and Network Security - ACNS 2005.
- Weirich, D., & Sasse, M. A. (2001a). *Persuasive Password Security*. Paper presented at the CHI2001 Extended Abstracts, Seattle, USA.
- Weirich, D., & Sasse, M. A. (2001b). *Pretty good persuasion: A first step towards effective password security in the real world*. Paper presented at the New Security Paradigms Workshop 2001, Cloudcroft, New Mexico, USA.
- Wetherell, M., & Potter, J. (1988). Discourse analysis and the identification of interpretative repertoires. In C. Antaki (Ed.), *Analysing lay explanations: a case-book*. London: Sage.
- Whitten, A., & Tygar, J. D. (1998). *Usability of security: A case study* (Technical Report No. CMU-CS-98-155). Pittsburgh: Carnegie Mellon University School of Computer Science.
- Whitten, A., & Tygar, J. D. (1999). *Why Johnny can't encrypt: A usability evaluation of PGP 5.0*. Paper presented at the 9th USENIX Security Composium, Washington.
- Whitten, A., & Tygar, J. D. (2003). *Safe staging for computer security*. Paper presented at the CHI 2003 Workshop on Human-Computer Interaction and Security Systems, Ft. Lauderdale.
- Willig, C. (Ed.). (1999). *Applied discourse analysis: social and psychological interventions*. Buckingham, UK: Open University Press.
- Wilson, D., & Zurko, M. E. (2003). *Lotus Notes and Domino contribution to the HCI and security systems workshop*. Paper presented at the CHI 2003 Workshop on Human-Computer Interaction and Security Systems, Ft. Lauderdale.
- Winkler, I. (1997). *Corporate Espionage. What it is, why its happening in your company, what you must do about it*. Rocklin, CA: Prima Publishing.
- Witte, K., Meyer, G., & Martell, D. (2001). *Effective health risk messages: A step-by-step guide*. Thousand Oaks, California, USA: Sage Publications Inc.
- Yan, J., Blackwell, A., Anderson, R., & Grant, A. (2000). The memorability and security of passwords - some empirical results (Technical report no. 500). Cambridge: Computer Laboratory, University of Cambridge.
- Yan, J. J. (2001). *A note on proactive password checking*. Paper presented at the New Security Paradigms Workshop 2001, Cloudcroft, New Mexico, USA.
- Yee, K.-P. (2003). *Secure interaction design and principle of least authority*. Paper presented at the CHI 2003 Workshop on Human-Computer Interaction and Security Systems, Ft. Lauderdale.
- Yurcik, W., Lakkaraju, k., & Haberman, M. (2003). *Two visual computer network security monitoring tools incorporating operator interface requirements*. Paper presented at the CHI 2003 Workshop on Human-Computer Interaction and Security Systems, Ft. Lauderdale.
- Zurko, M. E., Simon, R., & Sanfilippo, T. (1999). *A user-centered, modular authorization service built on an RBAC foundation*. Paper presented at the IEEE Computer Society Symposium on Research in Security and Privacy, Oakland, CA.
- Zurko, M. E., & Simon, R. T. (1996). *User-centered security*. Paper presented at the New Security Paradigms Workshop 1996, Lake Arrowhead, CA.

Zviran, M., & Haga, W. J. (1993). A comparison of password techniques for multilevel authentication mechanisms. *The Computer Journal*, 36(3), 227-237.

APPENDIX A: GROUNDED THEORY CATEGORIES AND PROPERTIES IN ALPHABETICAL ORDER WITH EXAMPLE QUOTES

1. Ability to satisfy proactive password checking criteria

A user who is faced with a proactive password checking mechanism has to be able to create passwords that satisfy the criteria of that mechanism, whether he wants his password to be strong or not. This category is a specific instance of the abstract category *user's abilities*.

Example 1: "Yes, except that most passwords require me to have digits as well, it would be more than [...] That's right, I just bang something in there. I have a few rules that I use as well. [...] Things like, ahm, well, things like I might replace an 'o' in a letter with a zero, or an 'I' with a one, that kind of thing."

Example 2: "I tend to use sort of stock passwords, to see if they work."

Example 3: "...so I do try and restrict myself to, to a pool of about, I suppose... I, I tend to have about three... three words and I tend to have about... basically, I suppose, three numbers. Ahm. That I use as often as I can, wherever I can. I mean, obviously sometimes it, you're not allowed to, because the systems don't allow that sort of thing. But, ahm, and then after that, ahm, if it's a number word combination, I try and sort of use eight, one of my three words with one-two-three-four sort of, that sort of combination."

2. Ability to connect passwords

A user who is subjugated to resources that employ change regimes and who does not want to choose perfectly unconnected passwords has to find a way of connecting his passwords in some way, e.g. by basing them on a theme or by indexing them. This category is a specific instance of the abstract category *user's abilities*.

Example 1: "Aahh, I use a sort of a stem, and I add different things to the end of it."

Example 2: "So I've got a basic password and I actually change bits of it."

Example 3: "I change the password as little as I can, often just the first, first digit. Ah, on the systems that require a change."

3. Ability to create password content

A user wishing to use a password of a certain strength and memorability must be able to create adequate password content. This category is a specific instance of the abstract category *user's abilities*.

Example 1: "They're words, or, ahm, sometimes it's where I've just been on holiday, or just something that comes, could be a TV programme I watch, you know, it could be anything. [...] Oh, well, I might put, ahm, a number in, so I say, if it was, for instance, Eastenders, I might put a 3, and then the rest of the word, so I do sometimes just put a number in."

Example 2: "Well, the NT one, it's my, what I actually got on my, I've got a mouse mat, and it happens to be a phrase connected to the current mouse mat. [...] it's something that, by looking at the mouse mat, it will remind me, sort of, at the moment it's a Dilbert mouse mat, so there is something."

Example 3: "I pick a word, and then abuse it in some way, maybe changing, ah, letters to numbers, to numerics, which look the same."

4. Ability to manage multiple passwords

A user wishing to use different passwords for at least some of his resources and wanting to change them regularly has to be able to manage these passwords correspondingly. This category is a specific instance of the abstract category *user's abilities*.

Example 1: "I have sort of three main passwords, and I use them in different contexts."

Example 2: "...most of the passwords I use will be part of a, a themed set."

Example 3: "I've got a standard phrase that I use, which is modified by the particular application. About half of them, and for others I pick a random word."

5. Ability to monitor

Users are fully aware that a lot of the behaviours set out in password regulations are difficult, if not impossible, to monitor. Various behaviours can actually be ranked according to the effort necessary to detect them. Ability to monitor is a property of the category available behaviour.

Example 1: "If they looked at the password and found out that it was a dictionary word or something."

Example 2: "Someone could come forward and say they heard you telling your password to someone else, or something like that."

Example 3: "I still don't see how anybody could ever tell whether you were doing it or not."

6. Attackers

This category refers to the specific perceptions of potential attackers of their computing resources that users have.

Example 1: "...spotty little s***s, basically, that have nothing better to do than to fascinate themselves by writing programs about how to get into things which they're not supposed to."

Example 2: "Very technically literate, very capable technically, with a devious mind."

Example 3: "Well, obviously employees, who may well have a, a grudge against a company."

7. Attack consequences

This category refers to the possible negative consequences users envisage as the result of a successful attack on one of their computing resources. It is an abstract category that groups together the categories consequences to user, consequences to others and consequences to organization. Examples can be found in the subsections dealing with these categories.

8. Attack likelihood

This category refers to a user's estimate of how likely it is that someone will attack one of his computing resources.

Example 1: "Pretty likely."

Example 2: "I actually think it is very unlikely that anybody has ever or would ever want to try and get into any of the systems that I log on to..."

Example 3: "Not very likely really."

9. Attack types

Users have specific perceptions of the types of attacks they may be exposed to. Attack types is an abstract category that groups together the categories personal attack, reward attack, group attack, cruising attack and opportunistic attack. Examples can be found in the subsections dealing with these categories.

10. Attacker's actions

This category refers to the actions users believe attackers to carry out once an attack has been successful.

Example 1: "Probably read personal information, like emails. I would think normal person would simply read, wouldn't delete anything. Probably the most common thing is probably reading emails, because with files, if it's technical. Could be some infor, if you've got some important file regarding your financial information, I don't know, your personal affairs, that might be also be interesting. Probably anything that would appeal to a sense of curiosity to other people. I think that's the ave, what the average person would do."

Example 2: "I'm assuming they deleted my files, which would be the worst possible case..."

Example 3: "Unless they use my ID to put a virus or anything like that on, actually physically doing something..."

11. Attacker's methods

This category refers to the methods users believe **attackers** to employ in order to get access to a computing resource.

Example 1: "You can know somebody's password... sorry, you can know somebody's username and you can then sort of try out random passwords or from a list of commonly used passwords. Things like that person's name, if you can access that at all..."

Example 2: "...they might simulate a situation where it seems that they actually need to know your password, for whatever reason, and they can go to a great extent to set up a, particular situation where you can be ill or whatever, or make you ill, things like that..."

Example 3: "Also, having interfaces on the machines which would simulate a full genuine interface, but it's just basically for capturing passwords."

12. Attacker's motivations

This category refers to the motivations users believe **attackers** to have when they are attacking specific computing resources.

Example 1: "Could be a challenge for them, just to do it. Could have some reason for getting the information from somebody's system. For any reason, financial, personal, emotional. He must have some reason to do it. So it's probably either you're gonna benefit from it in some material way, or it's just the challenge of doing it, rather than just actually get some information from the system, you want to prove to yourself or somebody else that you can do it."

Example 2: "Well, either they've got a grudge against society, or they just do it, they, for a, for humour or a laugh or something like that, I presume. Ahm, I can think of, that's probably it, or boredom, you know, they want to get their own back..."

Example 3: "The obvious one is directly criminal, people trying into somebody's bank account to steal money, or whatever..."

13. Attacker's targets

This category refers to users' perceptions of who gets targeted by **attackers**.

Example 1: "If I was rich and famous..."

Example 2: "They're high-profile. Some of those are supposed to be very secure, like the Pentagon is supposed to be unbeatable..."

Example 3: "... basically, some thing, some point, some personal reason as well, like, they didn't like me, basically. [laughs]"

14. Authority personality

This category refers to that aspect of a user's **personality** that determines the way in which he responds to orders in general, and to orders that contradict his own judgement in particular. **Authority personality** is a specific instance of the abstract category **personality**.

Example 1: "Because that's the sort of person I am - I do as I'm told."

Example 2: "...and I immediately put it into that category of statements that's come down from on high. You know, 'you must do this', 'you must do that', and 'it's mandatory that you do this', of which there are a lot. And because of that you tend to form some sort of judgment of your own as to whether that's really worth doing or not."

Example 3: "...and I am just one of the old guard. I see this as, right. this is what we got to do, this is how I'll manage it. It's not a problem, it's a management thing. Self-management."

15. Available behaviours

This category refers to password-related behaviours that a user is aware of and able to perform.

Example 1: "If I, ahm, changed my password regularly and didn't write it down..."

Example 2: "I write them down..."

Example 3: "I would immediately upon return into work, I would change it."

16. Awareness effort

This category refers to the efforts that users can see their organization make with respect to getting them to be more aware of password-related issues.

Example 1: "I can't think of any lecture I've been given or any email I've been sent about security. You know, I'd half expect to be sent some email periodically by the department saying 'You should be aware of security, change your password from time to time, here are the guidelines for it.' I can't remember ever receiving that."

Example 2: "But if they're really, really concerned about it then maybe they should have like an introductory lecture type thing..."

Example 3: "Just the amount of activity, the number of warnings, the information you see on it constantly."

17. Change regime

This category refers to the **change regime** that many organizations employ, which forces users to change their passwords at regular intervals.

Example 1: "...there's at least one BT system, I think it's one that prompts you to change it quite frequently..."

Example 2: "...the whole thing about changing passwords monthly is ludicrous..."

Example 3: "Everything else has to be changed reasonably frequently for security reasons. You just get told 'Your password's expired, enter a new one.'"

18. Choice

This is the central category of the grounded theory model, and it is understood to mean "choice of one of several available behaviours in a specific situation by a specific user".

Example 1: "I think if you don't have a password which can be broken by brute force, for example, or a dictionary word or something, then that already makes it more difficult. If you have a particularly long password that makes it even more harder, so, so the choice of password does help as well."

Example 2: "I mean you can hardly say a password which is entirely random, very difficult to remember, because you're not allowed to write it down either, so, I mean, you've got a balance between setting a secure password, setting one that you actually remember on a day to day basis..."

Example 3: "I think using things like past addresses or places where I lived years ago or using book titles or authors or whatever is fairly secure because somebody couldn't guess unless they knew a lot of my history or books I read. Which is quite a hard thing to guess. Using [password from above] I know is stupid, sorry, it's insecure... ahmm... Whether it's sensible or not depends on how much you think about security. Like I said, I'm not that fussed if somebody does break into the system. Fine. I normally don't get paranoid about that. So, it's easy for me to remember [password above] [...] which is why I'm doing this. And also it's easy to type."

19. Clarity of regulations

Regulations can be more or less clear about the exact behaviours users are expected to perform and the kind of punishment they will incur if caught breaking the regulations. **Clarity of regulations** is a specific instance of the abstract category **regulation characteristics**.

Example 1: "I think, whatever they were gonna do, you should know about it at the beginning, there should be, you know, something written down that says, you know, 'if you do this, then this is gonna happen to you.'"

Example 2: "They don't mention anything, there's no strict guidelines, so, if there are no strict guidelines set down..."

Example 3: "...it's not clear what the individuals actually have to do to stay within the rules..."

20. Computing resource

This category refers to the actual resources a user gets access to by authenticating himself successfully to the password mechanism.

Example 1: "Iris is the finance system."

Example 2: "WEBsystem, which is being fazed out, but that's a purchase requisition system..."

Example 3: "...and then it automatically gives me access to my Unix filestore as well."

21. Consequences to organization

Users have perceptions of the possible negative consequences to their organization that a successful attack on one of their **computing resources** might bring about. This category is a specific instance of the abstract category **attack consequences**.

Example 1: "Because what I have is basically some design documents, which, in the scope of the whole of BT is probably insignificant. Ahm, I don't have anything which a patent is pending on at the moment, so there isn't any direct commercial benefit by finding things I'm doing. I mean, okay, some organisation may find out a few technical stuff, a few technical issues of designs, but that is not a major thing."

Example 2: "...that would, that would be a pain for BT, 'cause that is work that's done for the whole of BT, so..."

Example 3: "Well, the impact can, can go to the point of the regulator, which is OFTEL, being able to, to make BT pay a big fine..."

22. Consequences to others

Users have perceptions of the possible negative consequences to other individuals that a successful attack on one of their **computing resources** might bring about. This category is a specific instance of the abstract category **attack consequences**.

Example 1: "Well, first of all I've got access to files created by other people, so, ahm, in that respect..."

Example 2: "The other aspect is, I keep some confidential, ah, mainly people matters, information in my computer, and I don't wish for that information to arrive at someone's desk, ahm, which shouldn't happen, that information on his desk, basically. But that deals mainly with people matters, so, ahm, performance reviews."

Example 3: "...and it would give that person a relatively easy way to cause a lot of trouble for everyone else on the same network and in the same internet and email system by clogging it all up."

23. Consequences to user

Users have perceptions of the possible negative consequences to themselves that a successful attack on one of their **computing resources** might bring about. This category is a specific instance of the abstract category **attack consequences**.

Example 1: "Well, I guess if they do something, if they delete some files or some vital information, something I've produced which is of importance to me, then obviously that would be, ah, that would be important, that would hurt, that would damage something. Whether it's work reputation, or whatever."

Example 2: "If I go to your account and send an email, let's say to your supervisor, saying, you know, bad things, nobody would know that you didn't send the email."

Example 3: "You could also, say, a potential employer sends you an email saying 'reply to this if you can make it for a particular interview', they might reply abusively or whatever, and that could, and therefore defamation of character as well, so you won't be able to go to your employer and say 'no, that wasn't me who sent the email', they will say well prove it or they could just say 'well, we've got 500 applicants, you've been, we can get rid of you and carry on, really'. It's not big deal to them."

24. Consistency of regulations

This category refers to the extent to which password regulations are consistent across systems, as well as over time. This is a specific instance of the abstract category regulation characteristics.

Example 1: "...which is about the proliferation of different systems with different requirements on passwords, and different times to change them, that's the big issue about it."

Example 2: "...because I think there are more than three different sets of rules flowing about..."

Example 3: "Is it gonna be a consistence of the rules across all those different domains."

25. Cruising attack

A cruising attack targets just about any resource and is usually carried out by attackers with indiscriminate aims, such as young hackers trying to break into anything they can get access to. This category is a specific instance of the abstract category attack types.

Example 1: "...there must be millions of people out there targeting random people..."

Example 2: "I would have thought if anything, it would be more likely to be a hacker just spamming, you know, doing thousands of people, of which I was just one."

Example 3: "Yeah, a bit like the situation if someone just, I believe people just, they get a program that pings IP addresses, until they find one that is online and then they'll try and hack that machine, it's not, it's nothing personal, it's just..."

26. Detection effort

This category refers to a user's perception of the effort his organization makes to detect violations of the password regulations.

Example 1: "I don't think they're gonna be, I don't think they'd have the time or the inclination to follow it up."

Example 2: "...you're gonna have to put things in there that audit it, and actually detect if people aren't following it..."

Example 3: "... if they made more of an effort to catch students who break, who violate these rules..."

27. Detection likelihood

This category refers to a user's estimate of the likelihood of being caught when performing an available behaviour that violates password regulations.

Example 1: "I mean, it's incredibly unlikely that any of the staff of, obviously, certain members of staff would turn a blind eye to it and of those there are some that wouldn't and it would be very, very unlikely, in fact I think there is no chance..."

Example 2: "Q ...do you think it would be likely that somebody would find out realistically, would you expect to be found out?"

A Not really, no.

Q Do you guys agree?

B Yeah.

Q What do you think? [directed at C]

C Yeah, I agree."

Example 3: "...it would be very, very unlikely, in fact I think there is no chance..."

28. Evidence of punishment

This category refers to any kind of evidence that users have of people in their environment having been punished in the past for breaking regulations.

Example 1: "A I think the only way people are gonna really listen is through other people being punished [...]"

D Yes unless you see some action taken, then they'll realise these guys mean business"

Example 2: “Not necessarily, though. Because, I mean, you know everyone, you know everyone a bit, anyway, in this department, so if someone is not there because they’ve been kicked out, it’s gonna have an effect on you.”

Example 3: “If you knew somebody to whom the punishment had been applied.”

29. Frequency of use

This category refers to the frequency with which a user uses a specific password.

Example 1: “I’ve got probably 4 or 5 housekeeping passwords that we use day in day out, well, you know, reasonably often...”

Example 2: “I have ones written down that I infrequently use, things I use every day, I don’t need to write them down.”

Example 3: “But a lot of those I use very infrequently.”

30. Group attack

A **group attack** tries to get into any one of a number of resources which belong to a group. The attacker wants to get access to one of them, but does not care which one. This category is a specific instance of the abstract category **attack types**.

Example 1: “...only insofar as I would be a point of entry to a wider world of BT research. [...] Because I have, I am a point of entry into a wider system that could be of interest to them.”

Example 2: “The chances are they’re probably looking for an easy route into a system. It’s just to find an obvious password that somebody somewhere using that system has got, rather than mine particularly, they just wanna get in, I should imagine.”

Example 3: “I would say it would, but again, it would still be a small risk, because you’re still one person within BT, and I assume everybody would have that, so again, there would be no particular reason for me over somebody else.”

31. Group characteristics

This category refers to characteristics of a group that a user is part of which may or may not entice **attackers** to target that group.

Example 1: “Don’t know. If I was working, for example, if I was working, say, you know, in a department where we did animal rights experiments, sorry, animal experiments and we were the subject of malicious attack by animal rights investigators, and I think that would increase my chances of being exposed to the mass, you know.”

Example 2: “...the, the more higher-profile, the better, so it’s probably, ahm, if you can hack into, ahm, a very security, ahm, an institution which is outwardly very security-conscious, then it’s a greater achievement, presumably.”

Example 3: “I think UCL is a really well-established university, it’s also renowned throughout the world, it was one of the first links outside America to actually have Internet access and all of these points actually highlight it as a real high-risk place, that’s why, you know, we’re in trouble several times with the Janet, I think it is, because we’re constantly being, you know, people are constantly seeing us as a source of attack and they will look for any way in whatsoever...”

32. Image benefit

Image benefit is a specific instance of the abstract property **user benefit** and refers to the benefit a user incurs by carrying out a specific **available behaviour** that reaffirms or improves his **self-image** or **public image**.

Example 1: “I would do. On the whole, I hope that I’m a helpful and trustworthy person. A trusting person. And that if I could help somebody out then I would say: ‘Yes, okay’.”

Example 2: “I would think a very, ahm, extremely conscientious...”

Example 3: “I’d just think they were very diligent...”

33. Image cost

Image cost is a specific instance of the abstract property user cost and refers to the cost a user incurs by carrying out a specific available behaviour that contradicts his self-image or public image.

Example 1: “Yeah, I would think that he’s not social person or, yeah, social he doesn’t do very well, or perhaps he’s, yeah, he’s even afraid of his shadow, that’s what I would say.”

Example 2: “Careless [laughs] and probably not too professional.”

Example 3: “That’s really a question for psychologists. What sort of people keep their desks tidy. What sort of people comb their hair in the morning. Probably the same sort of people who would not give their passwords away. People who are very sort of... either people who are very paranoid about breaches of security or whatever or people who were told ‘Never give your password away.’ without understanding why that would be and therefore they would never do it because they were told not to do it. People therefore who are obedient.”

34. Importance of security

This category refers to a user’s perception of how important security is for his organization.

Example 1: “It’s not, so to me it doesn’t seem as if, you know, that UCL really is that keen on security. “

Example 2: “...it just shows how the department cares and how important it is to the department...”

Example 3: “So, it seems to be a quite low priority.”

35. Incident evidence

This category refers to any evidence users have of security breaches having occurred in the past.

Example 1: “I think so too, yeah, but I’m basing that on just the fact that I haven’t heard about any serious problem, well any problems at all, actually, so I can’t really make a, a sensible judgement by me sitting there thinking about it, I just don’t know enough about it, I’ve never heard about any...”

Example 2: “I think if there were problems, you’d hear about them.”

Example 3: “That made me, this immediately, the fact that I heard about this made me immediately think quite a lot more seriously about security of the stuff. Because it was something that was the real events.”

36. Inconsistency of change regime

This category refers to cases of different resources having change regimes with different change intervals in place, with the timing of the changes not being synchronized.

Example 1: “One thing that is interesting from the grouping is one where I slip up, is by having some passwords, some systems that require me to change my password regularly, and others where I don’t have to change at all or very infrequently.”

Example 2: “...to be honest they do get out of step. But I have a day when I try to change as many of them as I can.”

Example 3: “When every day of the week you get prompted to change your password on this system or change it on that system and you get fed up with.”

37. Inconsistency of proactive password checking criteria

This category refers to cases of different computing resources having proactive password checking mechanisms in place that differ with respect to the criteria they use to determine valid password content.

Example 1: “...so you have to make sure that you create a password that’s acceptable to all of them because sometimes you can change a password, it’s perfectly acceptable for one system and you try and use the same password for another system and it’s not acceptable.”

Example 2: "I would expect the thing to understand you know, UNIX, VMS, and all of the minimum criteria that the major systems would use..."

Example 3: "I could only do it in certain systems where they could be the same, because other systems wouldn't accept the same - you know, the length was different, something like that..."

38. Known behaviours

This category refers to password-related behaviours that a user is aware of but not able to perform.

Example 1: "I'd be looking for examples of how you could create passwords for a given set of parameters. And perhaps how you could use a group password, and how you can adapt that. Because that could help you to remember things."

Example 2: "Some people find they have to write down passwords to cope with the number that they have and I was wondering if you could tell me if it is the same for you and, if so, where you might write them down."

Example 3: "I think that if it gave you hints about how to create things that are memorable to you. So at least the sort of basic passwords or root or whatever you want to call it would be easy for you to remember. And yet sufficiently complex so that it wasn't just a dictionary word or something that's weak like that."

39. Number of passwords

This category simple refers to the number of passwords a specific user employs.

Example 1: "I have four different passwords within BT and about five or six in non-BT-related work."

Example 2: "I'm trying not to confess it, it's probably only one."

Example 3: "I would say... eight."

40. Opportunistic attack

An opportunistic attack is not pre-planned but simply exploits an easy opportunity that presents itself. This category is a specific instance of the abstract category **attack types**.

Example 1: "...my level is more of the opportunist, who's actually there, looking over the shoulder, just seeing, and being a bit nasty, going 'I see what I can do with it'."

Example 2: "He watch me, you know, while I was typing my password."

Example 3: "... within a company like BT there's a lot of people who work here so could actually, you know, somebody a couple of characters away could actually get into your account."

41. Password regulations

This category refers to the regulations that are in place within an organization.

Example 1: "Having been here a year, I couldn't tell you what the current regulations are."

Example 2: "I'm not sure what the regulations are."

Example 3: "I think it's sensible, well, it's reasonable to have regulations regarding things like security..."

42. People security

This category refers to a user's perception of the extent to which other people in the organization are seen to make an effort to be security-conscious.

Example 1: "I mean, in the main, people don't take security very seriously at times."

Example 2: "I think most people that I know are reasonably conscious of security. They certainly wouldn't divulge a password just like that, for no obvious reason. If someone rang them up on the phone and said, you know, 'we're from systems group, could you give me your password', you know, I don't know, maybe they would actually say yes. I'd like to say I don't know anyone actually, but that's not entirely true, I'm sure there are some people that would."

Example 3: “I think the people I worked with on the project were security-conscious, yes. The people I generally sit next to aren’t always necessarily security-conscious.”

43. Perceived security

This category refers to a user’s perception of how secure the computing resources he accesses are.

Example 1: “...but I still don't think they're secure...”

Example 2: “Do I think it’s secure? To a certain extent, I suppose...”

Example 3: “...but the feeling is that it’s probably secure.”

44. Personal attack

A **personal attack** targets one specific computing resource, the owner of which will be known to the **attacker**. This category is a specific instance of the abstract category **attack types**.

Example 1: “...most of the scenarios I can imagine when someone would be trying to guess my password would be someone that, ah, you know, someone that I know personally.”

Example 2: “...friends messing about, just playing a little joke on another friend...”

Example 3: “They might want to particular, pick on particular people they might have a grudge against, ahm, either humiliate them by changing things about them, or make things happen that appear to have been generated by that, that person.”

45. Personal characteristics

This category refers to those characteristics of a user which may or may not entice **attackers** to target him.

Example 1: “I mean, if I was much more publicly known, maybe either politically inclined in some way or, ah... maybe I’d ah, upset a certain person or group of people.”

Example 2: “That I had more responsibility, in terms, job-wise, that I had, ahm, just, well, that would mean, that I had more responsibility job-wise or that I had more responsibility in terms, commercially, so if I had, say, more money, or if I had a, if something would make it worthwhile for them, basically, some thing, some point, some personal reason as well, like, they didn’t like me, basically...”

Example 3: “I think they would target high-profile people like directors of the company, people that they’d know...”

46. Personality

There are certain aspects of a user’s **personality** that can have an effect on his choice of password-related behaviours. **Personality** is an abstract category that groups together the categories **risk personality**, **responsibility personality** and **authority personality**. Examples can be found in the subsections dealing with these categories.

47. Physical security

This category refers to any physical barriers that an organization has put into place to deter **attackers**.

Example 1: “The whole site is much more secure. You can’t just walk in off the street to get in there. There are security guards around and that sort of thing. So, getting physical access to a machine is a difficult point there.”

Example 2: “I mean, for example, the lab’s door is unlocked a lot of times, certainly last year, anyone could walk in.”

Example 3: “You can get through the security gates with no security passcard a lot of the time, so, you know, I think they need to have, you know, better security...”

48. Possible behaviours

This category refers to password-related behaviours that a user is not aware of, even though they exist.

Example 1: "I don't really know enough about passwords to answer that. There probably is, but the way in which hackers work, they would find another way around it."

Example 2: "But I think maybe new entrants to the company, new graduates who are going to be hit with a lot of passwords to cope with at one time, it might be good for them to give them a system, if they are not already familiar with password management."

Example 3: "...if I needed really to read my email or something then I would find another way to do it, not by giving somebody my password to access the system."

49. Primary cost

Primary cost is a specific instance of the abstract property **user cost** and refers to the actual cost a user incurs by carrying out a specific **available behaviour**.

Example 1: "I hate changing my password, it's the worst thing ever, it's so hard to think of a new one."

Example 2: "...I wouldn't like the sort of system where you have to memorise a strong password once every 6 months, it's time-consuming, it's cumbersome..."

Example 3: "Yeah, rehearsal, I do a bit of rehearsal, but, no, it doesn't take long for me to remember."

50. Proactive password checking mechanism

This category refers to the **proactive password checking mechanism** put into place by an organizations in order to make it impossible for users to choose passwords that are considered to be weak.

Example 1: "...there are plenty of systems now that prevent me from doing, clearly the system is requiring a stronger password than I would naturally choose, because it takes me a while to pick one when it's happened..."

Example 2: "Yes, because a lot of the systems, I think it's Gatekeeper, I can't remember which one it is, it requests, well, for me, also to pick up my telephone messages, I need to put in a password of so many characters and so many letters..."

Example 3: "And some of the high-security ones tend to be the ones that force you to choose something bizarre, they won't let you, you know, you try typing in something like Gandalf, and it goes 'Oh no, that's far too easy.'. It must have letters and bizarre characters in it, all this kind of stuff."

51. Provided security level

Provided security level is a property of **available behaviours** and refers to a user's estimate of the level of security provided by that behaviour.

Example 1: "Because, for instance, I have one that I tend to use variants of more for things online. And I will never use that for things that are more important or local to BT. Because I figure that the online ones are pretty insecure. So, who knows where that's going to end up or whether someone can hack it."

Example 2: "...but I've used other people's passwords to log on to systems, where officially you shouldn't be, you know what I mean, but I think it's generally accepted that you do and I think within the same context it's still considered reasonably safe to do."

Example 3: "I mean yeah. Partly online versus, which I think of as insecure versus sort of more local or BT systems, or as you say some banking systems. I definitely wouldn't share the same between say Lloyds bank logging on and, and something like Amazon for instance. So that's a kind of grouping. So I think it's actually subconsciously about level of security more than anything else..."

52. Public image

This category refers to the image a user generates in public.

Example 1: "Somebody that has something which he's not supposed to have, or just very secretive by nature without having any reason for it."

Example 2: “Not qualified to do the job.”

Example 3: “They’re more worried about not to be seen to be breaching any security rules.”

53. Punishment risk

Punishment risk is a property of available behaviours and refers to a user’s estimate of the risk of punishment he incurs by carrying out a specific behaviour that is not in accordance with regulations.

Example 1: “I think the, nobody expects not following password regulations might affect you.”

Example 2: “Nothing is gonna happen if you’re not following these regulations...”

Example 3: “...people are generally quite flippant about these things and you are always going to have the feeling that ‘oh it’s never gonna happen to me’. The chances are even if it did nothing much would come of it.”

54. Regulation characteristics

There are certain characteristics that regulations have that have an effect on a user’s estimate of the severity of expected punishment. Regulation characteristics is an abstract category that groups together the categories clarity of regulations and consistency of regulations. Examples can be found in the subsections dealing with these categories.

55. Required security level

Required security level is a property of computing resources and refers to a user’s estimate of the security required by that resource.

Example 1: “...I think the system doesn’t actually warrant any security...”

Example 2: “This comes back to the security being appropriate for the system.”

Example 3: “I think it’s a high-security system, the system needs high security...”

56. Resource abilities

Access to a computing resource gives a user (and consequently an attacker) certain abilities. Resource abilities is a specific instance of the abstract property resource characteristics.

Example 1: “My Unix machine account has my Web page on it as well, so somebody clearly break into that and change my Web page, which is again my sort of public image to the world...”

Example 2: “They’ll be looking on the ones, obviously, on the boxes of servers because they can do much more damage at that level, so, ahm, that, in some ways, it’s those parts that gotta be the ones that we’re really gotta be concerned about...”

Example 3: “...where I can authorize a purchase order...”

57. Resource characteristics

Every computing resource has a number of characteristics. Resource characteristics is a property of computing resource. It is an abstract property that groups together the properties resource abilities, resource content and resource visibility. Examples can be found in the subsections dealing with these properties.

58. Resource content

Access to a computing resource gives a user (and consequently an attacker) access to the content of that resource. Resource content is a specific instance of the abstract property resource characteristics.

Example 1: “It depends what sort of data you keep in your account as well, in your student account, if you, for example, if you keep a file, say, with your phone numbers or passwords for other accounts, and they get into that, then from there they can spread into whatever.”

Example 2: “It could be an issue, I do work for [name of an organization that deals with criminal acts], so there’s a possibility, if a BT employer was making nuisance calls he could possibly want to get into the systems and be able to see what his notes are on the system, so there’s small, a very small chance,

very, highly unlikely, but it is quite sensitive data, no customer is allowed to see that data, so it's possible."

Example 3: "Email. Ah... It could cause me a problem because I've got a lot of data coming in from, you know, a lot of emails coming from customers, and if, say for example you've got appointments with a customer and they're wiped out, it could be pretty embarrassing, that, that they have deleted all your emails and accounts, so, in that respect, yeah, definitely, you know..."

59. Resource visibility

In order to instigate an attack on a specific computing resource, an attacker has to know that the resource has a certain resource content and/or certain resource abilities.

Resource visibility is a specific instance of the abstract property resource characteristics.

Example 1: "So, but often people outside don't know what we're doing until we publish, and by then it's too late, so, you know, it's probably very hard for someone to know what I'm doing at the moment, outside the company."

Example 2: "There's always the idea that you go to a conference, and you're giving a paper on security, that therefore they might just want to find out what you're doing now. So, you always have to consider."

Example 3: "No, they have to know that it's there and if they know that it's there then yeah, they can access it."

60. Responsibility personality

This category refers to that aspect of a user's personality that determines the extent to which he takes responsibility for protecting the resources provided by his organization.

Responsibility personality is a specific instance of the abstract category personality.

Example 1: "Personally, to me, I, it would be more important that my stuff is secure, not anything to do with UCL. My first consideration would be, would be my own personal account as opposed to UCL's reputation and all that. [...] Yeah, but, I mean, whatever, it might be really selfish but I would not think "oh, UCL is gonna lose money if my", you know."

Example 2: "So you take on board and make it an ingrained part of the way you work. And realise that it is important to the company. And the company values you for doing it."

Example 3: "I think, more people think it's sort of 'Hang on a minute. It's got nothing to do with me". They don't really care what happens to the department. [...] They've got the money to fight this if they have to."

61. Reward attack

A reward attack targets any resource access to which offers the successful intruder some kind of benefit. This category is a specific instance of the abstract category attack types.

Example 1: "...get some coursework which you have done, which they have yet to do, for example."

Example 2: "I mean they wanna change their exam marks or something like that..."

Example 3: "More sophisticated intruders probably are looking for something specific that could be anything. Technical work, design work, whatever, financial information, which is of some use for them."

62. Risk personality

This category refers to that aspect of a user's personality that determines the way in which he deals with low and uncertain risks. Risk personality is a specific instance of the abstract category personality.

Example 1: "I'd probably wait until something happened..."

Example 2: "I don't know. I'd take the risk."

Example 3: "...so, no worries, really. I would take the risk."

63. Safeguards

This category refers to the measures both an organization and individual users undertake to reduce the possible negative consequences of a successful attack on one of their computing resources.

Example 1: “Anything that’s personal, things like salaries, appraisals of people and so forth, what I’ve done, I always keep those on floppy disc, anyway...”

Example 2: “But then the issue of, say, personal emails coming in, you could always have a Hotmail account, which is like what I have done, I don’t have any emails which would be from a potential employer or something coming into the computer science account, it comes to another private account, which I think is more secure.”

Example 3: “Yeah, again, we use very strict backup procedures so they couldn’t destroy anything of value...”

64. Secondary cost

Secondary cost is a specific instance of the abstract property **user cost** and refers to the actual cost a user incurs as a consequence of carrying out a specific **available behaviour**.

Example 1: “It was, I was away at a conference, and had a portable, I couldn’t get mail to work, so I phone in and got a colleague in the office to check my mail.”

Example 2: “But I gave him that because I have certain permissions to do things that he can’t with regard to [name of a supervisor]’s website maintenance and it was just a lot easier for me to give him that rather than for me to keep doing, you know, his work for him.”

Example 3: “I have done as well, where to access the Network, a visitor who hasn’t got a Network login, we’ve logged him in on another computer. [...] So that they can access the net or whatever, and then when they’ve finished, log them off. But that’s because it takes so long to get somebody a Network password, and if somebody is only here for half a day it’s really not worth it.”

65. Self-image

This category refers to the image a user holds of himself.

Example 1: “Well, I don’t know really how you hack, I think that’s for, that’s for nerdier, nerdier people than me.”

Example 2: “Maybe I’m not paranoid enough, who knows, but I’m not terribly paranoid at all...”

Example 3: “...you see I don’t have much imagination when it comes to what people would be capable of doing and why they do it. Probably because I’m not at all technically minded...”

66. Severity of expected punishment

This category refers to the extent to which users expect to be punished for violating password regulations.

Example 1: “Probably close the account for a while...”

Example 2: “...I don’t think there would be any punishment for most of it...”

Example 3: “So you just, you get a slap on the wrist and told not to do it again.”

67. Social benefit

Social benefit is a specific instance of the abstract property **user benefit** and refers to the benefit a user incurs by performing an **available behaviour** that has a positive effect on his relationship with other people.

Example 1: “Yeah, he knew my main departmental, and I’m not supposed to divulge, but it was for a very good reason. [...] Well, I do trust X implicitly. Yeah, I trust X a lot, actually, I think I knew his password at the time as well, so there was a bit of a mutual trust going on, well, there still is.”

Example 2: “In the sense that, you know, it was a friend, it was doing a friend a favour, yes, no problem, I trust them.”

Example 3: “I know everyone down there better than I know people here. There are people here in my office whom I don’t really know very really well. And therefore I’d have less problems with trust. But this really comes down to personal trust. You can ask the same questions about ‘would I lend somebody money’. If they’re a friend and they need the money, then yes, no problem. “

68. Social context

This category refers to the social context a user finds himself in within his organization.

Example 1: “Well, they’re just in the same office. We all sort of work, we all work on the same program...”

Example 2: “But I can’t speak for any, anybody further than my own immediate area, really, can I?”

Example 3: “...like you know that all your friends are writing down their passwords, for instance...”

69. Social cost

Social cost is a specific instance of the abstract property **user cost** and refers to the cost a user incurs by performing an **available behaviour** that has a negative effect on his relationship with other people.

Example 1: “I would think they were being careless, and it might affect me on trusting them on more serious things...”

Example 2: “I think he would have been very upset. I think that would have caused some kind of crisis in the relationship somehow, I think.”

Example 3: “If they were a friend who I trusted or a colleague who I worked with then I’d be disappointed and I’d try to persuade them that I really needed it and that I’d do the same for them or whatever.”

70. Social norms

Social norms is a property of the **social context** a user finds himself in and refers to the norms and standards established there.

Example 1: “If I was joining a group, I would tend to adopt their standards. If they all changed their passwords and that were the rules then, okay, I’d change my password. I would say to them, to them, as a, in a conversation I would say ‘Don’t you think this is all a bit silly?’, but, if they had good reasons for it, or if that was just the way they do things, then fine, I’d go along with that.”

Example 2: “There is always one or two in every area where we have people who have just, do not focus on things like that. And it then just boils down to other people just nagging them.”

Example 3: “I’d think either they’re very new to the department and they don’t know anything about culture and sort of practices of computer scientists.”

71. Software security

This category refers to the security provided by the software running on specific computing resources.

Example 1: “It seems, in itself Unix seems to be more secure than Windows, therefore I spend more time guaranteeing that security. I’d say you can’t make NT secure, so I don’t bother with a hard-to-guess password. Whereas Unix is fairly secure, I think.”

Example 2: “I don’t think so, because you can choose a strong password, yeah, but that password is sent over the network, let’s say it’s unencrypted, no matter how strong it is, you can always get that password off wire right away...”

Example 3: “Sorry, but how much damage can you do from a CS machine. I mean, there’s more restrictions on the machine here than there would be somewhere else.”

72. Strike policy

This category refers to any mechanism the organization has put into place that allows users only a certain number of unsuccessful login attempts before they are locked out and need to have the password reset.

Example 1: “I see why they do it, it just annoys you, because you’re there thinking ‘okay, it’s one of these four, it’s always just one more than you’ve got’.”

Example 2: “You’re three times wrong, and you’re locked out.”

Example 3: “...so three strike might actually be more practical than it is at the moment.”

73. Usability

This category refers to the general usability of the password mechanism and password procedures.

Example 1: "...until the technology is supporting people's natural password behaviours, and making those behaviours more secure..."

Example 2: "...sufficient time and effort should be put into devising systems to make it easy for people to follow the regulations you're trying to make them follow"

Example 3: "...or you've got to make the process painless. So that it's actually easier to follow the process, than not to follow the process."

74. User benefit

User benefit is a property of available behaviours and refers to the benefits a user incurs by carrying out a specific available behaviour. It is an abstract property and groups together the properties social benefit and image benefit. Examples can be found in the subsections dealing with these properties.

75. User cost

User cost is a property of available behaviours and refers to the cost a user incurs by carrying out a specific available behaviour. It is an abstract property and groups together the properties primary cost, secondary cost, social cost and image cost. Examples can be found in the subsections dealing with these properties.

76. User's abilities

In order to perform certain password-related behaviours, users need specific abilities. User's abilities is an abstract category that groups together the categories ability to satisfy proactive password checking criteria, ability to connect passwords, ability to create password content and ability to manage multiple passwords. Examples can be found in the subsections dealing with these categories.

77. Visible detection effort

This category refers to the effort that an organization visibly makes to detect the violation of password regulations.

Example 1: "Or occasionally we will have a situation where there are spot checks..."

Example 2: "Yes, particularly here where they have got so many cameras about..."

Example 3: "...and I think we get security audits come around, people who are responsible for security in different organizational parts, in different buildings."

APPENDIX B: STATISTICAL TESTS

This appendix lists the statistical tests that have been carried out to support the claims made in section 4.3, section 4.6 and chapter 5. There were 4 groups in studies 2 and 5, which have been described in sections 4.3 and 4.6. The data that has been analyzed are the responses to the questionnaire in Appendix F. The reports listed here use a short form to refer to individual questions on this questionnaire. The first page following the title page of the questionnaire is referred to as 'p1', and the first question from the top on this page as 'p1q1'. Accordingly, 'p3q2' refers to the second question on page 3. On page 4, 'p4q1r1' refers to the first row of the first question.

Page 1

Kruskal-Wallis Test

	p1q1	p1q2	p1q3	p1q4	p1q5	p1q6
Chi-Square	4.169	5.793	3.828	4.357	3.814	2.414
df	3	3	3	3	3	3
Asymp. Sig.	0.244	0.122	0.281	0.225	0.282	0.491

Wilcoxon Signed Ranks Test

	p1q2 – p1q1	p1q3 – p1q1	p1q4 – p1q1	p1q5 – p1q1	p1q6 – p1q1
Z	-8.315	-3.222	-0.834	-8.182	-6.591
Asymp. Sig. (2-tailed)	0.000	0.001	0.404	0.000	0.000

p1q3 - p1q2	p1q4 - p1q2	p1q5 - p1q2	p1q6 - p1q2	p1q4 - p1q3	p1q5 - p1q3	p1q6 - p1q3	p1q5 - p1q4	p1q6 - p1q4	p1q6 - p1q5
-6.195	-6.924	-0.209	-1.479	-4.431	-5.538	-4.391	-7.304	-6.284	-2.047
0.000	0.000	0.834	0.139	0.000	0.000	0.000	0.000	0.000	0.041

Page 2

Kruskal-Wallis Test

	p2q1	p2q2	p2q3
Chi-Square	3.795	2.474	0.225
df	3	3	3
Asymp. Sig.	0.284	0.480	0.973

Wilcoxon Signed Ranks Test

	p2q2 – p2q1	p2q3 – p2q1	p2q3 - p2q2
Z	-3.315	-3.724	-1.705
Asymp. Sig. (2-tailed)	0.001	0.000	0.088

Page 3

p3q1 chi-square test

	Value	df	Asymp. Sig. (2-sided)
Pearson Chi-Square	1.572	3	0.666

p3q2 chi-square test

	Value	df	Asymp. Sig. (2-sided)
Pearson Chi-Square	2.476	3	0.480

p3q3 chi-square test

	Value	df	Asymp. Sig. (2-sided)
Pearson Chi-Square	0.019	3	0.999

Page 4

Kruskal-Wallis Test

	p4q1r1	p4q1r2	p4q1r3	p4q1r4	p4q2r1	p4q2r2	p4q2r3	p4q2r4
Chi-Square	3.666	2.151	1.780	3.780	1.562	6.907	4.277	1.596
df	3	3	3	3	3	3	3	3
Asymp. Sig.	0.300	0.542	0.619	0.286	0.668	0.075	0.233	0.660

Wilcoxon Signed Ranks Test

	p4q1r2 – p4q1r1	p4q1r3 – p4q1r1	p4q1r4 – p4q1r1	p4q1r3 – p4q1r2	p4q1r4 – p4q1r2	p4q1r4 – p4q1r3
Z	-6.077	-8.001	-6.955	-5.929	-4.333	-0.120
Asymp. Sig. (2-tailed)	0.000	0.000	0.000	0.000	0.000	0.905

	p4q2r2 – p4q2r1	p4q2r3 – p4q2r1	p4q2r4 – p4q2r1	p4q2r3 – p4q2r2	p4q2r4 – p4q2r2	p4q2r4 – p4q2r3
Z	-5.402	-7.930	-3.504	-6.079	-0.047	-4.449
Asymp. Sig. (2-tailed)	0.000	0.000	0.000	0.000	0.963	0.000

Page 5

p5q1 chi-square test

	Value	df	Asymp. Sig. (2-sided)
Pearson Chi-Square	4.775	6	0.573

p5q2 (stem choice) chi-square test

	Value	df	Asymp. Sig. (2-sided)
Pearson Chi-Square	21.329	15	0.127

p5q2 (alteration choice) chi-square test

	Value	df	Asymp. Sig. (2-sided)
Pearson Chi-Square	17.210	12	0.142

p5q3 chi-square test

	Value	df	Asymp. Sig. (2-sided)
Pearson Chi-Square	2.508	3	0.474

p5q4 chi-square test

	Value	df	Asymp. Sig. (2-sided)
Pearson Chi-Square	4.227	3	0.238

p5q5 chi-square test

	Value	df	Asymp. Sig. (2-sided)
Pearson Chi-Square	1.179	3	0.758

p5q6 chi-square test

	Value	df	Asymp. Sig. (2-sided)
Pearson Chi-Square	4.806	3	0.187

p5q7 chi-square test

	Value	df	Asymp. Sig. (2-sided)
Pearson Chi-Square	3.473	3	0.324

P5q8 chi-square test

	Value	df	Asymp. Sig. (2-sided)
Pearson Chi-Square	3.802	3	0.284

APPENDIX C: FLYER 1

Protect your account to protect yourself

It is vital for the reputation of UCL's computer science department as a teaching institution that it be perceived as highly security-conscious in both the commercial world and the research community. Support from industrial collaborators and research councils will be under threat if security breaches occur and are publicized to the outside world.

Obviously, we expect you not to make attempts to gain access to other people's accounts and not to use your own account to carry out criminal activities. However, other people might try to gain access to *your* account in order to do just that. They could then hack into major computer systems, carry out denial-of-service attacks or even just tamper with your own work.

You will never achieve perfect security, but you can do your bit to make sure your account is as secure as possible. The following is a list of instructions we would expect you to follow. We will investigate any case of students not following these instructions.

To protect your account:

- Choose a strong, yet memorable password. Details on how to do this can be found at <http://www.cs.ucl.ac.uk/teaching/advice.html>.
- Learn your password by heart instead of writing it down.
- Keep your password to yourself – nobody else needs to and should know it.
- Use this password for your computer science account only.
- Lock your screen every time you leave a computer for a short break (5-10 minutes). Log out if you leave the computer for longer than that.

APPENDIX D: WEBSITE

How to choose a strong password

You need to choose a combination of letters, numbers and symbols that should be memorable only to you. We recommend you make the password at least 8 characters long. Here' s a couple of easy ways to create strong passwords.

- **Put three or more words together.** This strategy uses the vast number of combinations of dictionary words in combination with changing some of the letters for numbers or symbols. You can also add symbols in random positions.

Examples:

hit a boss becomes *Hit@6o\$\$*, *give us a lolly* becomes *giveusalolly!*,
dosh yellow chockies becomes *(doshyellowchockies)*

- **Use the first letters of each word** in a line from a favourite song or poem. Do not use a well-known line (such as the first line or the chorus) - use a different one or miss out the first word, or use a rare poem or song that few people know. Change some of the letters for numbers or symbols.

Examples:

The second line from "Yesterday", by the Beatles, Now it looks as though they're here to stay, becomes *nilatthts* or *NilattHtS* or *N1l@ttHt\$*

APPENDIX E: FLYER 2

Protect your account to protect yourself

It is vital for the reputation of UCL's computer science department as a teaching institution that it be perceived as highly security-conscious in both the commercial world and the research community. Support from industrial collaborators and research councils will be under threat if security breaches occur and are publicized to the outside world.

Obviously, we expect you not to make attempts to gain access to other people's accounts and not to use your own account to carry out criminal activities. However, other people might try to gain access to *your* account in order to do just that. They could then hack into major computer systems, carry out denial-of-service attacks or even just tamper with your own work.

You will never achieve perfect security, but you can do your bit to make sure your account is as secure as possible. The following is a list of instructions we would expect you to follow. We will investigate any case of students not following these instructions. We will clearly not be able to monitor every student's actions all the time (nor would we want to do that). **However, if any illegal activities are undertaken from your account and you claim they were performed by somebody else, we will investigate your past behaviour with respect to these instructions. It will count strongly against you if you have ignored any of them – whether that actually led to the security breach in question or not.**

To protect your account:

- Choose a strong, yet memorable password. Details on how to do this can be found at <http://www.cs.ucl.ac.uk/teaching/advice.html>.
- Learn your password by heart instead of writing it down.
- Keep your password to yourself – nobody else needs to and should know it.
- Use this password for your computer science account only.
- Lock your screen every time you leave a computer for a short break (5-10 minutes). Log out if you leave the computer for longer than that.

APPENDIX F: QUESTIONNAIRE

Questionnaire

on

Password Security

Your Degree Programme:

Year of Study: 1st 2nd 3rd

Please indicate how strongly you agree or disagree with the following statements.

I put myself at risk if I

- do not protect my computer science account from unauthorized use.

Strongly disagree

Strongly agree

1 2 3 4 5 6 7

- write down my computer science password in order to remember it, even if I destroy the note once I have memorized the password.

Strongly disagree

Strongly agree

1 2 3 4 5 6 7

- keep my computer science password written down permanently.

Strongly disagree

Strongly agree

1 2 3 4 5 6 7

- give my computer science password to someone I don't know.

Strongly disagree

Strongly agree

1 2 3 4 5 6 7

- give my computer science password to someone I know and trust.

Strongly disagree

Strongly agree

1 2 3 4 5 6 7

- use my computer science password for other systems.

Strongly disagree

Strongly agree

1 2 3 4 5 6 7

Please indicate how likely it is that the following incidences will occur.

Someone tries to gain access to your computer science account

➤ in order to damage your work.

Very unlikely						Very likely
1	2	3	4	5	6	7

➤ in order to perform actions that would incriminate you.

Very unlikely						Very likely
1	2	3	4	5	6	7

➤ to use it as a base for carrying out criminal activities.

Very unlikely						Very likely
1	2	3	4	5	6	7

Please indicate whether you believe the following statements to be true or false.

Someone determined enough will ultimately find a way to get into my computer science account, even if I choose a strong password and keep it perfectly secret.

TRUE

FALSE

The department cannot penalize me if someone does get into my account and carries out criminal activities, as long as I have followed the departmental regulations on proper password practice.

TRUE

FALSE

The department can only penalize me for not following departmental regulations on proper password practice if my behaviour results in someone getting into my account and doing harm to other people or the department itself.

TRUE

FALSE

A scenario

The computer science account of an undergraduate student has been used to hack into the system of a major company in the UK. Information in that system was tampered with, causing considerable financial harm to the company. The department has approached the student, who clearly states that these actions were not carried out by him/her. Further investigations by the department were inconclusive as to whether the student or someone else carried out the criminal activities. However, it became evident that the student had chosen a password that was easy to crack, and had also shared it with other people, thus violating departmental regulations. It could not be established whether the failure to keep his/her password secure directly led to someone else being able to break into the account.

The members of the department sit down together and discuss whether they should punish the student. The four lines of argument presented below dominate the discussion.

A. The student should be punished –

- s/he cannot prove that s/he did not carry out the attack on the company's system.

B. The student should be punished –

- the department cannot prove that the student carried out the attack on the company's system, *but*
- the student cannot prove that her/his failure to keep the password secure did not directly lead to someone being able to break into the account.

C. The student should be punished –

- the department cannot prove that the student's failure to keep the password secure did lead directly to someone being able to break into the account, *but*
- the student broke the regulations on choosing a strong password and keeping it secure.

D. The student should not be punished –

- S/he might have broken the regulations on choosing a strong password and keeping it secure, *but*
- so do a lot of other students, and s/he cannot be punished for being unlucky enough for someone to get into the account, which was the only reason the bad password practice was found out.

How fair on the student is each of the arguments?

	Very unfair					Very fair		
A	1	2	3	4	5	6	7	
B	1	2	3	4	5	6	7	

C	1	2	3	4	5	6	7
D	1	2	3	4	5	6	7

For each of the arguments, how likely is it that it will become the position finally taken by the department?

	Very unlikely				Very likely		
A	1	2	3	4	5	6	7
B	1	2	3	4	5	6	7
C	1	2	3	4	5	6	7
D	1	2	3	4	5	6	7

In order to improve the password mechanism and training, we need to find out about the current state of affairs. The following set of questions will aid us in doing this and ask you about the way in which you have chosen your password, whether you have written it down, and whether you have shared it with other people. Please answer them honestly, always remembering that this questionnaire is totally confidential. If you feel uncomfortable with this, please skip the whole section rather than providing incorrect answers.

Have you made a conscious effort to choose a computer science password that is difficult to crack? YES NO

Which of the following methods did you employ to choose your computer science password – tick several, if you have combined some of these approaches. If you have used a completely different method, please use the last row to specify it.

Tick	Method	Examples
	a word that means something to you	“Ibiza”, “holiday”, “Arsenal”
	a concatenation of words	“dogshatecats”
	a word constructed from the first letters of each word in a poem, song, etc.	“yamtsffa” from “yesterday, all my troubles seemed so far away”
	any of the above, with numbers added to it	“Ibiza01”, “4dogshate8cats”
	any of the above, with characters or numbers replaced by symbols	“Ar\$en@l”
	other:	

Did you initially write down your computer science password in order to remember it, but then destroyed the note, once you’d memorized it?

YES NO

Do you permanently keep a written copy of your computer science password somewhere?

YES NO

Is your computer science password written down somewhere constantly?

YES NO

Have you at any point shared your computer science password with someone else?

YES NO

Has at any point someone else shared their computer science password with you?

YES NO

Do you use your computer science password for other systems?

YES NO

APPENDIX G: FLYER 3

Warning – violating password regulations puts jobs at risk

BT is turning itself from a monolithic company into one that constantly branches out into new and exciting business areas. One consequence of this is that we are actively inviting outsiders into our office space. We therefore have to be more security-conscious than ever. In addition, our external ventures involve other companies using our computer systems. This business relies strongly on BT being seen as a security-conscious company. As we all know, there will never be *perfect* security, but we will lose business if we suffer security breaches due to a lack of professionalism. Ultimately, this could result in severe job losses.

To be seen as security-conscious we need a clear set of regulations, which all employees understand and follow. We can't and don't want to monitor your every action, since we believe that most people understand the importance of these regulations for the company's and their own long-term success. However, in the interest of all of us we have to and will come down on the few offenders with serious disciplinary measures. These might result in damage to their professional reputation, dismissal and even criminal charges.

Obvious violations of the regulations will be investigated. More importantly, security breaches will result in an investigation of the involved parties' behaviour. Violations in the past will face repercussions even if they did not directly lead to the security breach in question. Adhering to regulations ensures that no charges at all will be faced..

This website contains a tutorial presenting the password regulations and the reasoning behind them. It also contains the "Password Manager" which makes it easy for you to keep to these regulations. Set aside 15 minutes to read the tutorial, and bookmark it for later reference. Setting up the "Password Manager" may take up to 30 minutes. The tutorial will guide you through this process.

Remember: Violating the regulations set out in the tutorial puts jobs and your own career at risk - adhering to them means that you are safe from disciplinary measures whatever happens.